

Vol. 1

The AI Behavior Index

How detected AI use and enablement
correlates with employee risky behavior
across environments and roles

Table of contents

01	Executive summary	15	Controlling for role: AI use still changes employee behavior
03	More AI than tools can see	17	Prioritize your AI behavior controls
07	Asymmetric risk: AI enablement correlates with changed behavior	20	About this report
09	Risk by role: There's no "AI user"	21	Appendix

Executive summary

We analyzed 90 days of detected employee behavior across Fable customer environments (approximately 290K employees' aggregated data over 30 organizations) comparing all detectable generative AI-enabled employees - whether sanctioned or shadow - against non-AI enabled employees.

In effect, we created actuarial tables for what risky behaviors organizations can expect from AI-enabled employees: both overall, and when compared to their non-AI using colleagues in the same role category.

290k

employees' behavior
analyzed over 90 day
period across +30 orgs

Most of the workforce is already using AI; we just can't see most of it.

- Across all Fable client environments with detections capable of finding AI or granting AI licensure, **34% of employees had at least one allocated or detected standalone AI application**, model, plugin, or agent ("AI-enabled users") out of about 117K total possible employees.
- Nearly **55% of all employees use at least one role-specific app with newly enabled AI capabilities**. These secret AI tools appear identical to ordinary, non-enabled SaaS tools from a monitoring seat.

55%

of employees have at least
one role-specific app with
new AI capabilities

AI-enabled employees aren't universally risky on all measures.

On average and compared to the general employee population, AI users are:

- 36% less likely to be missing MFA,
- 43% less likely to trigger a DLP alert, and
- 21% less likely to be a risky phishing link clicker.

However, AI-enabled users are:

- 3.2x more likely to browse unsafely (out of 1,772 relevant employees in observable environments with EDR browser-based detections) and
- 3.7x more likely to be overdue on training.

The controls that got budget years ago are holding; however, **the risk migrated to a behavioral layer most organizations never instrumented or now require different detection technology**.

34%

of employees have at least one
standalone AI app or license

Not all AI users show the same risky behavior trends.

We then separated out AI users by their approximate role, and discovered variances of behavior within each role. Among other findings:

Administration & Governance (legal, HR, risk, compliance) is:

- 24x more likely to be a risk phish clicker than other AI users,
- Carries the worst MFA gap, and
- Has the highest AI file-upload rate in the AI user data set.

Technology is 3.3x more likely to see a credential-theft alert.

Finance is the positive outlier, with among the highest proxy AI adoption rates and some of the best compliance discipline in the dataset: Evidence that risky behavior erosion isn't the price of AI enablement.

It's not just "different people tend to use AI".

When we compared AI-enabled employees in a given role to overall role-based behavior trends in the same time period, **behavioral gaps still exist**:

- AI-enabled Administration & Governance staff are 5.3x more likely to leave a threat briefing incomplete than their same-role peers... and still 23.5x more likely to click on risky email phishing links.
- In Finance, our "best behaved" cohort when just looking within the AI-enabled group, AI-enabled staff are 2.2x more likely to browse the internet unsafely than their same-role peers, and yet about 73% less likely to not have MFA turned on than those same peers.

Ultimately, mature security teams must analyze how AI enablement AND role interact with culture, internal incentives, and other variables to determine likely impactful risky behaviors at their organization - and prioritize them for remediation.

How to use this report

For any risky behavior, consider the following two questions:

1 How many people exhibit this behavior?

2 How unevenly is it spread across roles?

Those two answers will sort every finding into one of four plays for your organization: campaign to everyone, roll out worst-cohort-first, precision-target the outliers, or simply monitor.

This report is a reference, not your unique story. Take our baselines, gather your employee data and your internal threat model, and use our most relevant findings to secure your unique human attack surface.

All risky behavior participation rates, relevant population counts, and calculation explanations can be found in this report's Appendix.


More AI than tools can see

Before we could determine how AI-enabled employees behave, we needed to establish how many people use AI applications at all.

Media hype implied every employee now weaponizes AI to vibecode daily responsibilities; surveys from CISOs and employees alike report that “everyone” is using it.

However, when we looked across all Fable customer environments, we detected standalone generative-AI enablement - a login, an allocation, or a session with a named AI app - for roughly 34% of employees (40,156 of 116,732).

Overall generative AI application use, logins, or allocations



We were initially surprised at these figures, given media hype, when realized that this analysis only represents the floor of current generative AI use trends.

Hidden AI:

What security telemetry struggles to see

Our initial 34% of employees using generative AI apps only counts what AI security teams can proactively label as “AI” within the SOC detections: the ChatGPT tab, the Claude login, the Cursor session.

It does not count the AI riding inside the tools your SOC files under “productivity”: the Copilot in your Microsoft 365 tenant, the assistant in Workday, the Cortex queries in Snowflake, the Einstein layer over Salesforce, or the meeting notetaker in your video conferencing platform.

Those features process the same sensitive data through the same models, but the traffic resolves to a trusted first-party domain, and nothing about it trips an AI-tool detection.

From a security monitoring perspective, an employee using Copilot in Excel looks identical to an employee using basic spreadsheets without AI.

When you count that embedded layer, the picture inverts.

- **Over 91% of employees** in our dataset were enrolled in at least one application with built-in AI capability.
- **Nearly 55% of employees are enrolled in role-specific technologies with new AI features**, even once we stripped out workspace giants such as Google Workspace or Microsoft 365.

The “No AI app detected” column in the previous chart isn’t necessarily a population that isn’t using AI; it’s most likely a population using AI in ways traditional security telemetry can’t detect yet.

To your security team, an employee using Copilot in Excel looks the same as an employee using basic spreadsheets without AI.

ROLES	EXISTING TECH WITH AI FEATURES	WHAT THE TECH DOES
 ADMIN & GOVERNANCE Legal & Compliance Human Resources Business Admin Risk Management	Workday, SuccessFactors	HRIS (system of record for employees, compensation, benefits)
	Greenhouse, Lever	Applicant tracking (resume intake & hiring pipeline)
	Docusign	E-signature & agreement management
	Ironclad	Contract lifecycle management
	ServiceNow	ITSM / HR service management & workflow
 TECHNOLOGY Engineering & IT Developers Product Management Quality Assurance Data Analytics Design & UX	GitHub, GitLab	Source code hosting & collaboration
	Jira, Confluence	Ticket tracking & internal wiki
	Datadog	Infrastructure & app monitoring
	Snowflake, Databricks	Cloud data warehouses (central data stores)

ROLES	EXISTING TECH WITH AI FEATURES	WHAT THE TECH DOES
 OPERATIONS Procurement & Sourcing Facilities & Real Estate Operations Warehous & Logistics	SAP, Oracle Fusion	ERP (procurement) & supply-chain backbone
	Coupa	Spend management (vendor master, invoices, approvals)
	Blue Yonder (supply-chain suites)	Supply-chain planning & logistics
	UKG, Kronos	Workforce management (shifts, time, payroll)
 REVENUE Marketing Sales & Business Development Partnerships & Alliances	Hubspot, Salesforce	CRM (customer records, pipeline, pricing system of record)
	Gong	Revenue intelligence (records & analyzes sales calls)
	ZoomInfo, Clay	Sales prospecting
	Outreach, Salesloft	Sales engagement (sequencing & outbound communication)
	Canva, Adobe	Design tools
 FINANCE Finance & Accounting Mortgage & Lending Underwriting Claims	Microsoft 365 (Excel)	De facto financial model hosting & unstructured data
	NetSuite, Oracle ERP	Financial reporting
	Ramp, Brex, AppZen	Corporate cards & expense audit
	nCino, Guidewire	Lending & insurance platforms
 CUSTOMER SERVICE Customer Success Consulting & Professional Services	Zendesk, Intercom, Service Cloud	Ticketing & support platforms
	Zoom, WebEx	Meetings with clients & customers
	Notion, Obsidian	Internal knowledge bases

ROLES	EXISTING TECH WITH AI FEATURES	WHAT THE TECH DOES
☐ GENERAL	Microsoft 365 Copilot, Gemini for Workspace	Office productivity suites
	Slack, Teams	Internal chat platforms
	Grammarly	Writing assistance
	Otter	Automatic meeting transcriber

These findings influence every statistic that follows in this report, in two ways.

- 1. First, it sizes the problem.** Every behavioral finding in this report concerns AI-enabled employees. These numbers prove that even organizations without an official AI tool rollout underway are at risk of AI deployments within previously established tech stacks.
 - a. This share will only continue to grow every quarter, as vendors ship AI into tools organizations have already paid for – whether or not teams remember to tell Security that they’ve turned the feature on.
- 2. Second, it makes every calculation that follows conservative.** Throughout this report, we compare AI-enabled employees to the general population. But our detections route the embedded-AI users into the general group, because we can't see their AI use.
 - a. If AI use shifts behavior the way the next sections show, that misclassification pulls the comparison baseline towards AI-like behaviors – good and bad.

Ultimately, the multipliers in this report are floors, not ceilings. Our blind spots make the findings more conservative, not less, which means the problems are likely more severe than current security alerts can process.

So really, the question this report answers isn't “Should we worry about the few employees we’ve allowed to use AI?”

It's “Now that most of our workforce is exposed to AI we can’t see, what are they doing differently?”

“Now that most of our workforce is exposed to AI we can’t see, what are they doing differently?”

Asymmetric risk: AI enablement correlates with changed behavior


If AI use made employees uniformly riskier, this would be an easy report to write and an easy problem to budget against and control for.

However, when we compared AI-enabled employees to the general population and a pattern emerges that's stranger (and more actionable) than "AI is dangerous", per the chart below.

AI users are measurably better on the controls security teams have spent a decade funding, and measurably worse on a behavioral layer most organizations have never instrumented.

General employee
vs AI-enabled
employee risky
behaviors

OVERALL USER RATE
TOTAL AI USER RATE



See Appendix for graph figures and population counts.

We'll start with the good news. For technical controls – the ones that run quietly once configured – AI-enabled employees tend to outperform the general population, as detected AI users:

- **Have fewer detected MFA gaps:** AI-enabled users were about 36% less likely to be missing MFA than the general population (4.93% vs 7.74%).
- **Prompt at least one new DLP alert:** AI-enabled users were about 43% less likely to trigger a DLP alert in the report period (12.34% vs 21.52%).
 - Later analysis of the AI-enabled role-based cohorts show high rates file uploads to present AI applications, however, indicating that data loss may be occurring through non-detected channels.
- **Click fewer email phishing lures:** AI-enabled users were about 21% less likely to be a problematic live phish clicker (0.19% vs 0.24%).
 - This specific delta may seem immaterial, but many security teams have built their entire security awareness programs around this one number. So, we'll report it, too – and suggest security teams stop planting their forest from this single tree.
- **Have fewer new data breach exposure:** AI-enabled users were also modestly less likely to have a new data breach exposure in the reporting period (1.56% vs. 1.81%)

Read in isolation, that's a reassuring story: the people adopting AI are our more security-aware employees, and the controls we've invested in are holding. However, they're only half the chart, and the other half points the opposite direction.

On the volitional layer – the behaviors that require a person to actively do something, repeatedly, that no background control catches – the same AI-detected employees are dramatically worse at:

- **Unsafe browsing:** AI-enabled users were about 3.2x more likely to be caught browsing unsafely in the reporting period, when compared to the general population (46.33% vs. 14.39%)
 - That's nearly half of AI-enabled employees, versus one in seven of the general population.
- **Compliance training overdue:** AI users were also 3.7x more likely to have ignored or otherwise not completed their required training (10.00% vs. 2.69%)

These aren't marginal differences; they're the largest gaps in the dataset, and they run in the direction few security teams are watching.

The controls that got budget years ago (MFA enforcement, DLP tooling, email phishing programs) are doing their jobs. And now, AI seems to have moved human risk to a layer of discretionary behavior that most security stacks don't measure at all.

Correlation versus causation: Aren't AI users "just more technical"?

At this point, a fair reader will ask the obvious question: **aren't AI adopters just different people than the normal population?** Aren't they more technical, or more senior, or more curious; isn't that base reason what we're really measuring?

No. That question rests on an underlying and unproven assumption of "what sort of person uses AI", while also being a different question altogether than what this report is trying to answer.

That question asks us to determine the background cause of the relevant differences we've observed so far in our dataset; we are reporting on the quantitative observed application deployments and subsequent correlated likelihood of employees' risky behaviors.

To that end, this section establishes only that AI-enabled employees tend to behave differently from the general population. This section does not prove that AI use is what drives the observed difference, or that the type of person who uses AI apps is any better (or worse) than other employees when it comes to security hygiene.

(While we later control for general role-based differences within the observed behaviors, we still will not say that "AI causes risky behavior" or "All employees who use AI are inherently risky.")

Instead, the important finding is that whoever your AI users are, the given participation rates are the broad behavioral profile you can expect out of that group of people – which is generally not the profile your current controls were tuned for.

Risk by role: There is no “AI user”

The previous section treated AI-enabled employees as a single population. However, “AI users” are not one kind of person doing one kind of work. They’re your legal team and your warehouse floor, your engineers and your sales desk – each group carries different access, workflows, and cultural pressures that together create unique risk profiles and behavior trends.


Average them together, and you get a number. If you pull them apart, we get a map.

When we break the AI-enabled employee cohort out by general role categorization and compare each “bucket” against the rest of their AI-enabled colleagues in other roles, the risks don't spread evenly.

Each role-based group carries different access, workflows, and cultural pressures that drive unique human risk profiles and behavior trends.

Risky behavior indicators for AI-enabled employees by role

- ALL AI EMPLOYEES
- ADMIN & GOVERNANCE
- CUSTOMER SERVICE
- FINANCE
- OPERATIONS
- REVENUE
- TECHNOLOGY



See Appendix for graph figures and population counts.



Admin & Governance

Legal & Compliance

Human Resources

Business Admin

Risk Management

24x

more likely to be risky email phish link clicker than other AI-enabled employees

Why worry?

Mergers & acquisitions (M&A) materials and communications

Third-party vendor & procurement conversations

3.2x

more likely to not have MFA enabled than other AI-enabled employees

Workforce information (comp, SSNs, benefits, health insurance)

Payroll banking, direct deposit details

HRIS onboarding

E-signature platform administration

79.6%

Highest proxy AI adoption rate via file uploads

Vendor contract & procurement data

Benchmark compensation data

Attorney-client privileged communication via documents

Whistleblower investigation files

56.96%

Second highest unsafe browsing rate among all AI-enabled employees

</>

Technology

Engineering & IT

Developers

Product Management

Quality Assurance

Data Analytics

Design & UX

3.3x

more likely to have new credential theft alert than other AI-enabled employees

Why worry?

CI/CD pipelines

Package-publish rights

API & pipeline service account tokens

Production databases

Raw data feeds and datalake admin

System admin controls and access

Backup and server admin

Local admin permissions

Offboarding & onboarding permissions

Production & test environments

3.4x

more likely to not complete required threat briefings than other AI-enabled employees

Bug trackers & vulnerability disclosures

Stack traces

Product roadmaps & unreleased specs

User research recordings & participant data

Browser extensions (coding plugins)

Source repositories



Operations

Procurement & Sourcing
Facilities & Real Estate
Operations
Warehour & Logistics

3.8x

more likely to have
overdue training than
other AI-enabled
employees

Why worry?

Vendor master files
Invoice approval workflows
WMS / TMS (shipment control)

2.4x

times more likely to have
incomplete threat
briefings all AI-enabled
employees

Shared terminals
Deskless workflows
Banking-detail changes



Revenue

Marketing
Sales & Business
Development
Partnerships & Alliances

1.4x

more likely to have
a new DLP alert than
other AI-enabled
employees

Why worry?

Call & meeting transcripts
Prospect & customer exports

57.6%

Highest rate of unsafe
browsing practices among
all AI-enabled employees

Full CRM (customer & prospect data,
pipeline, pricing)
Partner roadmaps under NDA
Marketing automation platform
Browser extensions (sales plugins)

58.3%

Second highest rates of
proxy AI adoption via
file uploads among all
AI-enabled employees



Customer Service

Customer Success
Consulting &
Professional Services

1.5x

more likely to have a new
DLP alert than other AI-
enabled employees

Why worry?

Shared queue logins
Seasonal staffing patterns & external
contracting
Long-lived support tool sessions
Ticket history with customer data
& secrets

1.9x

more likely to not have
MFA enabled on core apps
than other AI-enabled
employees

“Log in as user” impersonation capability
Credentials to other organizations’
environments
Trained to be helpful over paper policy
process



Finance

Finance & Accounting
Mortgage & Lending
Underwriting
Claims

2x

more likely to be risky
email phish link clicker
than other AI-enabled
employees

Why worry?

Budgets & enterprise resource planning
(ERP)
Wire initiation & approval
Banking portals
Pre-release financials
Customer financial files

The team writing your AI policy is its worst offender

The standout group is Administration & Governance, which includes all legal, compliance, risk, HR, and administrative employees. This cohort drafts the acceptable-use policy, runs the vendor reviews, and owns the words “appropriate use of AI” in your handbook.

And yet, this same cohort is also, by a wide margin, the riskiest AI cohort in the dataset,

- **24x more likely to be a problematic email phishing link clicker** than other AI-enabled employees (2.41% vs. 0.10%): the single largest concentration of any behavior in any cohort we measured.
- **Highest file-upload-to-AI rate of any role at 79.6%:** the heaviest-intensity AI adopters in the company, as evaluated by this proxy metric.
- **Worst MFA gap among AI users** at 11.97%, 3.2x the rest of the AI cohort (3.72%).

The people with privileged access to workforce PII, M&A material, attorney-client communications, and whistleblower files are simultaneously your heaviest AI adopters, your weakest on multi-factor authentication, and (by a factor of twenty-four) the group most likely to click the wrong email link. The cohort governing AI risk is the cohort most exposed to it.

The cohort governing AI risk is the cohort most exposed to it.

Technology: Strong on tooling, exposed on credentials

Technology (engineering, IT, product, QA, data, and design) looks reassuring on the surface and isn't always. This cohort tends to be at or better than the AI average on most passive controls; we assume they trust technical controls more than behavioral.

However, AI-enabled technology employees were:

- **3.3x more likely to have a new credential-theft alert** than other AI-enabled users (5.96% vs. 1.81%).
- **3.4x more likely to leave a required threat briefing incomplete** (4.98% vs. 1.48%).

The blast radius here is different in kind, not just degree. A stolen credential from this cohort isn't a user account; it's a cloud console, a CI/CD pipeline, a signing key, or a production database. And, the employees holding that access are the ones most likely to skip the briefing that would tell them how the current attacks work.

They're strong on the controls they own, but opt out of the awareness that protects what those controls can't cover.

Revenue, Operations, Customer Services: Real risk, less scrutiny

The middle of the pack carries risk that tends to not get executive investment because they don't carry executive titles:

- **Revenue posts the highest unsafe-browsing rate of any qualifying cohort** (57.6%) – unsurprising for people whose job is the open internet, but they sit on the full CRM, pipeline, and pricing. They get the scrutiny of neither Tech nor of the exec suite; perhaps they should.
- **Operations is 3.8x more likely to be overdue on training** (26.79% vs. 7.02%), which means the cohort approving daily invoices and changing vendor banking details, is falling behind on exactly the briefings that flags the latest invoice fraud formats.
- **Customer Services carries the highest new DLP-alert rate in the AI cohort** (17.47%), sitting on ticket histories full of customer-pasted secrets and “log in as user” impersonation powers.

Finance: the proof that erosion isn't destiny

And then there's Finance: the cohort that breaks the doom narrative, and the most important single bucket in this report for that reason.

Finance adopted AI at scale (52.3% file-upload rate, one of the highest in the AI enabled cohort) and held the line:

- **Lowest training-overdue rate in the entire dataset** by a substantial margin at 0.85%, compared to a non-Finance AI user average rate of 12.11%.
- **Generally intact behavior discipline across the board** - especially for briefing completions, new credential theft alerts, and compliance training completions - despite AI adoption rates that rival the riskiest cohorts.

Even their twice-as-likely to be a risky email phish link clicker than the average AI user is still relatively low, at a Finance AI-user average participation rate of just 0.36%.

Finance proves the thing the rest of this section could be misread to deny:

AI enablement does not have to mean behavioral erosion. Employees can lean in hard, and still stay disciplined.

So the next step for security teams everywhere is, what do Finance employees with AI have that other AI-enabled employees don't have?

(We hypothesize they possess a historic awareness of and investment in more human-based compensating controls, due to increased regulatory compliance requirements and insider-threat risk awareness; however, that is close to establishing a cause when we can only safely present correlative findings.)

What do Finance employees with AI have that other AI-enabled employees don't have?

Controlling for role: AI use still changes employee behavior

Here is the objection that everything so far has been waiting for. AI adopters, a skeptic will say, are simply different people: more technical, more senior, more online. So, a report comparing them to everyone else is really just measuring who they already were.

That logic made sense to us, too. So, we compared AI users to their own colleagues in the same role. By holding the role constant, the compared populations now do the same work, sit in the same systems, and carry the same access.

If the behavior still moves, something other than job function is correlated to the movement.... And it still moved.

When AI enablement makes things worse

Within a matched role, AI enablement can still predict behavioral shifts... just not in the same directions we might expect or otherwise predict.

The clearest case is Administration & Governance. Compare AI-enabled legal, compliance, risk, and HR staff to overall risky behavior rates of peers of their peers in the same functions, and the gaps are stark:

- **Risky email phish link clicker: 23.5x** (0.10% vs 2.41%)
- **Briefing incomplete: 5.3x**
 - AI-enabled A&G staff are more than five times as likely to leave a required threat briefing unfinished as their same-role counterparts (7.69% vs. 1.45%).
- **Unsafe browsing: 2.9x** (56.96% vs. 19.92%)

AI-enabled administrators and governance employees are in the same department, doing the same job, with the same data access as their non-AI enabled counterparts.

The same pattern shows up in Technology-based roles, as **AI-enabled technology employees are 3.8x more likely to be overdue on compliance training** (10.47% vs. 2.77%) and **3.7x more likely to browse the internet unsafely** (42.90% vs. 11.73%) than their non-AI peers in the same roles.

Additional study is required to see if these patterns sustain themselves across more environments, cultural pressures, and AI adoption.

5.3x

AI-enabled legal, compliance, risk, and HR staff are more likely to have incomplete threat briefings than peers in the same role

3.7x

AI-enabled engineers are more likely to browse the internet unsafely than peers in the same role

23.5x

AI-enabled legal, compliance, risk, and HR staff are more likely to be labeled a problematic email phish clicker than peers in the same role

Where AI isn't associated with worse behavior

Not every role shows a gap. In Operations, Revenue, and Customer Services, several volitional behaviors barely change when you hold the role constant.

On looking more closely at the data, we believe the general population in these roles already exhibits the behavior at high rates, making AI-based behavior changes harder to detect. Revenue already browses the open internet in risky ways; Operations already runs behind on training.

The within-role likelihoods of risky behaviors are comparatively smaller, precisely because the baseline risk was large to begin with.

It's a different mechanism with the same exposure – and a useful caution against assuming AI is the only thing worth instrumenting.

Mature security teams must analyze how AI enablement and role interacts with culture and internal incentives to determine likely impactful risky behaviors at their organization.

Finance: When AI enablement (sometimes) helped

And then there is Finance: the AI user cohort that did well compared to other AI enabled peers, yet now perform inconsistently against their own role.

When we examine Finance AI user behavior against other users in finance, accounting, underwriting, and claims roles, their risky behaviors deteriorate:

- **Unsafe browsing:** 2.2x more likely among AI-enabled Finance staff than the others in their role (45.12% vs. 20.07% general-role baseline).
- **Incomplete briefings:** 1.5x more likely among AI-enabled Finance staff than their role-based peers (1.11% vs 0.72%) – though still an excellent completion rate, compared to many roles across multiple environments.

... And improve:

- **MFA enabled:** AI-enabled Finance employees are 72.68% LESS likely to be missing MFA than the average Finance employee (4.13% vs 15.12%).
- **New DLP alert:** Again, AI-enabled Finance users are 72.57% LESS likely to have a new DLP alert in the 90-day period than the average Finance employee (12.21% vs 44.52%).

These users have similar roles, often similar permissions, and intense regulatory pressures – yet, AI enablement has correlated with both better AND worse behavioral outcomes.

Why would AI-enabled employees have MFA enabled more often or fewer DLP alerts than their peers in the same department? Why would they miss briefings more often, or browse unsafely (though perhaps their AI is browsing badly for them)?

But why is AI enablement associated with BETTER employee behavior outcomes in Finance, and not in other roles (such as Administration & Governance or Technology)?

That question requires further study and larger sets before we can hypothesize on cause, instead of reporting on correlation.

Ultimately, mature security teams must analyze how AI enablement and role interacts with culture, internal incentives, and other variables to determine likely impactful risky behaviors at their organization – and prioritize them for remediation.

Prioritize your AI behavior controls

A finding you can't act on is trivia. So here is the part that turns the preceding pages into a plan.

Every risky behavior in this report can be sorted by two questions, and only two.

1. How many of your AI-enabled employees exhibit the behavior?


- That is, the average participation rate of how many users on any given day use AI and signal risky behaviors.

2. How unevenly is it spread across roles?

- We calculated spread by dividing the highest role-based participation rate by the lowest

Plot those two axes against each other, and every behavior lands in one of four quadrants, each with a different behavior intervention strategy to prioritize work.

The chart below places the behaviors from this report; the thresholds (i.e., 10% participation, 3x spread) are deliberate judgment calls, not statistical cutoffs, chosen so that "one in ten employees" reads as a population-level problem rather than an individual one.



See Appendix for quadrant data.

Four AI behavior plays

■ UNIVERSAL CAMPAIGN

Blanket the (AI) org

High participation

Tight spread

What happens

When a behavior touches a large share of employees and every role does it at roughly the same rate, role-targeting just delays coverage. File uploads to AI services, unsafe browsing, and DLP alerts live here: nearly everyone, everywhere, regardless of department.

What you can do

The play is org-wide: clear acceptable-use guidance on what does and doesn't belong in a prompt, browser-level protections, and DLP tuned for AI-service channels rather than only the legacy exfiltration paths. You are not hunting outliers; you are raising a floor.

■ SEQUENCED ROLLOUT

Worse cohorts first

High participation

Wide spread

What happens

A behavior common enough to matter but concentrated unevenly enough that a blanket campaign wastes effort on the already-compliant. Compliance training overdue is the case in point: Operations runs at 27% overdue while Finance sits under 1%.

What you can do

Don't nag the whole company! Start with the worst cohort, fix it, and work down the ranking. The intervention is the same everywhere; the sequence is the strategy.

■ OUTLIER HUNTING

Precision-target

Low participation

Wide spread

What happens

A small affected population, but heavily concentrated in one or two cohorts. This is where your highest-leverage targeting lives. The 24x phish-clicker concentration in Administration & Governance, the credential-theft signal in Technology, the briefing gaps - these behavior trends don't call for a company-wide message that burns everyone's attention. They call for a surgical one aimed at the cohort that carries the risk, and silence everywhere else.

What you can do

Precision here is not just efficient; it's how you avoid the alert fatigue that makes the next campaign land flat.

■ MONITOR

Watch, don't message

Low participation

Tight spread

Nothing concentrated enough to target and nothing prevalent enough to campaign on. Data breach alerts sit here for this report. The play is to watch the baseline and spend your scarce attention elsewhere, while noticing if anything migrates into this quadrant over time, because that's what a control working looks like. This risky behaviors can "graduate" into this quadrant, as mature awareness programs watch for recidivism.

Use our chart to inform your internal threat model,

The quadrants tell you the shape of each behavior; your own threat intelligence tells you which behaviors matter this quarter. So, work from the intersection! For example:

- **If you're seeing voice-phishing against your support desk**, start where Customer Services ranks worst and ignore the rest of the map for now.
- **If leadership is worried about prompt injection attacks**, then the file-upload and unsafe-browsing cohorts are your universal-campaign starting line.

The chart is a way to sort options, not a substitute for knowing which fight you're in.

And if you can't measure spread yet, because you haven't instrumented these behaviors across roles, default to the universal plays while you build that visibility.

Because the data suggests that the risky behaviors in the high-participation quadrants are happening in your environment right now - whether or not you're looking for them.

The data suggests that risky behaviors are happening in your environment right now - whether or not you're looking for them.

About this report

About the data

1 Everything in this report comes from Fable clients' security detections - which vary by environment.

We didn't run a survey, we didn't ask employees what they thought they were using (or what they'd admit to), and we didn't ask CISOs what they feared is happening with AI.

So, when this report says a group of employees browsed somewhere unsafe or uploaded a file to an AI service, it's because a Fable customer's detection caught it, not because an employee volunteered it in a survey.

We analyzed 90 days of employee behavior across Fable customer environments: roughly 290,000 total employees from over 30 organizations across industries in Q2 2026, each instrumented to varying depths. Not every org has every tool, which causes population variances per analysis. (e.g. the 'total possible' AI user population is a ~117K subset of the 290K, representing the environments where AI detection was available.)

2 "AI-enabled" means we detected an AI app, agent, plugin, or model - including shadow AI.

An employee lands in our AI-enabled cohort if we detected them with any AI license allocated to their account, or if clients' telemetry detected any AI application on the employee's endpoint or account (e.g., standalone app, browser-based model, extension, or agent), regardless of whether their organization sanctioned it. We did that on purpose, because we wanted to study how AI users behaved, not necessarily who the organization sanctioned to use AI worked. After all, shadow AI users behave like AI users, because they are AI users.

To reference Michael Crichton for a moment: If you only count the dinosaurs you expected to see, you'll miss the ones breeding in the corners of the park.

3 These numbers are floors, not estimates.

Our detections caught standalone AI use, but could not see AI embedded inside sanctioned SaaS, like the Copilot in your spreadsheet, the assistant in your CRM, or the model summarizing your support tickets server-side.

If AI use shifts behavior the way this data suggests, that contamination pulls our non-AI baseline toward AI-like behavior, making the true relativities at or above our printed ones.

Our blind spots make these findings conservative, not inflated.

If you think we're wrong, please scan your own environment and check, then tell us what you find! We'd genuinely like to see your numbers, and would be happy to include you in the next benchmarking report.

4 This report looks at correlations, not causations.

Basically, we've created actuarial tables of AI behavior use: People with these attributes tend to do these things more often than other people without these attributes. We are not claiming AI use causes risky behavior.

In insurance, actuaries don't care why the cohort is riskier; they care that it measurably is, and they price accordingly. Treat every relativity in this report the same way: if you are enabling AI, this is the behavioral baseline walking through your door, whatever its causes.

5 Our counting rules

Every percentage is a participation rate: employees exhibiting a behavior, divided by the employees we could observe for that behavior. Populations vary by behavior because detections vary by customer, because not every organization has every detection; so each statistic carries its own n (total population).

With that in mind:

- Nothing is presented as a headline finding unless its base possible population is at least 100 employees;
- Flag anything between 25 and 100 as directional;
- Below 25 we don't publish, full stop;
- AI role relativities are computed against all other AI-enabled employees excluding the named role;
- Cross-role multipliers should not be ranked against each other beyond what that convention supports;
- AI-enabled employee comparisons versus general or role-based user populations are drawn from the same environment set, excluding any environments lacking required detections; and
- Some users double-counted as crossover roles between buckets, leading to sums greater than the total AI user population.

These rules unfortunately exclude conclusions about executive leadership populations who use AI, because so few cleared the $n > 100$ threshold. While we found some interesting directional statistics, we've withheld them for this report and hope to share them in the future. We'd rather show you less with confidence than more with doubt.

6 This is a starting point, not a verdict.

This dataset will grow as more customers connect more detections, and these relativities will be re-measured. Some will hold; some will turn out to be early-adoption artifacts that fade as AI use normalizes. That's the point of a baseline: it's the thing you measure drift against. Consider this report the first table, not the final word.

Appendix

All data referenced in this appendix relates to observable employee behavior across Fable customer environments, covering approximately 290,000 employees in over 30 organizations with varying levels of integrated or available telemetry, as observed from March 1, 2026 through May 31, 2026 (Q2 2026).

Throughout the report, to determine how more or less likely a population of interest was to exhibit a risky behavior signal than another population, we calculated relativity scores by dividing the

population of interest's average participation rate by the compared control population. Relativity scores above 1.0 meant the population of interest was more likely to exhibit the behavior signal than the control population; scores below 1.0 indicated that a population is less likely to exhibit the behavior than the control population, subtracting the relativity score from 100% to determine how much less likely a given behavior would be.

For extremely long AI-role vs AI non-role comparison rate tables and role-based relativities (pg. 11-16), email stryker@fablesecurity.com

APPENDIX 1.1

Overall general AI application use, logins, or allocations

METRIC	MEASUREMENT
Number of employees with any AI application detected Includes installation, use, login, or license allocation	40,156
Total relevant employee population For environments in which AI detection was possible	116,732
Percentage of total detected AI-enabled employees during the 90 day period	34.40%

APPENDIX 1.2

Hidden AI application access

METRIC	MEASUREMENT
Percentage of all employees enrolled or using an application with potential hidden AI capabilities Including workspace suites such as Copilot in Office 365 or Gemini in Google Workspace	91.30%
Percentage of all employees enrolled in or using a role-based application with potential hidden AI capabilities EXCLUDING workspace suites	54.80%

General employee vs AI-enabled employee risky behaviors

RISKY BEHAVIOR	GENERAL USER RATE	OVERALL AI USER RATE
Unsafe browsing detected by client EDR(s)	14.39% 13,609	46.33% 1,770
New data loss alert from client DLP(s)	21.52% 37,256	12.34% 6,359
Compliance training overdue per client LMS and/or Fable	2.69% 136,747	10.00% 2,500
MFA not enabled per client IdP(s), IAM(s), and/or ITDR(s)	7.74% 197,808	4.93% 1,998
Incomplete briefing per client LMS and/or Fable	2.90% 223,650	2.92% 27,243
New data breach alert from client dark web alert(s)	1.81% 278,586	1.56% 29,899
Risky email phish link clickers per client email security gateways (ESGs)	0.24% 54,390	0.19% 25,259

Risky behavior indicators for AI-enabled employees by general role

ROLE-BASED CATEGORY	Detected file uploads to AI	Unsafe browsing	New DLP alert	Compliance training overdue	MFA not enabled	New credential theft alert	Incomplete briefing	New data breach alert	Risky phish link clicker
Administration & Governance	79.63% 108	56.96% 158	14.94% 415	4.21% 1,165	11.97% 142	0.11% 887	7.69% 104	1.68% 119	2.41% 953
All AI-Employees	48.89% 1,929	46.33% 1,770	12.34% 6,359	10.00% 2,500	4.93% 1,998	3.52% 1,874	2.92% 27,243	1.56% 29,899	0.19% 25,259
Customer Service	40.53% 190	39.67% 121	17.47% 744	3.90% 231	7.50% 200	2.20% 182	0.38% 7,321	1.46% 7,676	0.10% 6,900
Finance	52.32% 453	45.12% 164	12.21% 696	0.85% 468	4.13% 460	0.67% 450	1.11% 1,707	1.82% 1,920	0.36% 1,677
Operations	45.41% 207	41.10% 73	6.41% 1,170	26.79% 377	4.72% 212	1.48% 203	6.36% 1,996	2.06% 2,335	0.40% 1,516
Revenue	58.33% 84	57.55% 106	17.44% 407	10.08% 119	-	1.35% 74	0.73% 3,690	1.16% 3,863	0.35% 3,409
Technology	45.57% 1,207	42.90% 725	12.69% 2,790	10.47% 1,222	3.78% 820	5.96% 772	4.98% 11,222	1.46% 12,552	0.15% 10,524

Risky behavior indicators for all employees by general role

ROLE-BASED USER CATEGORY	Unsafe browsing	New DLP alert	Compliance training overdue	MFA not enabled	New credential theft alert	Incomplete briefing	New data breach alert	Risky phish link clicker
Administration & Governance	19.92% 830	30.03% 1,709	2.93% 6,520	6.46% 12,518	2.88% 1,326	1.45% 15,032	3.48% 17,132	0.10% 2,011
Customer Service	5.59% 4,793	25.13% 8,019	1.27% 10,228	8.86% 10,228	2.13% 1,027	1.03% 1,770	2.71% 37,554	0.04% 17,136
Finance	20.07% 4,094	44.52% 7,683	0.89% 15,365	15.12% 24,119	1.08% 3,299	0.72% 29,227	2.10% 31,773	0.48% 8,465
Operations	12.68% 707	8.75% 6,435	2.92% 59,345	5.03% 65,438	1.45% 1,517	2.04% 72,067	1.44% 84,243	0.13% 5,813
Revenue	20.46% 288	9.61% 1,932	3.40% 7,696	7.24% 10,226	2.46% 603	1.56% 15,806	4.54% 18,110	0.15% 4,103
Technology	11.73% 1,959	10.34% 5,386	2.77% 27,108	6.13% 35,942	4.95% 1,476	2.96% 55,175	1.68% 54,939	0.23% 11,595

AI behavior prioritization quadrant

DETECTED BEHAVIOR	AI USER RATE	BEHAVIOR SPREAD MAX ÷ MIN
Uploaded file to AI app	48.89%	1.96x
Unsafe browsing detected	46.33%	1.45x
Compliance training overdue	10.00%	31.5x
New DLP alert	12.34%	2.7x
MFA not enabled	4.93%	3.2x
New credential theft alert	3.52%	54.2x
Incomplete briefing	2.92%	20.2x
New data breach alert	1.56%	1.8x
Risky email phish link clicker	0.19%	24.1x

Ready to see Fable
in action?

Request a demo to see how Fable surfaces
shadow AI and helps shape behavior.



FABLESECURITY.COM
HELLO@FABLESECURITY.COM

