

Aletheia

WHITE PAPER SERIES

White Paper No. 05/26
September - 2026

Principal Investigator

Mike Levy, CIA, CRMA, CISA, CISSP, CDPSE

Authors

Dr. Eric Redpath
Dr. Clark N. Banach
Kelly Crous
Ashley Hedges
Bryce McCarthy
Justin Morabito
Bella Peters

Completed by AUDIT Lab in Collaboration with



CANISIUS
UNIVERSITY

and



CherryHill
ADVISORY

The 30% Blind Spot

**AI ASSURANCE IN INTERNAL AUDIT IS RELATED
TO GOVERNANCE MATURITY, NOT BELIEF IN AI**

Table of Contents

Executive Summary	4
Introduction	5
PROJECT SETTING	
PROJECT DETAILS	
Methodology	6
RESEARCH QUESTIONS	
MOTIVATION	
ADOPTION DOMAINS	
TECHNOLOGY READINESS	
COMPOSITE MODEL	
PERSONAS	
Research Design	13
SAMPLE	
DOMAIN ANALYSIS	
BEHAVIORAL ANALYSIS	
BLIND SPOT ANALYSIS	
PERSONA ANALYSIS	
Domain Analysis	21
Behavioral Analysis	22
WILLINGNESS TO ACT	
TECHNOLOGY EXPOSURE	
READINESS DIMENSIONS	
SENIORITY	
Blind Spot Analysis	27
BLIND SPOT	
OVERSIGHT	
GOVERNANCE	
MATURITY	
Personas	34
ADOPTION SEGMENTS	
PERSONA APPLICATIONS	
HIDDEN CHAMPIONS	
PERSONA LIMITATIONS	
TAILORED INTERVENTIONS	
Recommendations	48
FIRM LEVEL FOCUS AREAS	
AUDITING AI	
INDIVIDUAL INTERVENTIONS	
INDUSTRY INTERVENTIONS	
ADOPTION ROADMAP	
Retrospective	52
WHAT WE LEARNED	
PROJECT LIMITATIONS	
PROJECT FUTURE	

Tables

Table 1: Disposition of Survey Responses (n = 3,673)	13
Table 2: Regional Distribution	14
Table 3: Persona Assignment Rules in Priority Order	20
Table 4: Relationships Between Individual Factors and Willingness to Present	22
Table 5: Mean Responses to Technology Readiness Items	23

Figures

Figure 1: Adopter Categorization on the Basis of Innovativeness	7
Figure 2: Technology Readiness Dimensions	8
Figure 3: Technology Readiness Means and Correlations	8
Figure 4: Composite Model of AI Adoption	10
Figure 5: Technology Readiness Profiles	11
Figure 6: Distribution of Adult Internet Adoption	12
Figure 7: Sample Size by Analysis	13
Figure 8: Seniority Distribution	14
Figure 9: AI Governance and Assurance Across the Three Lines	17
Figure 10: Perceived Maturity by Adoption Domain	21
Figure 11: Technology Readiness Responses and the Confidence Gap	24
Figure 12: Confidence Gap by Seniority	26
Figure 13: Adjusted Blind Spot Probability by Governance Maturity	27
Figure 14: Response Patterns Within the Blind Spot	28
Figure 15: Governance Maturity Distribution	30
Figure 16: Contradictory Coverage Answers by Governance Maturity	31
Figure 17: Factors Associated with Willingness to Present and the Blind Spot	32
Figure 18: Adoption Segments and Their Personas	34
Figure 19: Persona Landscape by Normalized Readiness	35
Figure 20: Persona Conditions and Proposed Interventions	37
Figure 21: Adoption Roadmap Across Time	51

Executive Summary

WHO

AUDIT Lab at Canisius University, in collaboration with Cherry Hill Advisory, leveraged the business psychology expertise of Eric Redpath, JD, PhD, and the internal audit expertise of Mike Levy, CIA, CRMA, CISA, CISSP, CDPSE, to assess 3,089 retained survey responses from internal audit professionals across 12 industries and 66 countries.

WHAT

The 'Ethics in the Age of Algorithms' survey was used to identify gaps between AI adoption and organizational adaptation. From these findings, adoption personas were developed using rule-based sorting to help leaders identify appropriate support for teams with varying constellations of experience. Recommendations were then developed from project outcomes.

WHEN

The survey was administered at an internal audit training event in December 2025. The research and reporting processes took place between January 2026 and September 2026.

HOW

Rogers' diffusion of innovations theory and Parasuraman's technology readiness framework are combined to explore how professional attitudes, technology exposure and governance relate to willingness to present AI-assisted work and AI audit coverage at represented firms.

WHY

This report has been developed to assess how internal auditors are using AI and to advance evidence based approaches to AI adoption. The project helps internal audit leadership address assurance gaps by identifying where teams need practical support, clearer accountability and governance maturity.

FINDINGS

Belief in AI's potential to improve audit quality averaged 3.99 on a scale from 1 to 5, compared with 3.28 for willingness to present AI-assisted findings. This demonstrates a gap between optimism and stated willingness to act. Only 15% of respondents reported advanced AI ethics governance maturity within their organizations (scores of 4 or 5). Among the respondents who answered the AI audit coverage question, 30% indicated either no audit coverage of AI within their organizations or uncertainty about whether such coverage existed. Of the factors tested, the strongest predictor of reported AI audit coverage was ethics governance maturity (see Figure 17). Belief in AI's potential and trust in AI results reviewed by humans, on the other hand, were consistent predictors of willingness to present AI-assisted findings and had no meaningful relationship with increases in reported audit coverage. Results indicate that audit coverage of AI may be constrained by governance maturity rather than by belief.

Introduction

PROJECT SETTING

The internal audit profession is not resisting AI. Internal auditors are adopting AI faster than their organizations are establishing governance. On average, internal audit professionals are likely to believe that AI can improve the quality of their work and they expect the upcoming pace of change to accelerate. At the same time, audit functions face the challenge of incorporating AI tools while maintaining reliable evidence, clear accountability and credible reporting to boards and audit committees.

AI adoption is not a one-size-fits-all process. A leader who conflates individual readiness with institutional readiness will experience higher rates of failure than one who adopts a multidimensional approach. A person may believe in the potential of AI but never use it, use it daily without appropriate governance or work in a well-governed organization without practical approved use cases. Comprehensive adoption is achieved by meeting people where they are at and guiding them into actions that are authorized, trusted and governed.

AUDIT Lab's first briefing in this series depicted AI adoption as a broad phenomenon across accounting and upstream disciplines (Redpath et al., 2026). The second briefing narrowed that lens to focus on the internal audit discipline, specifically, noting preliminary evidence of a governance gap between the rate of personal adoption and the rate of institutional adaptation (Banach et al., 2026). This report presents updates to the preliminary findings, as well as several tools internal audit leaders can use to develop effective interventions for their teams.

PROJECT DETAILS

This research explores whether professional confidence in AI processes corresponds to practical use by independently assessing belief, technology exposure, governance, oversight and willingness to act. We propose that these domains of adoption evolve independently and generate a range of readiness conditions. We then examine governance and oversight as distinct but interdependent paradigms, drawing on MIT's Mapping AI Risk Mitigations: Evidence Scan and Preliminary AI Risk Mitigation Taxonomy (Saeri et al., 2025) and IIA documentation (2020, 2023, 2024, 2026). This distinction separates establishing evaluation structures and protocols from conducting evaluations.

The findings are then used to develop ten rule-based intervention profiles to help leaders identify appropriate governance support, training and practical experience for teams with varying constellations of experience. These profiles provide a practical lens for organizational planning and a foundation for continued evaluation. We then provide recommendations to establish standards for approved use, human review, accountability, audit planning and ongoing monitoring, including an implementation roadmap.

Methodology

RESEARCH QUESTIONS

The question audit leaders keep asking is not whether to adopt AI, but rather how to cultivate teams that can adapt to it. To provide an evidence-based response, we set out to understand who is ready to use AI and what kind of support is required for pragmatic AI adoption. Before we address these broader questions, we explore more specific questions about domains of adoption, adoption behavior and organizational conditions.

MOTIVATION

This project responds to the need to understand how readiness is distributed and what forms of support may be required. Individuals with similar job titles or levels of technical access may face substantially different obstacles to responsible adoption. Some may have tools without supporting governance; others may have a strong control environment but limited practical exposure. Still others may report favorable beliefs while remaining hesitant to use or present AI-assisted work. These differences make it difficult to assess readiness through a single measure of adoption. Each constellation of readiness necessitates a different intervention, which is the rationale behind the personas developed for this report.

Although a number of approaches to AI adoption are being explored across industries, this report operationalizes the IIA literature on AI assurance to measure readiness against the roles it defines for governing, managing and assuring AI. The IIA separates governance, which directs the use of AI and whose scope includes ethical, social and legal responsibilities, from the independent assurance that internal audit provides (IIA, 2020, 2024). By establishing belief, exposure, governance, oversight and action as separate domains, audit leaders can identify gaps, set defined targets in each domain and develop adoption initiatives that address the specific needs of their teams.

ADOPTION DOMAINS

In this report AI adoption is modeled across five distinct domains. 1) Belief, the first domain, includes perceptions about AI and whether it can improve internal audit processes and work product; 2) Exposure, the second domain of adoption, is a record of use cases; 3) Governance, the third domain, is policies that address controls, accountability and review protocols for approved use cases; 4) Oversight, the fourth domain, is extending audit coverage to AI use cases; and 5) Action, the final domain of adoption, is a willingness to present, recommend and defend AI-assisted work. An individual or organization achieves comprehensive adoption by achieving stated objectives in each domain. Although conceptualized as a progression, maturity in the domains is often a non-linear process, meaning the order of adoption can be iterative, concurrent, asynchronous or non-sequential.


Methodology

TECHNOLOGY READINESS

The model used for this report builds on Rogers' (1962, 1995) 'Adopter Categorization on the Basis of Innovativeness' and Parasuraman's (2000) 'Technology Readiness Index.' Rogers' model categorizes adopters into five groups along a timeline of adoption (innovators, early adopters, early majority, late majority and laggards; see Figure 1). This sorting is based on how early someone is willing to adopt an innovation relative to other members of their group.

Parasuraman's (2000) framework separates readiness into drivers (i.e., optimism and innovativeness) and inhibitors (i.e., discomfort and insecurity). Parasuraman and Colby (2001) then use these dimensions to segment respondents by their overall disposition toward technology and to profile when each segment is likely to adopt (see Figures 2, 5 and 6). Our model integrates the adoption curve but has been developed to account for non-linear progression in consideration of asynchronous drivers of technology readiness in varying domains of adoption.

Figure 1: Adopter Categorization on the Basis of Innovativeness (Rogers, 1995)



Methodology

Figure 2: Technology Readiness Dimensions (adapted from Parasuraman, 2000)





Figure 3: Technology Readiness Means and Correlations (adapted from Parasuraman, 2000)



Methodology

Figure 3 summarizes the findings from Parasuraman's (2000) original study. Points show mean scores and horizontal spans show one standard deviation of his original sample. The connecting lines show correlations between dimensions. Innovativeness and optimism are positively associated ($r = 0.52$), as are discomfort and insecurity ($r = 0.56$). Correlations between drivers and inhibitors range from -0.40 to -0.28 . These relationships are consistent with enthusiasm and anxiety coexisting within the same person, which helps explain why favorable attitudes alone cannot provide a complete picture of readiness.

The four dimensions capture related but distinct dispositions, allowing a person to recognize technology's benefits while remaining uncomfortable with its use or uncertain about its performance. If readiness were a single continuum running from resistant to enthusiastic, the correlation between drivers and inhibitors would be closer to -1.00 . Instead, the dimensions are distinct constructs that vary semi-independently within the same person. Knowing that someone is optimistic about technology tells us remarkably little about whether that same person also trusts it (see Figure 3).

COMPOSITE MODEL


Rogers' adoption categories describe when people adopt an innovation relative to others; however, for organizational planning, adoption timing provides only part of the picture. Describing someone as an early adopter may capture their willingness to try AI while revealing little about the concerns that limit its use in other contexts. An organization may establish policies before developing practical use cases, while another may extend audit coverage to AI before formalizing its governance.

Our model therefore considers belief, exposure, governance, oversight and action as distinct domains of adoption, accounting for both personal attitudes and the organizational conditions in which people work. The model combines Rogers' focus on adoption timing with Parasuraman's recognition that drivers and inhibitors can coexist. Together, these frameworks allow us to represent adoption at different rates across domains (see Figure 4).

Rather than AI readiness being the outcome of passing time, it is assumed to be the outcome of passing maturity milestones in each domain. Belief, exposure and action describe individual attitudes and experience, whereas governance and oversight describe organizational conditions. In Figure 4, the vertical axis lists the adoption domains. The readiness dimensions across the top illustrate where each construct is expected to exert the most influence on adoption timing across the bottom, not a sequence a person is expected to pass through.

Methodology

Figure 4: Composite Model of AI Adoption



At an organizational level the rows represent institutional maturity rather than personal disposition. The distance between a pair offers information for planning. A short distance means domains of adoption are progressing together and require general support. A long distance means one domain has outrun another, and the intervention has to target the lagging domain specifically. The appropriate intervention depends on the domains involved and the conditions producing the gap.

A gap between governance and exposure, for example, directs attention to the relationship between organizational controls and practical use. Strong governance with limited exposure may call for opportunities to gain experience, while extensive use with limited governance may call for clearer controls and accountability. A person does not have to resolve one attitude toward technology before experiencing another. An auditor may believe that AI can improve audit quality while remaining unwilling to present AI-assisted findings. This coexistence complicates planning based on a single adopter category.

Methodology

PERSONAS

At the individual level, a person may be an innovator, with an extensive record of AI exposure, but a laggard in the action domain. When employing the framework in Figure 4, this combination would be represented by red dots in column 1 of the exposure row and column 5 of the action row. This illustrates that a person could be an innovator in one domain and a laggard in another. Reading the pair together establishes the building blocks of a persona. Belief at column 2 joined to exposure at column 4 describes someone whose belief in AI is established while their practical experience is not.

Parasuraman and Colby (2001) addressed a similar complexity. Noting that drivers and inhibitors can coexist, they segmented technology consumers by combinations of their attributes (see Figure 5). The five profiles range from explorers, who score high on both drivers and low on both inhibitors, to laggards, who show the reverse. Between them, pioneers are drawn to technology while remaining uncomfortable and insecure, and paranoids are optimistic but slow to experiment.

Figure 5: Technology Readiness Profiles (adapted from Parasuraman and Colby, 2001)

Profile	Contributors to readiness		Inhibitors of readiness	
	Innovativeness	Optimism	Discomfort	Insecurity
	Inclination to pioneer	Belief in benefits	Feeling overwhelmed	Distrust of technology
Explorers Enthusiastic about technology with relatively few reservations	High	High	Low	Low
Pioneers Enthusiasm and strong concerns coexist in the same person	High	High	High	High
Skeptics Limited enthusiasm alongside relatively few reservations	Low	Low	Low	Low
Paranoids See benefits but feel uneasy and are less inclined to pioneer	Low	High	High	High
Laggards Limited enthusiasm alongside strong reservations	Low	Low	High	High


The same person can be drawn to technology and held back by it.

Profiles describe combinations of beliefs. Their prevalence depends on the population studied.

High and Low describe each factor relative to the general population in the original analysis.

Methodology

Figure 6: Distribution of Adult Internet Adoption (adapted from Parasuraman and Colby, 2001)



The area under each band equals that profile's share of the adult population.
Published shares sum to 98% through rounding and are normalized to 100% here.
Adapted from Parasuraman & Colby (2001), Techno-Ready Marketing, Exhibit 4-1.

Parasuraman and Colby (2001) also connected these profiles to adoption timing. Ordering the profiles by when each entered the market redraws Rogers' adoption curve in terms of attitudes and timing (see Figure 6). The first segment takes up the same 16% that Rogers had separated into innovators and early adopters. The middle 68%, which Rogers divides evenly at 34% on either side of the mean, is split into three groups. Although a familiar shape, the curve is flatter, indicating that adoption was spread more evenly over time rather than concentrated around the mean.

Because drivers and inhibitors coexist, people do not move along the adoption curve at the same pace or for the same reasons, and what helps one group adopt may do little for another. The personas apply the same reasoning to AI in internal audit. In this study, beliefs and willingness are measured by five sentiment items, and conditions by AI and GRC technology use, governance maturity and AI audit coverage. Because the data are cross-sectional, maturity in each domain indicates how far along the curve respondents were when surveyed and is used to develop segments and personas.

Research Design


SAMPLE

The raw data contained 3,673 responses. Observations were verified using a series of data-quality procedures, leaving 3,089 records (see Table 1). A small portion of respondents (34) gave identical ratings to sentiment and audit coverage prompts, a pattern suggesting they clicked through without engaging with the content. Robustness checks do not change the governance findings when these observations are omitted, although unease about AI technology adoption becomes more significant. Sample sizes vary by analysis (see Figure 7).

Table 1: Disposition of Survey Responses (n = 3,673)

Status	n	Reasoning
Excluded	497	Flagged as ghosts (non-substantive responses)
Excluded	82	Flagged as duplicates (repeat session)
Excluded	5	Flagged not-session (incomplete or test session)
Retained	34	Flagged as serial flat (no variation in responses)
Retained	3,055	Cleaned and verified records


Figure 7: Sample Size by Analysis



Research Design

The sample was gathered by Cherry Hill Advisory in December 2025 via the 'Ethics in the Age of Algorithms' survey, administered to internal audit professionals at a training event. Respondents represent six seniority levels, from staff auditors to the C-suite, with the largest identifiable geographies being the United States, Canada and the Caribbean (see Table 2 and Figure 8). Regional analysis was omitted from the final report given the uneven distribution of respondents across regions and a large share of records with undisclosed locations.

Figure 8: Seniority Distribution



Note: Figures sum to 101% as a result of rounding.

Table 2: Regional Distribution

Region	n	% of total sample
United States	1,760	57.0%
Did not disclose	918	29.7%
Canada	170	5.5%
Caribbean	108	3.5%
Africa	41	1.3%
Europe	41	1.3%
Latin America	23	0.7%
Middle East	16	0.5%
Asia-Pacific	12	0.4%

Research Design

DOMAIN ANALYSIS

Comprehensive adoption includes the belief that AI can improve the quality of work, exposure to AI use cases, integrated or leading AI ethics governance, awareness of AI audit coverage and a willingness to confidently present AI work product to a board. We evaluate these factors in a topline analysis of domain maturity to offer provisional insights into how perceived maturity varies across adoption domains. Although each item was purposefully selected, a single measure cannot capture the full scope of a domain; however, the items align with the definitions established in this report and provide a reasonable basis for comparison.

Representing the belief domain are responses to the statement "AI can improve the quality of IA work," on a scale from 1 to 5, where 5 indicates total agreement and 1 indicates no agreement. The statement reflects the domain's definition and corresponds to optimism in Parasuraman's (2000) framework. Responses of 4 or 5 are considered representative of domain maturity. Exposure is estimated from self-reported technology use by assessing whether respondents have experience with at least one AI-adjacent tool. The measure captures the maturity of what respondents do rather than what they believe.

Governance maturity is estimated using responses to the question: How would you rate your organization's AI ethics governance? Responses of 4 or 5 are considered representative of governance maturity.¹ Although the question is framed around ethics, its levels describe the policies, controls, accountability and review that define the domain. Our assumption is that if there is continuous monitoring of ethics responsibilities, then related dimensions of governance, including legal and regulatory compliance, are also likely to be supported by established policies and controls.

Oversight maturity is represented by responses to a question about known components of AI audits performed at respondent organizations. If a respondent rated "none/I don't know" as a 4 or 5, it is assumed that either there is no AI audit coverage at their organization or the respondent is unaware of it. This is the blind spot referenced throughout the report. Responses of 1, 2 or 3 are classified as outside the blind spot for this analysis; they do not independently verify audit coverage. Perceived awareness of AI audit coverage is identified as a better general indicator of oversight maturity than reporting on the individual components an organization chooses to audit.

¹ Level 5 - Leading (continuous monitoring, ethics by design); Level 4 - Integrated (formal reviews, clear accountability); Level 3 - Managed (controls exist, ethics not fully integrated); Level 2 - Basic (policy exists but not implemented); Level 1 - Ad Hoc (no defined oversight).

Research Design

Action was estimated from responses to “I would present AI-assisted findings to the board/audit committee,” on a scale from 1 to 5, where 5 indicates total agreement and 1 indicates no agreement. Presenting to the board or audit committee is the most demanding form of the action, because the auditor must be prepared to defend the work. Responses of 4 or 5 are considered representative of domain maturity.

BEHAVIORAL ANALYSIS

The behavioral analysis estimates the relationship between action and individual conditions by treating the 'willingness to present' item as an ordinal outcome. The estimation employs 2,291 complete cases and robust standard errors. Seniority is used as a control and as a means to assess variation between roles, with staff auditor serving as the reference category for comparison. Governance maturity and awareness of AI audit coverage are included as controls, but they are not variables of interest or statistically significant (see Figure 17).

The analysis focuses on the willingness to present and its relationship with four beliefs aligned with Parasuraman's (2000) readiness dimensions. These are perceived AI risk (insecurity), unease about advanced technology (discomfort), belief that AI can improve audit quality (optimism) and trust in AI results under human review (innovativeness). Respondents rated their agreement with the statements "AI poses unacceptable risks to IA," "I feel uneasy about relying on advanced technology in audit work," "AI can improve the quality of IA work" and "I trust AI results when reviewed by humans." The relationship between exposure and willingness to present is also estimated for AI tools and for governance, risk and compliance (GRC) tools, each reflecting hands-on experience with at least one tool of that type.

We report odds ratios for each factor. An odds ratio of 2 means that a one-point increase in a belief doubles the odds of reporting a higher level of willingness to present. An odds ratio of 1 means no association, and a value below 1 means lower odds. Estimates are also included for the change in probability of reporting a 4 or 5 for willingness to present AI-assisted work product. All estimates are conditional associations that hold the other factors constant, and the cross-sectional design does not support causal conclusions. The outcome measures stated willingness rather than documented behavior. The survey did not ask whether respondents had previously presented AI-assisted findings, how often they had done so or how those findings were received.

The analysis also calculates each respondent's confidence gap by subtracting willingness to present from belief that AI can improve audit quality, so a positive gap means belief exceeds stated willingness. Gaps are compared across seniority levels (see Figure 12). A separate analysis estimates the share of the gap's variation that seniority explains.

Research Design


BLIND SPOT ANALYSIS

Where the behavioral analysis examines individual conditions, the blind spot analysis assesses organizational conditions. It estimates the relationship between oversight and governance maturity by treating the blind spot as a binary outcome. Respondents who rated "none/I don't know" 4 or 5 when prompted about known components of their organization's AI audits form the blind spot, and those who rated it 1, 2 or 3 form the comparison group. The final estimation employs 2,016 complete cases and robust standard errors. Governance maturity is measured by self-reported AI ethics governance levels, with Level 1 (ad hoc) serving as the reference category for comparison in the governance maturity analysis (see Figure 13).

Individual beliefs, exposure and seniority are included as controls; however, they are not variables of interest and appear only in Figure 17. Whether internal audit has extended coverage to AI is expected to depend on the audit plan and governance maturity rather than on how a respondent feels about AI or which tools they use. A respondent can be optimistic about AI and use it daily in an organization that has not brought AI into its audit coverage. Individual conditions are therefore not expected to explain the blind spot, although they may shape how respondents perceive their organization's coverage.

We are interested in whether AI ethics governance maturity is associated with a higher probability of audit coverage. Figure 9 maps the analysis onto a framework developed from the IIA's Three Lines Model (2020, 2026) and AI Auditing Framework (2024). In this structure, the governing body is accountable to stakeholders and helps ensure that AI use is consistent with ethical, social, legal and strategic responsibilities. The highlighted cells mark the variables of interest and depict ethics governance upstream of internal audit assurance on AI use cases.

Figure 9: AI Governance and Assurance Across the Three Lines (based on IIA, 2020, 2024)



Research Design

Results are reported as the adjusted probability of falling into the blind spot at each governance level, so a lower probability indicates that respondents at that level are more likely to report at least some AI audit coverage. The survey did not measure how management plans, executes and controls AI use cases, so the analysis can show whether governance and coverage are associated but not the mechanics of how one leads to the other.

Governance maturity was measured with a single self-reported item, so the analysis cannot show whether specific governance practices produce stronger audit coverage. Respondents assigned their own organization to a level, and the measure was not validated against documented policies, control inventories or audit workpapers. Two organizations with identical control environments may be recorded at different levels, so the measure reflects perceived maturity rather than an independently verified control assessment.

PERSONA ANALYSIS

The same auditor can be optimistic about AI's potential, eager to experiment with it, uncomfortable when unable to explain its outputs and uncertain whether those outputs are reliable. These dispositions can coexist and become more or less relevant depending on the use case. Someone may confidently use AI to organize information while hesitating to rely on it for an audit judgment or to present AI-assisted findings.

The persona framework operationalizes the adoption model using rule-based intervention profiles that account for these mixed dispositions. Respondents are characterized by their use of AI and GRC technology and differentiated by willingness to present and the four belief items that align with Parasuraman's (2000) readiness dimensions. Where further differentiation is required, governance maturity and expectations about future technology change are employed to refine assignment. A Readiness Score is then reported for each persona as a multi-measure composite of belief and action maturity. The approach was developed to avoid treating AI exposure as the default indicator of adoption maturity.

Persona shares and Readiness Scores differ from preliminary values reported in the June briefing (Banach et al., 2026). Respondents were reassigned using a more rigorous rule-based approach and stricter sample criteria, and Readiness Scores are now reported on a normalized scale from 1 to 10. Most persona shares differ only slightly. Respondents with AI or GRC exposure can no longer be assigned to Observer, which is now smaller. Explorer and Bureaucrat are larger because they include all respondents with AI exposure only or GRC exposure only who do not meet the conditions of a more specific profile.

Research Design

Assignment is limited to the 2,847 respondents who answered all five sentiment items. Each respondent is first segmented by one of four exposure conditions based on reported AI and GRC technology use (1. both, 2. AI only, 3. GRC only or 4. neither). The exposure condition determines which profiles a respondent can receive, and the remaining conditions determine which of those applies.

Governance maturity has two roles in the analysis. It enters the definitions of Complacent and Overwhelmed, and it is analyzed separately as an organizational domain in the blind spot analysis. Audit coverage plays no part in assignment. It is compared across personas only after assignment, as an organizational outcome.

Profiles are tested in a fixed priority order, and each respondent receives the first profile whose conditions they meet (see Table 3). Anxious is tested first and is the only profile without an exposure condition. It captures respondents who score 4 or 5 on belief that AI can improve audit quality, but also 4 or 5 on uneasiness and perceived risk.

Respondents with AI exposure are then tested for Vanguard (high belief, trust in AI results under human review and willingness to present), and the remainder are assigned to Explorer if they report no GRC technology or to Equipped if they report both. Respondents with GRC exposure only are tested for Foundation (high belief and trust), and the remainder are assigned to Bureaucrat.

Respondents with neither type of exposure are tested for Believer (high trust with willingness to present of 3 or lower), then Complacent (governance Level 1 with perceived risk of 1 or 2) and then Overwhelmed (high expected change over the next three years with governance at Level 1 or 2 or not reported). Those who meet none of these conditions are assigned to Observer. Where a respondent meets the conditions of more than one persona, the priority order determines the primary persona.

The Readiness Index provides Readiness Scores for each respondent as a composite of the five sentiment items. The raw score adds belief that AI can improve audit quality, trust in AI results under human review and willingness to present, then subtracts perceived risk and unease (Improve + Trust + Present – Risk – Unease), giving a theoretical range of –7 to 13. Raw scores are scaled to range from 1 to 10, which does not change the rank order of respondents or personas. Higher scores indicate stronger readiness drivers and willingness to act relative to readiness inhibitors.

Research Design

To test whether the personas represent distinct groups or positions along a continuum, separation was assessed using standardized measures of AI and GRC exposure, the five Readiness Score items and governance maturity. Silhouette coefficients compare how close each respondent is to members of their own persona relative to members of the nearest other persona and range from -1 to 1, with values near zero indicating overlapping groups. A centroid is a persona's average score on each of these measures, and the distance between two centroids shows how alike two personas are overall. The ten centroids were clustered using Ward's method to group similar personas, resulting in four higher-order segments (Tooled Adopters, Governance-First, Untooled Middle and Risk-Conflicted). Because the segments group respondents with similar profiles, they can be read as a guide to general adoption conditions, with nested personas offering specific interventions.

Table 3: Persona Assignment Rules in Priority Order

Priority	Persona	Exposure	Assignment Conditions
1	Anxious	Any	Optimism, Insecurity and Discomfort 4 or 5
2	Vanguard	AI, with or without GRC	Optimism, Innovativeness and Willingness to Present 4 or 5
3	Explorer	AI only	Remaining Respondents
4	Equipped	AI and GRC	Remaining Respondents
5	Foundation	GRC only	Optimism and Innovativeness 4 or 5
6	Bureaucrat	GRC only	Remaining Respondents
7	Believer	Neither	Innovativeness 4 or 5 and Willingness to Present 3 or Lower
8	Complacent	Neither	Governance Level 1 and Insecurity 1 or 2
9	Overwhelmed	Neither	Expected Change Over 3 Years 4 or 5 and Governance Level 1, Level 2 or not Reported
10	Observer	Neither	Remaining Respondents


Domain Analysis

Figure 10 reports the share of respondents meeting the maturity criterion in each domain. Only 14.9% of respondents reported a governance level of 4 or higher at their organizations. By contrast, 69.8% of the 2,557 respondents who answered the audit coverage question were outside the reported blind spot (30.2%). The shares of respondents reporting a strong belief in the potential of AI (72.6%), exposure (49.1%) and a strong willingness to act (43.7%) range between 40% and 75%. Governance specifically lags behind. Although individual belief is strong and exposure is close to half the sample, governance maturity has not kept pace.

Most respondents reported some AI audit coverage, but the remaining 30.2% reported either no coverage or uncertainty about whether it exists. For a profession responsible for assurance, that share is still too high. This is the blind spot examined in the Blind Spot Analysis, where ethics governance maturity is identified as a consistent predictor of reported coverage. A majority of people believe in AI; the challenge is to build governance around use.

Because each domain is reported independently, a respondent can meet the criterion in one domain without meeting it in another. Belief exceeds willingness to act by about 29 points, reflecting the confidence gap explored in the Behavioral Analysis.

Figure 10: Perceived Maturity by Adoption Domain



Note: Domain phases are reported independently. Shares are of respondents who answered each item (n = 2,557 to 2,924).

Behavioral Analysis

WILLINGNESS TO ACT

The reported odds ratios should be read as averaged estimates across varying response thresholds rather than as constant effects. They establish which domains are most strongly associated with willingness to act and in what order, though the strength of each association varies depending on the portion of the scale being compared. A partial proportional odds specification, which allows the association to differ at each threshold, is the next recommended step in the research. This is because what moves someone from hesitant to neutral may not be what moves them from neutral to confident, and a single common odds ratio cannot tell the two apart.

Table 4: Relationships Between Individual Factors and Willingness to Present AI-Assisted Findings

Predictor	Category	Willingness to Present	Odds Ratio	p-value	Interpretation
Belief that AI improves audit quality	Belief Driver	+18.5	2.55	< .001	Strong positive association
Trust under human review	Belief Driver	+7.4	1.45	< .001	Positive association
Perceived AI risk	Belief Inhibitor	+0.2	1.01	.829	No meaningful direct association
Unease about advanced technology	Belief Inhibitor	-1.4	0.93	.064	Small, borderline negative association
AI-tool exposure	Exposure	Binary contrast	1.16	.054	Small, borderline positive association
GRC-tool exposure	Exposure	Binary contrast	0.96	.628	No meaningful direct association
Director compared with Staff	Seniority	Role contrast	1.94	< .001	Higher stated willingness
C-Suite compared with Staff	Seniority	Role contrast	2.06	< .001	Higher stated willingness

Behavioral Analysis

TECHNOLOGY EXPOSURE

Following a series of robustness checks, the role of AI-tool exposure appears more modest than reported in the preliminary findings briefing (Banach et al., 2026). Respondents who had used at least one AI tool had about 16% higher odds of reporting greater willingness to present (odds ratio = 1.16, $p = .054$), holding the other variables in the model constant. In other words, AI-tool exposure has a small and only marginally significant association with willingness to act.

GRC-tool exposure, however, showed no statistically significant association with willingness to present (odds ratio = 0.96, $p = .628$). Together, these results narrow the role of technology exposure. Prior experience with industry technology such as GRC platforms does not appear to be a factor distinguishing respondents who are willing to present AI-assisted findings.

READINESS DIMENSIONS

Table 5 illustrates the above average technology readiness of internal audit professionals. Professional optimism about AI's potential is high, with respondents rating the belief that AI can improve audit quality at an average of 3.99 out of 5, while willingness to present AI-assisted findings received a lower average rating of 3.28. The mean difference between these two items was 0.708, with a 95% margin of error of about ± 0.04 . Mean scores above the scale midpoint for drivers and below it for inhibitors indicate a generally favorable orientation toward AI; however, the sample is notably less confident about presenting AI-assisted work (see Figure 11). This difference indicates the measurable gap between belief in AI capabilities and stated willingness to act. We have identified this as the confidence gap.


Table 5: Mean Responses to Technology Readiness Items

Technology Readiness Beliefs	Dimension	Mean
AI can improve the quality of audit work	Driver	3.99
I trust AI results when reviewed by humans	Driver	3.79
I feel uneasy about relying on advanced technology in audit work	Inhibitor	2.68
AI poses unacceptable risks to internal audit	Inhibitor	2.40

Technology Readiness Action	Domain	Mean
I would present AI-assisted findings to the board/audit committee	Action	3.28

Behavioral Analysis

Figure 11: Technology Readiness Responses and the Confidence Gap



Belief that AI can improve audit quality was the strongest predictor of willingness to present AI-assisted findings. A one-point increase in this belief was associated with 2.55 times the odds of reporting a higher willingness to present. Holding all other variables constant, the estimate was statistically significant at 0.1% ($p < .001$).

Expressed as an estimated probability change, an increase in belief corresponded with an approximately 18.5 percentage point increase in the probability of meeting the high-action threshold for willingness to act. Positive beliefs do not ensure that a professional will feel comfortable presenting AI-supported conclusions, but respondents who see a stronger connection between AI and audit quality report substantially greater willingness to do so.

Trust in AI results, when reviewed by humans, was also positively associated with willingness to present. A one-point increase in the trust item was associated with approximately 1.45 times the odds of reporting a higher action level. The estimate was statistically significant at the 0.1% level as well ($p < .001$). The associated increase in the probability of meeting the high-action threshold was approximately 7.4 percentage points. Meeting the threshold in this case means rating willingness to present at 4 or 5 and is an indicator of domain maturity.

Behavioral Analysis

The wording of the survey item is important to clarify. Respondents were asked about trust when human review was present. The results, then, reflect confidence in AI-supported and professionally reviewed processes. The item does not measure trust in fully autonomous decision making. This finding supports the inclusion of human review within the design of AI-enabled work. The results also suggest that human review may contribute to trust when it is defined as an observable process.

Unease showed a small, negative association. Its estimated odds ratio was 0.93 ($p = .064$), corresponding to an estimated reduction of approximately 1.4 percentage points in the probability of meeting the high-action threshold. Perceived AI risk was not meaningfully associated with willingness to present after controlling for other variables. The estimated odds ratio was 1.01 ($p = .829$), corresponding to an approximate 0.2 percentage point change in the predicted probability.

These findings suggest concern and positive orientation can coexist. A professional can recognize AI's value while still feeling uneasy about its safety or reliability. Although the survey's risk and unease measures capture related concerns, they represent distinct experiences. The lack of a strong direct relationship between perceived risk and willingness to present also limits interpretations based on general risk aversion. More specific measures would be required to evaluate concerns involving professional liability, accountability, reputational consequences or client acceptance.


SENIORITY

The confidence gap appears to narrow among more senior respondents. Staff auditors show the widest gap and the C-Suite the smallest (see Figure 12). Professional level was observed to be associated with the action measure. Compared to staff auditors or the staff level, Directors had 1.94 times the odds of reporting a higher level of willingness to present. Likewise, C-Suite respondents had approximately 2.06 times the odds. Both comparisons also showed statistically significant results at the 0.1% level (see Table 4).

Senior professionals may possess greater decision authority, more experience communicating with boards, more control over review processes or greater familiarity with defending use of professional judgment. They may also face a different level of career and reputation risk when presenting new forms of evidence. Further analysis is required to determine whether the pattern remains valid within firms, or whether it is an attribute of the sample in attendance at the training event where the survey was conducted. Although these findings are informative, a separate analysis found that seniority alone explained only about 0.6% of the variation in the confidence gap.

Behavioral Analysis

Figure 12: Confidence Gap by Seniority



Note. Error bars show 95% confidence intervals. Colors group levels by gap size and do not indicate statistically significant differences.

The findings also revise the earlier claim that tooling matters approximately 3.5 times as much as seniority. The revised ratio is approximately 1.2. The earlier comparison divided a binary tool contrast by a one-level change in an ordinal seniority measure. When both predictors were standardized, their estimated effects were 0.146 for AI-tool exposure and 0.123 for seniority. The model only explained about 3.8% of the variation in the Readiness Score. The updated evidence supports a modest association between exposure and readiness without establishing that providing a tool creates professional capability.

Another finding is that VPs and directors were more likely to perceive governance as mature at their firms, whereas C-Suite and staff auditors were more likely to perceive governance as ad hoc rather than concrete. This means that the two ends of the seniority ladder have a different perception of governance than the middle.

The survey collected responses on professional levels but not decision rights, permission, communication experience or perceived consequences; therefore the mechanisms behind the observed relationships remain unknown. Many junior staff reported high willingness to present, while some senior professionals remained hesitant. Efforts to encourage responsible AI use should therefore extend beyond assumptions based on seniority alone.

Blind Spot Analysis


BLIND SPOT

Of the 2,557 respondents who rated "none/I don't know" as a part of the audit coverage question, 773 (30.2%) rated it 4 or 5. Because none and don't know share a single prompt, a high rating could reflect a genuine absence of audit activity or assurance processes that exist but are not known to the respondent.

Governance maturity was a consistent predictor of the blind spot. While holding beliefs, tool exposure and seniority constant, the adjusted probability of falling into the blind spot declined from 40.9% at Level 1 to 22.2% at Level 5 (see Figure 13). The progression is not perfectly linear and the largest difference falls between basic and managed governance. Although governance maturity was not associated with willingness to present, respondents reporting higher governance levels were less likely to fall into the blind spot.

This could mean that an organization's institutional infrastructure matters more than individual perception in determining whether respondents are aware of AI within the audit process. Governance appears to be a high-value intervention point, although these cross-sectional data do not establish causation.

Figure 13: Adjusted Blind Spot Probability by Governance Maturity



Note. Adjusted probabilities (n = 2,016) holding beliefs, tool exposure and seniority at observed values. Error bars show each level's margin of error at the 95% level. The dashed line marks the overall 30.2% (2,557 respondents). Colors group levels by size and do not indicate statistically significant differences.


Blind Spot Analysis

The 30.2% should be read as a warning signal rather than a precise count. Of the 773 respondents in the blind spot, 527 also rated at least one component 4 or 5, including 164 who gave the same 4 or 5 to every row. These contradictory answers are most common at Level 1 and carry much of the decline across governance levels (see Figures 14 and 16).

With respondents who gave identical ratings to all eight rows removed, the blind spot still falls from 33% at Level 1 to 20% at Level 5. A stricter blind spot includes 190 respondents (7.4%) and also declines, from 10.9% at Level 1 to 1.4% at Level 5. The stricter blind spot is limited to respondents who also rated every audit component 1 or 2 or answered only the "none/I don't know" row. The relationship with governance therefore holds under both definitions.

The stricter blind spot also points to an awareness gap. Under the stricter definition, staff auditors make up 22% of the blind spot, compared with 12% of those who reported consistent coverage. If employees cannot identify whether AI is governed or audited, the control environment may lack transparency, responsibility may sit with too few people or assurance results may not reach the people who rely on them.

Figure 14: Response Patterns Within the Blind Spot



Note. Each cell is the average rating that group gave the prompt, from 1 (not at all) to 5 (extensively). Shares are of the 2,557 respondents who answered None/I don't know and sum to the 30.2% blind spot.

Blind Spot Analysis

OVERSIGHT

Respondents were asked how extensively seven components were reviewed in their organization's AI audits, rated from 1 (not at all) to 5 (extensively). The components were LLM prompt logic, training data bias controls, model version and change control, output validation, human override processes, governance and ethics. In total, 2,629 respondents rated at least one component and 2,415 rated all seven.

Every component averaged close to the midpoint, from 2.93 for governance to 3.18 for LLM prompt logic, and between 35.1% and 42.2% of respondents rated each one 4 or 5. Technical controls were rated slightly higher than governance and ethics, but the spread from first to last is only a quarter of a point.

Ethics was often the component left out. About a quarter of respondents who rated another component 4 or 5 rated ethics 1 or 2, so ethics review cannot be assumed to accompany technical review. Governance and ethics are statistically indistinguishable ($p = .614$), so neither should be described as the least reviewed component on its own. The governance component refers to governance reviewed within AI audits and is a separate measure from organizational governance maturity.

Nearly half of respondents described coverage as all or nothing. Among those who rated all seven components, 31.4% rated none of them 4 or 5 and 14.6% rated all seven 4 or 5. The remaining 54.0% reported extensive review of some components but not others. Many of the all-or-nothing respondents gave every component the same rating (see Figure 14). Robustness checks identified and omitted these responses to control for bias.

Believing that AI improves audit quality provides no certainty on whether an organization is actually auditing it. The two measures are not meaningfully related ($r = -0.02$). These forms of adoption are widely assumed to move together. The responses demonstrate that this is not always the case. The domains of belief (people believe in the potential of AI), exposure (people are using the tools), governance and oversight (governance and review protocols exist) appear to mature at varying rates and in no particular order.

The respondent's answer does not reflect whether an organization's AI audit procedures are comprehensive, effective, well designed or supported by appropriate evidence. A respondent who reported that at least one AI area is audited may still work in an environment with inadequate coverage. Conversely, rating "none/I don't know" a 4 or 5 could indicate either a genuine absence of audit activity or a failure in which assurance processes are not known.

Blind Spot Analysis


GOVERNANCE

Of the 2,924 respondents who answered the governance question, approximately 85% reported governance maturity within Levels 1 through 3. Only about 15% reported Level 4 or Level 5 governance, leaving a majority of respondents in environments ranging from ad hoc to managed governance instead of integrated or leading governance. Level 3 was the most common response (35.4%), and only 3.0% reported Level 5 (see Figure 15).

Governance refers to the organizational structures through which AI use is identified, authorized, reviewed, documented and monitored. This makes the concentration at the lower and middle levels so much more significant. Without operational governance, an employee can believe AI is valuable without knowing whether its use is permitted. A firm can deploy an AI tool without assigning accountability for its outputs, and an audit group can be aware of AI use cases without maintaining a complete inventory of approved applications.

Higher governance maturity was associated with a smaller blind spot but not with more extensive audit coverage. Ratings for all seven components declined slightly as governance maturity increased. The association did not differ significantly across the seven components ($p = .648$). It held after removing respondents who gave every component the same rating, clustering responses by email domain and controlling for response style.

Figure 15: Governance Maturity Distribution




Note. Self-reported AI ethics governance maturity among the 2,924 respondents who answered the item. Descriptions follow the survey wording.

Blind Spot Analysis

Across ethics governance levels, respondents differed in how reliably they described coverage, not in how extensive they reported it to be (see Figure 16). The share of respondents who gave the same 4 or 5 on all eight prompts fell from 13.8% at Level 1 to 2.4% at Level 4. The share who contradicted themselves, rating “none/I don’t know” 4 or 5 while also rating a component 4 or 5, fell from 16.8% to 11.6%. Both patterns made coverage at Level 1 look more extensive. With both groups set aside, the share of 5s still fell from 18% to 10%, indicating respondents at more mature levels may have a higher bar for calling audit review extensive. Respondents who reported more mature AI ethics governance also gave more consistent, and therefore more reliable, accounts of their own audit coverage.

Figure 16: Contradictory Coverage Answers by Governance Maturity



Note. Respondents who answered all eight prompts (n = 2,265). Both groups rated “none/I don’t know” and at least one component 4 or 5. Level 5 is small.

Ethics governance was not linked to ethics review in particular. The governance item asked specifically about AI ethics, yet the ethics component moved with governance in the same way as the technical components. Relative to each respondent’s other ratings, ethics did not rise as governance rose. Setting aside respondents who gave identical ratings to every prompt, between 26% and 29% of those who rated another component 4 or 5 rated ethics 1 or 2 at every governance level. That includes 28.4% of those at Levels 4 and 5, where governance is described as integrated or leading.


The survey measured governance maturity through a single self-reported ethics governance item. Respondents assigned their own organization to a level, and the measure was not validated against documented policies, control inventories or audit workpapers. Two organizations with identical control environments may therefore be recorded at different levels. Governance should be interpreted as organizational context rather than an independently verified control assessment.

Blind Spot Analysis

MATURITY

The behavioral analysis examined individual conditions associated with willingness to act on AI-assisted work. The blind spot analysis examined organizational conditions associated with whether AI falls within the audit process. Read side by side, the two point to different conditions. Greater odds indicate more willingness to present in the left panel and a higher likelihood of falling into the blind spot in the right panel (see Figure 17).

Figure 17: Factors Associated with Willingness to Present and the Blind Spot



Note. Odds ratios from the willingness-to-present model (n = 2,291) and the blind spot model (n = 2,016), each holding the other factors constant. Belief items are per scale point, seniority is compared with Staff and governance enters both models as a single trend, so its odds ratio is per level rather than a comparison with Level 1. Awareness of AI audit coverage (at least one audit component reported) is a control in the willingness model only, since coverage is the outcome of the blind spot model. Values above 1 mean higher odds of the outcome. Shading marks the strongest association in each model.

At the organizational level, governance maturity was the most consistent correlate of the blind spot, with lower odds at every level above Level 1. Belief that AI can improve audit quality was unrelated to reported coverage ($r = -0.02$). It had only a small association with the blind spot, in the direction of a larger blind spot rather than a smaller one. Tool exposure was not associated with the blind spot once governance and the other factors were held constant. Respondents who used AI tools were in the blind spot at almost the same rate as those who did not (29.1% compared with 31.4%).

Blind Spot Analysis

At the individual level, optimism and innovativeness were the strongest of the four readiness dimensions in predicting willingness to present AI-assisted findings (odds ratios of 2.55 and 1.45). Willingness was also higher among Directors and the C-Suite than among staff auditors. AI-tool exposure had a small and only marginally significant association. GRC-tool exposure, perceived risk and governance maturity showed no meaningful association.

Individual action and organizational assurance are therefore related but distinct outcomes. One organization may have employees willing to advocate for AI-supported conclusions while lacking mature governance and audit coverage. Another may maintain formal controls while its employees remain hesitant to present AI-assisted work. Belief, exposure and governance do not move together in these responses and can mature at different rates.

While AI adoption describes whether a technology is available or used, assurance describes whether the risks, processes, outputs and controls surrounding that use are subject to evaluation. An organization may acquire an AI tool quickly while needing considerably more time to establish approved-use policies, validation procedures, documentation standards and escalation protocols.

Most AI initiatives are framed and measured as adoption initiatives, with success defined by licenses purchased, employees trained, pilots launched or time saved. None of these measures provide much information about whether AI-generated work is being adequately controlled and monitored or whether its risks have been incorporated into an audit plan. Measures of assurance, such as the share of AI use cases that have been inventoried, rated for risk and audited, would show whether oversight is keeping pace with adoption.

Conclusions About the Blind Spot:

1. Three in ten respondents were in the blind spot. Although most respondents believed AI can improve audit quality, belief was unrelated to reported coverage.
2. Tool use is not a substitute for governance. Introducing an AI tool does not ensure it is used and monitored under auditing standards. AI-tool users were in the blind spot at the same rate as non-users, so about three in ten were not clear on audit coverage.
3. General risk awareness is not operational assurance. Perceived AI risk was not associated with willingness to present or with the blind spot. Professionals may recognize that AI involves risk while lacking the authority to initiate processes.
4. Governance maturity was the clearest organizational predictor of the blind spot. Higher governance did not come with more extensive review of individual components, and ethics review did not keep pace with ethics governance.

Personas


ADOPTION SEGMENTS

Grouping similar personas by shared attributes reveals broader patterns that can help leaders decide where to begin. As described in the research design, Ward hierarchical clustering grouped the ten persona averages into four adoption segments. Tooled Adopters account for 45.4% of eligible respondents and include Vanguard, Explorer and Equipped. Governance-First accounts for 20.3% and includes Foundation and Bureaucrat. The Untooled Middle accounts for 29.9% and includes Believer, Observer, Overwhelmed and Complacent. Risk-Conflicted accounts for 4.4% and consists of Anxious (see Figure 18).

The segments provide broad planning categories. The personas retain the distinctions needed to consider an intervention within each category. Neither level should be interpreted as a set of fixed psychological types. Blue identifies Tooled Adopters, green Governance-First, orange the Untooled Middle and red Risk-Conflicted. The colors illustrate persona segments, rather than a ranking or sequence of adoption. Governance-First refers to GRC exposure without reported AI exposure; the name does not establish mature AI governance. Risk-Conflicted is defined by high belief alongside high perceived risk and unease, regardless of tool exposure.

Figure 18: Adoption Segments and Their Personas

Share of eligible respondents · n = 2,847 complete-sentiment cases



Colors link each segment to its constituent personas.

All percentages use the same denominator. Segment totals sum to 100.0%.

Personas

PERSONA APPLICATIONS

The ten personas operationalize the report's broader adoption framework as an intervention-oriented tool. The persona landscape is color coded by segment and arranged by normalized readiness, demonstrating that members of the same segment can have varying degrees of readiness (see Figure 19). The framework supports the use of secondary conditions to modify primary personas. A respondent may resemble one intervention profile while displaying an additional feature that matters for implementation. In practice, this would begin with assigning a primary persona, then identifying a secondary condition that is materially relevant to select a limited intervention.

Figure 19: Persona Landscape by Normalized Readiness

n = 2,847 complete-sentiment cases · Sorted by normalized readiness

● Tooled Adopters ● Governance-First ● Untooled Middle ● Risk-Conflicted

Persona	Readiness 1-10 scale	Share of sample	AI exposure	GRC exposure	Selected defining conditions
Vanguard	8.29	15.9%	Present	Mixed 32.2% present	Belief, trust and willingness to present each rated 4-5
Foundation	7.78	10.4%	None	Present	Belief and trust both rated 4-5
Complacent	7.02	3.4%	None	None	Governance Level 1; perceived risk rated 1-2
Equipped	6.62	9.1%	Present	Present	Remaining respondents with both AI and GRC exposure
Believer	6.58	9.1%	None	None	Trust rated 4-5; willingness to present rated 1-3
Observer	6.58	12.2%	None	None	Remaining respondents after all earlier assignment rules
Explorer	6.48	20.4%	Present	None	Remaining respondents with AI exposure only
Overwhelmed	6.41	5.2%	None	None	Expected change rated 4-5; governance Level 1, 2 or missing
Anxious	5.78	4.4%	Not specified*	Not specified*	Belief, risk and unease each rated 4-5
Bureaucrat	5.74	9.9%	None	Present	Remaining respondents with GRC exposure only

Conditions summarize the first-match rules in Table 3. Item ratings use a 1-5 scale.

*Anxious can include any tool combination; exposure is not an assignment criterion.

Readiness is descriptive and does not determine persona assignment.

Personas

HIDDEN CHAMPIONS

As an example, Hidden Champions were identified using seniority as a secondary condition. The segmentation flagged 71 staff auditors who reported AI-tool exposure and scored at least 4 on both belief that AI could improve audit quality and trust in AI results under human review. Despite this, their average confidence gap was 0.957, exceeding the mean for all staff auditors as well as the mean for every other level of seniority.

A plausible interpretation is that staff auditors are the most likely to require explicit permissions and operate under the closest review, with the greatest exposure to consequences. At this level, willingness to present could reflect the perceived cost of being wrong as much as confidence in the tools. If so, staff auditors with practical exposure and favorable attitudes would benefit from institutional structure and support to increase their willingness to employ approved use cases in their work product.

In this case, the confidence gap does not arise from low willingness to act but from above average scores of optimism. The average willingness-to-present score was 3.56, above the neutral midpoint, and less than half of the 71 respondents (46.5%) scored below the high-action threshold.

PERSONA LIMITATIONS

The persona framework is expected to have weak external validity beyond the internal audit profession; however, it is a practical, rule-based consulting tool. The assignment rules are tested for overlap between personas before the priority order is applied, for stability when the score thresholds are loosened or tightened and for how much of the variation in reported AI audit coverage the personas explain.

Under more lenient thresholds, 92.0% retained the same label. Under stricter thresholds, only 69.9% retained their original assignment. These results indicate moderate stability under the lenient conditions and substantially weaker stability under stricter conditions. Persona assignment is based on varying domain maturity and statistically explains less than 1% of the variation in reported AI audit coverage.

The segmentation uses a complete-sentiment base of 2,847 respondents. Before applying a priority rule, 28.4% of eligible respondents met the conditions for at least two personas, with some meeting as many as three. This means that the chosen priority materially influences which primary persona a respondent receives. The persona framework has practical value, but the results show that assignment remains sensitive to rule design. It should not be used for employee evaluation, hiring, promotion or psychometric assessment.

Personas

TAILORED INTERVENTIONS

Despite the limitations, these personas can be used to organize a multidimensional adoption problem into recognizable profiles that can be addressed through different forms of support. The personas can guide team-level planning, training design, pilot selection and governance conversations. The intervention overview illustrated in Figure 20 is an evidence-informed operationalization of the profiles. Each intervention targets one of the five domains. Organizations should evaluate these interventions through bounded pilots and subsequent outcome measurement.

Figure 20: Persona Conditions and Proposed Interventions

Defining conditions inform practical actions to evaluate through bounded pilots

● Tooled Adopters ● Governance-First ● Untooled Middle ● Risk-Conflicted

Persona	Measured profile conditions	Proposed intervention
Vanguard	AI exposure; belief, trust and willingness to present each 4–5	Extend advanced use with formal validation and documented review.
Explorer	AI exposure without GRC exposure; remaining after earlier rules	Add proportionate governance around practical experimentation.
Equipped	Both AI and GRC exposure; remaining after earlier rules	Review workflows, handoffs and accountability for AI output.
Foundation	GRC exposure without AI exposure; belief and trust both 4–5	Launch a controlled first AI use case with review and clear ownership.
Bureaucrat	GRC exposure without AI exposure; remaining after earlier rules	Identify barriers to a first pilot and clarify risk-based approvals.
Believer	Neither AI nor GRC exposure; trust 4–5 and willingness 1–3	Provide supported experience with human-reviewed AI output.
Observer	Neither AI nor GRC exposure; residual after earlier rules	Diagnose the relevant barrier and choose a role-specific next step.
Overwhelmed	Neither tool; expected change 4–5; governance Level 1, 2 or missing	Start with one manageable use case and build governance alongside it.
Complacent	Neither AI nor GRC exposure; governance Level 1 and risk 1–2	Calibrate risk awareness and establish minimum controls.
Anxious	Belief, risk and unease each 4–5; any AI/GRC exposure combination	Use evidence, validation and a bounded, reversible pilot.

The survey did not test intervention effects. Evaluate each action through a bounded pilot. Conditions summarize Table 3; item ratings use a 1–5 scale. Willingness refers to presenting AI-assisted findings. Expected change refers to the next three years.

Personas

Vanguards report AI exposure, strong belief in AI's contribution to audit quality, substantial trust under human review and high willingness to present AI-assisted work. They represent a comparatively activated form of adoption in which practical exposure and professional orientation are accompanied by stated action. Their average reported governance level is about 2.6 out of 5, so high willingness here does not come with mature governance.

Recommended Intervention: Vanguards may be appropriate candidates for advanced use cases, controlled peer coaching, champion networks and participation in formal pilots. Their work should remain subject to defined review requirements, documentation standards and escalation procedures. Before extending advanced use, confirm that each use case has a named reviewer and a documented validation step.

Central Organizational Question: How can the organization use these professionals as experienced participants while preserving validation, documentation and review standards?

The Vanguard

AI adoption persona

• **Tooled Adopters**



SELECTED PROFILE CONDITIONS

 **AI exposure**

Present

All respondents report AI exposure.

 **GRC exposure**

Mixed

32.2% 67.8%
32.2% report GRC; 67.8% are AI-only.

 **Belief, trust and action**

Rated 4 or 5

1 2 3 4 5
Favorable belief and trust with stated action.

 **Governance maturity**

Mean 2.6 / 5

1 ——— 5
Exposure and action do not imply mature governance.

8.29 / 10

Normalized readiness
1-10 scale; descriptive score

15.9%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Advanced use with formal validation

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

Explorers report AI-tool exposure without GRC-tool exposure and do not meet the earlier Anxious or Vanguard rules. They are the largest persona (20.4%), and their mean willingness to present is 2.83, below the sample mean of 3.28. Their tool use is not evidence of audit coverage, since AI-tool users fell into the blind spot at almost the same rate as non-users (29.1% compared with 31.4%).

Recommended Intervention: Explorers may benefit from controlled sandboxes, approved-tool lists, defined review gates, GRC onboarding and clear criteria for moving a use case from experimentation into their routine. The goal is to make activity visible and reviewable. Before expanding a pilot, confirm who approves the use case, reviews the output and records exceptions.

Central Organizational Question: How can useful experimentation be preserved while introducing the controls needed for responsible use?

The Explorer

AI adoption persona

• **Tooled Adopters**



SELECTED PROFILE CONDITIONS

 **AI exposure**

Present

All respondents report AI exposure.

 **GRC exposure**

None

No reported GRC-tool exposure.

 **Willingness to present**

Mean 2.83 / 5

1 ————— 5
Stated action is more moderate than exposure.

 **Assignment condition**

AI-only respondents

Assigned after earlier priority rules are applied.

6.48 / 10

Normalized readiness
1-10 scale; descriptive score

20.4%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Introduce governance around experimentation

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

Equipped respondents report the presence of both AI and GRC-related tools, but their stated action is moderate. The necessary technological components may be available without producing a willingness to present AI-assisted work. GRC-tool exposure showed no association with willingness to present (odds ratio 0.96) or with the blind spot (0.89), which fits a group whose tools have not produced action (see Figure 17).

Recommended Intervention: The Equipped profile calls for workflow redesign rather than additional tool acquisition alone. Relevant steps include assigning use-case ownership, clarifying role responsibilities and measuring whether the available tools are producing documented and reviewable outcomes. In the first redesigned workflow, record who owns each AI-assisted output and whether it was reviewed and documented.

Central Organizational Question: Why are existing tools and control infrastructure not producing stronger workflow integration or responsible action?

The Equipped

AI adoption persona

• **Tooled Adopters**



SELECTED PROFILE CONDITIONS



AI exposure

Present

All respondents report AI exposure.



GRC exposure

Present

All respondents also report GRC exposure.



Willingness to present

Mean 2.93 / 5

1 ————— 5
Stated action is more moderate than exposure.



Assignment condition

Remaining dual-tool users

Assigned after earlier priority rules are applied.

6.62 / 10

Normalized readiness
1-10 scale; descriptive score

9.1%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Redesign workflows and accountability

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

Foundation respondents report GRC-tool exposure and a positive orientation toward AI, but no direct AI exposure. Their organizations may possess a control base that could support adoption, while practical use remains underdeveloped. Foundation has the second-highest readiness of the ten personas (7.78) and meets the same belief and trust conditions as Vanguard, but without AI tools.

Recommended Intervention: Foundation respondents would likely benefit from a guided pilot built around an accounting or audit task with limited risk and clearly defined success criteria. Existing governance structures should be used to specify approval, evidence, human review and stop conditions before the pilot begins. After the pilot, ask participants the survey's willingness item again to see whether belief has turned into stated action.

Central Organizational Question: Which bounded, low-risk use case can activate the existing control environment?

The Foundation

AI adoption persona

• Governance-First



SELECTED PROFILE CONDITIONS



AI exposure

None

No reported AI-tool exposure.



GRC exposure

Present

All respondents report GRC exposure.



Belief and trust

Rated 4 or 5

Favorable responses distinguish this GRC-only profile.



Technology combination

GRC only

A control-tool foundation without direct AI exposure.

7.78 / 10

Normalized readiness
1-10 scale; descriptive score

10.4%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Activate a controlled first AI use case

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

Bureaucrat respondents report GRC-related tool exposure, but low activation and no direct AI exposure. The surrounding control environment may be more developed than the organization's practical use of AI. Bureaucrats have the lowest readiness of the ten personas (5.74). They have the same limited exposure as the Foundation persona but did not rate both belief and trust 4 or 5.

Recommended Intervention: Relevant actions may include risk-tiered permission, simplified approval for low-risk applications, bounded pilots and a clearer distinction between prohibited, restricted and permitted uses. The objective is to retain appropriate control while reducing unnecessary barriers to limited and reviewable experimentation. For an initial pilot, track approval times and record which requirements delayed or prevented participation.

Central Organizational Question: Are current approval requirements or control processes making responsible experimentation unnecessarily difficult?

The Bureaucrat

AI adoption persona

• Governance-First



SELECTED PROFILE CONDITIONS



AI exposure

None

No reported AI-tool exposure.



GRC exposure

Present

All respondents report GRC exposure.



Willingness to present

Mean 2.60 / 5

1 ————— 5

Infrastructure is present without direct AI exposure.



Assignment condition

Remaining GRC-only users

Assigned after earlier priority rules are applied.

5.74 / 10

Normalized readiness
1-10 scale; descriptive score

9.9%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Simplify risk-tiered approval

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

The Complacent label refers to a combination of low reported concern, weak governance and no reported AI or GRC exposure. The term describes the rule-based survey pattern and does not constitute a judgment about an individual's character or diligence. All Complacent respondents report governance at Level 1, the level with the highest adjusted blind spot probability (40.9%, see Figure 13).

Recommended Intervention: This profile may benefit from scenario-based risk calibration, baseline governance education, minimum control expectations and examples of realistic failures involving confidentiality, unsupported conclusions, incomplete validation or inappropriate reliance. Before any tool is introduced, confirm that the team has a written acceptable-use rule and a named reviewer.

Central Organizational Question: Is the low level of concern based on informed experience, or has the respondent had little opportunity to evaluate the relevant risks?

The Complacent

AI adoption persona

Untooled Middle



SELECTED PROFILE CONDITIONS

 **AI exposure**

None

No reported AI-tool exposure.

 **GRC exposure**

None

No reported GRC-tool exposure.

 **Governance maturity**

Level 1

1 2 3 4 5
All respondents report the lowest governance level.

 **Perceived risk**

Low reported risk

Low concern is a defining profile condition.

7.02 / 10

Normalized readiness
1-10 scale; descriptive score

3.4%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Calibrate risk and establish minimum controls

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

Believers report substantial trust under human review, but express low willingness to act and report no direct exposure to AI tools or GRC technology. Their positive orientation has not yet developed into regular, authorized or presentable professional use. Their rule pairs high trust (4 or 5) with low willingness (3 or lower), and they report no tools, so what they lack is practice rather than conviction.

Recommended Intervention: Believers may benefit from hands-on demonstrations, role-specific training, low-risk practice opportunities, explicit acceptable-use guidance and clearly defined review and escalation procedures. The emphasis should remain on converting general confidence into task-specific and supported experience. Ask participants the willingness to act question before and after supported practice with reviewed AI output.

Central Organizational Question: What practical condition is missing between favorable belief and responsible use?

The Believer

AI adoption persona

• Untooled Middle



SELECTED PROFILE CONDITIONS

 **AI exposure**

None

No reported AI-tool exposure.

 **GRC exposure**

None

No reported GRC-tool exposure.

 **Trust under human review**

Rated 4 or 5

1 2 3 4 5
High trust in AI results when reviewed by humans.

 **Willingness to present**

Rated 1 to 3

1 2 3 4 5
Limited stated willingness to present AI-assisted work.

6.58 / 10

Normalized readiness
1-10 scale; descriptive score

9.1%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Provide supported practical experience

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

Observer is the residual profile for respondents without AI or GRC tools who do not meet more specific criteria. At 12.2%, Observers are the third-largest persona, and their readiness (6.58) sits in the middle of the ten. The absence of tool exposure does not necessarily reflect similar beliefs, concerns or support needs.

Recommended Intervention: Observers may benefit from peer examples, concise demonstrations, visible leadership support, role-specific next steps or time-limited participation in a pilot. Because this profile is residual, the recommended response should begin with further diagnosis rather than an assumption that every Observer faces the same barrier. A short follow-up question asking which of the four barriers in the central question applies would turn the residual label into a specific next step.

Central Organizational Question: What evidence, permission, experience or role-specific relevance would support the next step?

The Observer

AI adoption persona

• Untooled Middle



SELECTED PROFILE CONDITIONS



AI exposure

None

No reported AI-tool exposure.



GRC exposure

None

No reported GRC-tool exposure.



Assignment condition

Residual profile

No earlier persona rule takes precedence.



Interpretation

Further diagnosis needed

The residual label does not identify one shared barrier.

6.58 / 10

Normalized readiness
1-10 scale; descriptive score

12.2%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Diagnose the next role-specific step

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

Overwhelmed respondents report no AI or GRC-related tools, weak governance and high expectations concerning future technological change. The profile combines substantial anticipated change with limited present infrastructure. Their governance is at Level 1 or 2 or was not reported, and Levels 1 and 2 have the highest adjusted blind spot probabilities (40.9% and 34.0%, see Figure 13).

Recommended Intervention: A sequenced implementation plan is more appropriate than a broad deployment. Organizations may begin with one narrow use case, foundational training, basic acceptable-use guidance and a small number of review requirements. Limiting the initial scope may help the organization build practical capacity. For the first use case, record which governance requirement was in place before it went live.

Central Organizational Question: Which minimum capability and governance requirement should be introduced first?

The Overwhelmed

AI adoption persona

• Untooled Middle



SELECTED PROFILE CONDITIONS



AI exposure

None

No reported AI-tool exposure.



GRC exposure

None

No reported GRC-tool exposure.



Governance maturity

Weak or missing

Limited governance is part of the profile condition.



Expected future change

High

Substantial anticipated change with limited infrastructure.

6.41 / 10

Normalized readiness
1-10 scale; descriptive score

5.2%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Sequence adoption and reduce initial scope

Rule-based intervention profile. Readiness is not the assignment cutoff.

Personas

Anxious respondents report high unease and perceived risk, alongside favorable beliefs about AI's value. The label reflects high responses on the survey's concern items and should not be interpreted as a personality assessment. In the blind spot model, unease was associated with higher odds of the blind spot (odds ratio 1.10, see Figure 17), so their concern may partly reflect not knowing what assurance exists.

Recommended Intervention: Anxious respondents may benefit from transparent risk evidence, defined validation protocols, reversible pilots, explicit human-review requirements and documented incident and escalation processes. The intervention should distinguish valid concerns from generalized uncertainty. Share each pilot's validation results with participants so concern can be checked against evidence.

Central Organizational Question: Is the reported concern calibrated to the use case, and which control would make participation acceptable?

The Anxious

AI adoption persona

• Risk-Conflicted



SELECTED PROFILE CONDITIONS

★ **Belief in AI value**

Rated 4 or 5

1 2 3 4 5
Favorable belief that AI can improve audit quality.

⚠ **Perceived risk**

Rated 4 or 5

1 2 3 4 5
High reported concern about AI risk.

⚠ **Unease with technology**

Rated 4 or 5

1 2 3 4 5
High reported unease with advanced technology.

🛡 **AI and GRC exposure**

Not an assignment criterion

The rule does not require a particular tool combination.

5.78 / 10

Normalized readiness
1-10 scale; descriptive score

4.4%

Share of respondents
Complete-sentiment sample n = 2,847

INTERVENTION PRIORITY

Provide evidence, validation and reversible use

Rule-based intervention profile. Readiness is not the assignment cutoff.

Recommendations

FIRM LEVEL FOCUS AREAS

Recommendations at the firm level are organized into five focus areas that address roadblocks identified in each domain. Decision makers are advised to assess their teams and organizations and begin with 1 or 2 domains that are furthest behind.

Focus Area 1: Human-Reviewed Trust (Cultivate Belief Maturity)

Trust under human review was associated with greater willingness to present (odds ratio 1.45, see Table 4). Be specific about what must be independently verified, and assign it to named reviewers. Maintain an audit trail of prompts, inputs, outputs and conclusions as part of a streamlined routine. Train professionals to challenge AI output and to document the process. Trust arises when review is present, not when it is assumed.

Focus Area 2: Responsible Exposure (Cultivate Exposure Maturity)

AI-tool exposure had a borderline association with willingness to present (see Table 4). Give people accounting-specific demonstrations, not just generic AI training. Prioritize a few confined, high-value use cases over unrestricted experimentation. Create a designated environment for approved tool use, with set windows to explore new use cases.

Focus Area 3: Governance and Permission (Cultivate Governance Maturity)

Adjusted blind spot probability fell from 40.9% at governance Level 1 to 22.2% at Level 5 (see Figure 13). Document when staff may engage in AI-assisted work and under what conditions they may present it to others. Give people an easy way to escalate cases they are unsure about. Have managers communicate acceptable use in workshops, meetings and documentation, and bring AI-related judgment into normal review discussions.

Focus Area 4: Risk and Assurance (Cultivate Oversight Maturity)

Three in ten respondents reported no AI audit coverage or did not know whether any existed (see Figure 10). Inventory where AI is used and classify each use by risk. Bring AI use into the internal audit plan instead of a separate program. Standardize reviews of the data flows behind each AI tool. Set up an incident and exception process before it is needed.

Focus Area 5: Persona-Tailored Change Management (Cultivate Action Maturity)

Directors and the C-Suite reported greater willingness to present than staff auditors (odds ratios of 1.94 and 2.06, see Table 4). Do not run a one-size-fits-all training program. Match the intervention to where an individual or team sits on the adoption curve in each domain, adjusting for seniority. Use the personas to guide a conversation, not to label individuals, and track movement between personas after interventions and over time.

Recommendations

AUDITING AI

Findings indicate a strong belief in AI's value, but belief does not reliably predict whether AI has entered the assurance process. If anything, a respondent with a stronger belief in the potential of AI had slightly higher odds of falling into the blind spot. This creates a problem, because most AI initiatives are framed and measured as adoption initiatives, not assurance initiatives. Our recommendation, after assessing the state of AI adoption in internal audit, is to structure an initial AI assurance program around six interlocking components.

1. Establish Organizational Visibility

Firms need to identify where AI is being used, by whom, for what purpose and with what data. Assurance cannot be scoped when AI activity remains informal or undocumented.

2. Assign Accountability

Each AI use case should have a business owner, a control owner and a designated party responsible for review or assurance. Accountability should never rest solely with the individual user.

3. Classify Risk and Materiality

AI uses must be differentiated by their potential impact on financial reporting, audit evidence, client communication, regulatory compliance, confidential information and professional judgment. Low-risk productivity uses warrant less scrutiny than applications that influence accounting estimates, control conclusions or audit opinions.

4. Define Evidence and Review

Governance must specify which inputs, prompts, outputs, validations, approvals and professional judgments should be documented. Human review should be a defined control rather than a general expectation.

5. Preserve Domain-Level Assurance

Audit planning should consider each relevant AI risk or control area separately. Organizations should avoid concluding that AI is 'audited' merely because one component, such as output validation, has been reviewed. Ethics needs particular attention, since about a quarter of respondents who rated another AI audit component 4 or 5 had rated ethics 1 or 2.

6. Monitor Change and Exceptions

AI tools and use cases will change after initial approval, particularly as adoption expands. Governance should therefore include periodic reassessment, change controls, incident reporting, exception review and escalation to appropriate leadership.

Recommendations

INDIVIDUAL INTERVENTIONS

A persona can identify the conditions that appear most relevant at a particular time for an individual, but it should not be interpreted as a permanent description of abilities, personality or motivation. Changes in belief, exposure, governance, oversight or actions are the intended outcomes of interventions; changes to the measures used for assignment would then change the profile that best describes a team member.

The personas are rule-based consulting profiles, built on emergent trends, that provide a practical vocabulary for navigating adoption objectives. The approach does not assume that the personas represent naturally occurring psychological types or stable categories within the wider internal audit population. Their principal purpose is to help organizations connect observable adoption conditions with more appropriate questions and interventions.

The adoption model, personas and focus areas provide a foundation for structured conversations. An organization can then identify conditions affecting its teams, select bounded interventions and monitor outcomes in targeted domains. Because focus areas 3 and 4 address organizational conditions, individual interventions draw on the practices outlined in focus areas 1, 2 and 5.

INDUSTRY INTERVENTIONS

Among internal audit professionals, this paper's main audience, only 14.9% report mature governance at their organizations, making focus area 3 a clear priority. Public accounting firms face direct risks as AI-assisted findings are the closest thing to a corporate reporting output; focus areas 2 and 4 would likely matter most here. In prior briefings we found corporate accounting and finance departments, on average, tend to be earlier on the curve and may benefit most from focus areas 2 and 3 (Redpath et al., 2026). These departments can channel their existing beliefs into safe exposures before governance maturity follows.

Professional accounting associations are tasked with standardizing what acceptable use and reviewed AI output actually mean across firms rather than leaving it open to interpretation. As such, education programs built around demonstrating AI's acceptable use cases would be more beneficial than programs built to demonstrate potential. Tech vendors should note that access to an AI tool predicted a small increase in willingness to present. Although the results were only marginally significant, they suggest that decisions about who gets first access to a tool can shape adoption conditions. Staff auditors with AI exposure and high belief and trust, such as the 71 Hidden Champions described on p. 36, are ready candidates for early pilots once they have explicit permission and a named reviewer.

Recommendations

ADOPTION ROADMAP

The roadmap places the five focus areas on a timeline, with one row for each domain (see Figure 21). Firms should develop objectives for each domain and monitor them separately to accommodate the finding that domains can mature at different rates.

First 90 Days:

Inventory current AI use within your area of jurisdiction. Document a governance baseline, even if it's rough, so there is something to measure against later. Choose a small number of low-risk pilots rather than a single firm-wide rollout. Communicate acceptable use clearly and formally in writing. Begin profiling teams using seniority and the persona framework.


3 to 12 Months:

Expand pilots that worked and be willing to end the ones that don't using formal checkpoints. Train managers directly on how to evaluate AI-supported judgments. Integrate GRC and internal audit functions around AI specifically. Implement persona-specific learning paths and continuously track domain maturity and incidents as ongoing metrics.

Beyond 12 Months:

Integrate AI into the actual audit workflow. Review whether interventions moved the needle on domain maturity. Revisit persona assignment and definitions for continued validation. Benchmark against other firms or professional groups.

Figure 21: Adoption Roadmap Across Time



Retrospective

WHAT WE LEARNED

Positive attitudes toward AI should not be mistaken for institutional readiness. Most respondents (72.6%) strongly believed AI can improve audit quality, but only 43.7% were strongly willing to present AI-assisted findings, and only 14.9% reported governance at Level 4 or higher (see Figure 10). Widespread employee use of AI does not, by itself, signal institutional maturity. Maturity requires evidence that use is authorized, controlled, reviewable and within a defined scope of assurance.

Greater governance maturity was associated with a lower chance of being in the blind spot. The adjusted probability fell from 40.9% at Level 1 to 22.2% at Level 5 (see Figure 13), though governance showed no association with willingness to present (see Figure 17). AI-tool users were in the blind spot at almost the same rate as non-users (29.1% compared with 31.4%), so tool presence does not guarantee audit coverage.

The blind spot is partly a visibility problem. More than two-thirds of those in it (68.2%) also rated at least one audit component 4 or 5 (see Figure 14). Respondents in more mature governance environments described their audit coverage more consistently, and the share who contradicted themselves fell from 16.8% at Level 1 to 11.6% at Level 4 (see Figure 16). Organizations should first determine whether AI audit activity exists, how complete it is and whether it is consistently visible to those using the technology.

For individuals, belief in AI's professional value was the strongest measured correlate of willingness to present (odds ratio 2.55, see Table 4). Trust under human review was also a moderate predictor (1.45), while AI-tool exposure had only a small, marginally significant association (1.16). Adoption programs that connect AI to clear improvements in evidence gathering, analysis, documentation or communication may therefore provide a stronger foundation than broad messaging about technological change.

Directors and C-Suite executives also reported greater willingness to present than staff auditors (odds ratios of 1.94 and 2.06, see Table 4), and the confidence gap narrowed with seniority (see Figure 12). The survey measured stated willingness, not decision rights, permission or perceived consequences, so how seniority matters remains an open question.

The segments and personas demonstrate how these conditions manifest in individuals. Toolled Adopters made up 45.4% of respondents, the Untooled Middle 29.9%, Governance-First 20.3% and Risk-Conflicted 4.4% (see Figure 18), so no single adoption program is likely to suit everyone. Belief in AI was already widespread. What lagged was the governance, visible assurance and confidence to present that turn belief into institutional readiness.

Retrospective

PROJECT LIMITATIONS

The sample is event-based and drawn from a single cross-section, so nothing in this paper should be interpreted as a prediction over time. The responses were reported by internal audit professionals attending a training event on ethics and emerging industry technologies. Attendance is self-selected and attendees could hold more favorable views of AI than the internal audit profession as a whole. Moreover, the survey instrument did not capture an employer identifier outside of email domains, so it is not possible to accurately account for clustering within organizations. The observed relationships are helpful for understanding the state of AI adoption in internal audit, but do not imply causation.

The action analysis employs 2,291 of the 3,089 retained responses, or 74%, while the final blind spot model retains 2,016, or 65%. Respondents who skipped any variable are therefore excluded from those estimates. No test of the missingness mechanism was performed, and the analysis does not presume that data are missing at random. It should also be noted that self-reporting captures perceptions that are subject to bias and that responses are not always consistent with how people behave.

PROJECT FUTURE

The current survey instrument measures the “belief that AI can improve audit quality” with a single item. Multiple items capturing different aspects of perceived value would improve understanding of outcomes. In the same way, “trust in AI results under human review” compresses reliance and verification behavior into a single item. Separating them would help distinguish whether someone trusts an AI output because a human reviewed it or because they trust the process. A multi-measure approach to each dimension would improve our understanding of the mechanics behind the observed relationships.

Beyond self-reporting, manager or peer observations could provide further validation of perceived realities. Analyses of governance and risk assurance could be validated by indexing policy documentation and process artifacts. Finally, repeated measures of the same professionals and firms over time would allow for direct longitudinal analysis of progress in each domain.

An analysis that combines self-report with manager or peer evidence, real-life testing and organizational artifacts would give a fuller and more defensible picture than any one of these methods can offer alone. Future surveys could also measure psychological safety, fear of reputational loss, permissions, AI competence, quality of use, client acceptance, regulatory pressure and actual presentation behavior.

Works Cited

- Banach, C., Boyer, R., Crous, K., Hedges, A., Langefels, K., McCarthy, B., Morabito, J., Peters, B., Redpath, E. (2026). The State of AI Adoption in Internal Audit in 2026. Aletheia Research Briefing 06/2026 (M. Levy, Ed.).
- Parasuraman, A. (2000). Technology Readiness Index (TRI): A Multiple-Item Scale to Measure Readiness to Embrace New Technologies. *Journal of Service Research*, 2(4), 307–320.
- Parasuraman, A., Colby, C. L. (2001). *Techno-ready marketing: How and why your customers adopt technology*. Free Press.
- Redpath, E., Hedges, A., Langefels, K., McCarthy, B., Morabito, J. (2026). The State of AI Adoption in Audit, Accounting, and Upstream Disciplines in 2026. Aletheia Research Briefing 05/2026 (I. J. Redpath, Ed.).
- Rogers, E. M. (1962). *Diffusion of Innovations*. New York: Free Press of Glencoe.
- Rogers, E. M. (1995). *Diffusion of Innovations* (4th ed.). New York: Free Press.
- Saeri, A. K., Lloyd George, S., Graham, J., Lacarriere, C. D., Slattery, P., Noetel, M., Thompson, N. (2025). Mapping AI Risk Mitigations: Evidence Scan and Preliminary AI Risk Mitigation Taxonomy. <https://doi.org/10.48550/arXiv.2512.11931>
- The Institute of Internal Auditors. (2020). The IIA's Three Lines Model: An update of the Three Lines of Defense. <https://www.theiia.org/en/content/position-papers/2020/the-iias-three-lines-model-an-update-of-the-three-lines-of-defense>
- The Institute of Internal Auditors. (2023). The IIA's Artificial Intelligence Auditing Framework. <https://www.theiia.org/en/content/tools/professional/2023/the-iias-updated-ai-auditing-framework/>
- The Institute of Internal Auditors. (2024). The IIA's Artificial Intelligence Auditing Framework. <https://www.theiia.org/globalassets/site/content/tools/professional/aiframework-sept-2024-update2.pdf>
- The Institute of Internal Auditors. (2026). Statement of position: Three Lines Model, assurance and advice in support of effective governance. <https://www.theiia.org/en/resources/statements-of-position/>

Disclosure Statement: This white paper was developed in collaboration with Aletheia's AUDIT Lab at Canisius University and received nonfinancial support from Cherry Hill Advisory. All opinions expressed are those of the authors and not those of the institutions they represent in their professional capacities. Generative AI tools were utilized to support exploratory research, data processing and transcription of interpretations. Authors reviewed, synthesized and revised all final content, performing targeted verification on key findings and robustness checks. The authors take full accountability for the published work.