BEYOND GOVERNANCE

MAKING AI USE VISIBLE IN SMALL AND MID-SIZED BUSINESSES

August 2026

Presented By

DataedX Group

Before Governance: Making AI Use Visible in Small and Mid-Sized Businesses

© 2026 by DataedX Group is licensed under CC BY-NC 4.0. To view a copy of this license, visit https://creativecommons.org/licenses/by-nc/4.0/

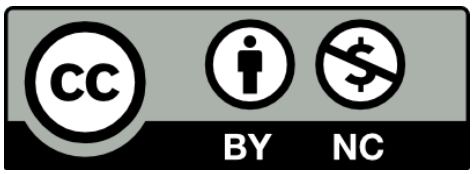

Suggested citation: Marshall, B.,(2026) 'Before Governance: Making AI Use Visible in Small and Mid-Sized Businesses', DataedX Group® Publication, Creative Commons BY-NC 4.0. Retrieved from https://dataedx.com/beyond-governance

Executive Summary

...whether lightweight discovery can provide reliable visibility into AI use, whether targeted controls improve oversight without creating disproportionate administrative burdens, ...

Small and mid-sized businesses (SMBs) cannot govern what they cannot see. As artificial intelligence (AI) enters everyday business operations, a practical starting point for responsible adoption is not necessarily more governance infrastructure. It is sufficient visibility into AI use to identify consequential applications, understand their business implications, and direct leadership attention where it matters most.

AI is entering SMBs through multiple channels: employees experimenting with generative AI, teams adopting new tools, and AI capabilities embedded in software businesses already use. These developments can improve productivity, decision-making, and service delivery. They can also introduce risks involving sensitive information, business decisions, customer relationships, employees, vendor dependencies, and operational continuity.

Leadership may recognize that AI is being used without having a reliable picture of where, how, by whom, or with what consequences. Approved-tool lists and written policies provide useful starting points, but they do not establish whether actual practices align with documented expectations. Nor do they necessarily reveal every AI-enabled activity on which the business depends.

This report presents an exploratory pilot involving nine SMBs that examined how organizations might develop a practical understanding of AI adoption and translate that understanding into proportionate governance decisions. The original approach focused on existing governance artifacts, including policies, procedures, vendor documentation, and meeting notes. The pilot encountered limitations in obtaining and interpreting these materials and in using them to establish how AI was being used in practice. This prompted a shift toward AI adoption visibility: the ability to identify AI use, understand its business context, assess its significance, and determine where further oversight may be warranted.

The revised approach incorporated three early-stage prototypes: an employee-facing decision aid, an AI tool disclosure mechanism, and a business-oriented AI use-case inventory. These prototypes were used to explore potential governance

needs and elicit participant feedback. They were not validated interventions, and the pilot did not establish their effectiveness.

The pilot supports further investigation of a practical approach to SMB AI governance built around four activities:

1. **Establish visibility:** Identify AI tools, embedded capabilities, experimental uses, and the business activities they support.
2. **Understand business context:** Examine the information, decisions, people, processes, and dependencies involved.
3. **Prioritize consequential uses:** Determine which applications warrant greater attention based on their potential consequences and operational significance.
4. **Apply proportionate controls:** Introduce targeted oversight that reflects the organization's risks, resources, and capacity

Key Findings and Insights

The pilot's documented design and methodological evolution point to five important areas for analysis and further validation:

- **Governance documentation does not establish implementation.** Policies and assigned responsibilities describe intended arrangements but cannot independently demonstrate understanding, consistent practice, or enforcement.
- **Employees may need practical decision support.** Accessible guidance may help employees interpret governance expectations in everyday situations that formal policies do not fully anticipate.
- **Disclosure is a discovery mechanism, not necessarily a complete inventory.** Employee reporting can surface AI use but may not capture every tool, embedded capability, or workflow.
- **Operational dependency deserves attention alongside adoption.** The significance of an AI use may depend on the business activity that relies on it, not simply the number of tools or frequency of use.
- **Business context makes AI inventories more actionable.** Understanding what AI supports and what could happen if it fails can help leaders determine where oversight is most valuable.

These insights should be interpreted in light of the pilot's exploratory design. Their empirical strength depends on the extent to which participant accounts and research records substantiate them.

Priority Recommendations

SMB leaders should consider three initial actions:

1. **Conduct a focused AI discovery exercise.** Work with department heads and relevant employees to identify AI tools, embedded features, experimental uses, and associated business activities.
2. **Assess material risks and dependencies.** Prioritize uses involving sensitive information, consequential decisions, customers, employees, critical processes, or significant reliance on AI-enabled services.
3. **Implement targeted governance controls.** Address material gaps through appropriate training, vendor reviews, human oversight, technical safeguards, policy updates, or escalation procedures.

The objective is not to impose enterprise-scale governance on every instance of AI experimentation. It is to develop enough understanding to make informed, proportionate decisions.

Scope and Contribution

This is an exploratory pilot, not a statistically representative study, a validated maturity model, or a definitive assessment of organizational readiness. It does not establish that the proposed prototypes improve governance outcomes or that the findings generalize across the broader SMB population.

Its contribution is to identify a practical direction for further investigation: whether lightweight discovery can provide reliable visibility into AI use, whether leaders can consistently prioritize consequential applications, and whether targeted controls improve oversight without creating disproportionate administrative burdens.

The Governance Challenge

Why SMBs Represent a Distinct AI Governance Challenge

Small and mid-sized businesses occupy a different position within the AI governance landscape from large corporations and government agencies. Many established governance frameworks, organizational maturity models, and

regulatory implementation approaches assume access to resources that smaller organizations may not possess.

These resources can include dedicated compliance personnel, formal information governance programs, enterprise risk management functions, specialized legal counsel, internal audit capabilities, AI governance committees, and documented procedures for technology procurement and use.

SMBs often operate with lean staffing models. Business owners, department managers, and IT personnel may fulfill multiple operational roles, leaving limited time and expertise for additional governance responsibilities. AI oversight therefore competes with other business priorities for attention and resources.

Applying enterprise-oriented frameworks without adapting them to these circumstances risks producing approaches that are comprehensive in theory but difficult to implement. Excessive administrative requirements may discourage engagement with governance altogether.

At the same time, AI tools can introduce meaningful business risks and opportunities. Relevant considerations include:

- Confidentiality of organizational information.
- Intellectual property protection.
- Customer privacy and data handling.
- Accuracy and reliability of AI-generated outputs.
- Human accountability for AI-assisted decisions.
- Transparency about AI use where appropriate.
- Dependence on AI-enabled vendors, software, and workflows.
- Operational continuity when AI systems become unavailable or unreliable.

The challenge is to develop governance that is practical enough for SMBs to use and sufficiently robust to address consequential applications.

This suggests a starting point that emphasizes organizational awareness: understanding where AI is being used, what business activities it supports, and where additional oversight may be warranted.

The AI Adoption Readiness Gap

Adopting responsible AI principles is not the same as demonstrating the ability to implement them. An organization may have an acceptable-use policy, vendor requirements, and stated commitments to responsible AI without having the

visibility, accountability, or operational processes needed to put those commitments into practice.

For SMBs, this distinction is particularly important because AI adoption may develop across departments, through employee experimentation, and within software products already used in daily operations. Formal governance may evolve more slowly.

Leadership may recognize the importance of responsible AI while lacking a reliable picture of where AI is being used, which business activities it influences, and which applications create meaningful risk or operational dependency.

The resulting gap is not simply a lack of policy. It concerns the organization's capacity to identify AI use, assign accountability, assess exposure, monitor outcomes, involve affected stakeholders where appropriate, and provide opportunities for review or correction when something goes wrong.

Existing governance documents can help establish what an organization expects to happen, who is nominally responsible, and which risks have been recognized. They are less conclusive about whether expectations are understood, followed, or enforced.

A policy does not establish that employees can interpret it in ambiguous situations. A vendor agreement does not, by itself, demonstrate that the organization understands the implications of the vendor's AI practices. An assigned owner does not necessarily have the information or authority needed to intervene.

Effective governance therefore requires more than formal commitments. It requires sufficient visibility into actual practices to connect AI use with business context, potential consequences, and decisions about oversight.

For SMBs, the practical challenge is to develop these capabilities without importing the administrative burden of an enterprise-scale governance program. Governance must be lightweight enough to use, specific enough to matter, and flexible enough to fit the business.

Purpose of the Pilot

The pilot began with the premise that examining existing policies, procedures, vendor requirements, and related governance materials might reveal meaningful information about organizational AI risk and inform proportionate interventions.

The initial objective was to determine whether a low-touch engagement could establish a useful baseline of governance practices and identify practical next steps without requiring a comprehensive organizational assessment.

The work evolved as limitations in the original approach became clearer. Governance materials were difficult to obtain, and their contents did not necessarily establish how AI was being used or whether documented expectations reflected everyday practices. The inquiry consequently expanded toward questions of AI adoption visibility, employee understanding, accountability, implementation barriers, and responsibility for vendor-related risks.

The revised pilot pursued three objectives:

1. **Examine governance artifacts.** Explore what policies, procedures, and related materials can establish about intended controls, recognized risks, and assigned responsibilities, and identify important questions they leave unanswered.
2. **Explore lightweight approaches to AI visibility.** Investigate whether focused conversations and practical prototypes can help identify AI tools and embedded capabilities, understand their business applications, and surface potentially consequential uses.
3. **Inform proportionate leadership action.** Explore how greater visibility into AI use, business exposure, and operational dependency might help SMB leaders prioritize risks and identify appropriate responses.

The pilot was designed to explore organizational conditions and possible intervention points, not to certify governance maturity or determine the effectiveness of a particular intervention.

Research Questions and Assessment Framework

Research Questions

The pilot was guided by three research questions:

1. What can existing governance materials reveal about AI oversight in participating SMBs, and what remains unknown?

2. How do participating SMBs understand and identify AI use across their business activities?
3. What lightweight approaches might help leaders prioritize consequential AI uses and determine appropriate oversight?

Together, these questions examine the relationship between governance design, implementation, AI adoption visibility, and organizational capacity.

The first considers what formal materials establish about intended arrangements. The second focuses on the visibility of actual AI use and its business context. The third explores potential interventions that could help organizations translate greater understanding into practical governance decisions.

These questions reflect the pilot's evolution from document-centered assessment toward a broader examination of AI adoption readiness.

Assessment Framework

The pilot used five analytical dimensions to organize evidence, examine governance capabilities, and identify gaps between documented expectations and operational practice.

Leadership, Accountability, and Institutional Responsibility. Examines ownership, decision-making authority, and accountability for AI-related risks and interventions.

Policies, Governance Artifacts, and Operational Implementation. Considers what documented governance arrangements establish and whether they reflect actual AI practices.

AI Discovery, Risk Assessment, and Oversight. Examines visibility into AI use, business context, material risks, and operational dependencies.

Transparency, Stakeholder Participation, and Contestability. Considers opportunities for affected stakeholders to raise concerns, challenge consequential AI-assisted outcomes, and seek review or correction.

Organizational Resources, Expertise, and Operational Capacity. Examines the capacity to implement and sustain governance proportionate to organizational needs and risks.

These dimensions are intended to structure inquiry rather than provide a validated scoring system. Their value lies in helping researchers and practitioners examine different aspects of governance without treating documentation, discovery, or organizational capacity as interchangeable.

Cross-Cutting Analytical Approach

AI adoption visibility cuts across all five dimensions rather than constituting an inventory exercise alone.

Leadership needs to recognize consequential uses and act on available information. Governance documentation must be considered in relation to actual practices. Risk assessment must connect AI use to business exposure and dependency. Affected stakeholders need appropriate opportunities to raise concerns and seek review. Organizations must also have sufficient capacity to sustain these activities.

The framework therefore distinguishes four related aspects of AI adoption readiness:

- **Governance design:** What policies, responsibilities, expectations, and controls are formally established?
- **Governance implementation:** To what extent are these arrangements understood, applied, and supported in practice?
- **AI adoption visibility:** Can the organization identify where AI is being used and understand the business activities, exposures, and dependencies involved?
- **Governance capacity:** Can the organization act on what it learns and sustain oversight proportionate to the significance of its AI use?

These distinctions help prevent the presence of policies, assigned roles, or tool inventories from being treated as sufficient evidence of effective governance.

They also provide a practical progression from discovery to prioritization and intervention. Rather than applying uniform controls to every instance of AI experimentation, SMB leaders can investigate which uses warrant greater attention and why.

Scope and Limitations of the Framework

The five dimensions constitute an exploratory analytical framework for structuring evidence collection, interpreting pilot findings, and identifying questions for further investigation. They are not a validated maturity model, a certification of compliance, or a definitive measure of organizational readiness.

The pilot does not establish that the framework provides a complete account of AI adoption or that the proposed discovery and assessment approaches will generalize across SMBs.

The distinction between observations, interpretations, and claims requiring further investigation is central to interpreting the results. Documented pilot activities and participant accounts can support conclusions about the experiences captured in the study. Broader claims about prevalence, effectiveness, or causation require additional evidence.

Further research should examine whether lightweight discovery methods reliably identify consequential AI use, whether leaders can consistently assess risks and dependencies, and whether targeted controls improve oversight without creating unnecessary administrative burdens.

Pilot Design and Methodology

Research Design and Study Rationale

The pilot initially sought to establish organizational baselines, assess governance maturity, identify gaps and risks, and examine the usefulness of an Explainability AI (XAI) assessment tool.

As the pilot progressed, a mismatch emerged between the planned maturity-assessment pathway and the governance arrangements participants described. The original approach depended in part on obtaining documentation and establishing sufficiently developed governance practices to support the intended assessment sequence.

Weresponded by redirecting the inquiry toward current AI use, organizational visibility, oversight limitations, implementation barriers, and practical interventions that participants might consider feasible.

The revised approach placed organizational experience ahead of a predetermined assessment sequence. It also introduced three early-stage prototypes to elicit feedback about possible ways of addressing the problems under investigation.

The resulting study combined structured organizational assessments, questionnaires, guided assessment sessions, semi-structured interviews, and prototype-informed feedback.

The design was exploratory and iterative. Initial activities provided a basis for understanding participating organizations, while subsequent engagement informed changes to the research emphasis and the types of evidence sought.

The prototypes were not treated as finished products evaluated against predetermined performance criteria. Instead, they served as research instruments for investigating organizational needs, potential intervention points, and perceived implementation conditions.

This approach allowed us to connect descriptions of organizational problems with participant responses to possible solutions while recognizing that stated preferences do not necessarily predict actual use or effectiveness.

Participant Cohort and Organizational Context

The pilot engaged nine SMBs representing a range of organizational circumstances.

Participating organizations ranged from businesses employing a handful of people to those with more than 100 employees. Reported annual revenues ranged from just under $1 million to more than $20 million.

Participants differed in their levels of AI adoption, dependence on AI-enabled tools, and existing governance practices. This variation provided an opportunity to explore governance needs across different organizational contexts.

In SMBs, AI-related responsibilities may arise before organizations have established dedicated roles, formal procedures, or consistent documentation for managing them. Oversight may depend on individual employee judgment, informal management practices, or arrangements distributed across several business functions.

These circumstances raise important questions about whether organizational leaders have a sufficiently complete account of the AI tools their employees use, how those tools influence business activities, and which applications warrant additional scrutiny.

The cohort's variation provided a useful context for exploratory inquiry, but it does not establish how prevalent these conditions are across the wider SMB population. Nor does the sample establish that any observed differences are representative of businesses of comparable size or revenue.

The cohort should therefore be understood as a contextually varied group through which we explored organizational experiences, potential governance challenges, and the relevance of practical interventions.

Study Structure and Data Collection

The pilot was organized into three light-touch engagements intended to accommodate the operational constraints of participating businesses while gathering information over time.

Engagement 1: Introducing the Pilot and Establishing the Baseline

The first engagement lasted approximately 60 minutes. It introduced the pilot's objectives, participation process, and anticipated activities.

The session covered the pre-qualification questionnaire, documentation requested for a baseline Explainability AI assessment, participant questions, and next steps. Participants were given approximately one week to follow up on the requested information.

This stage established the initial assessment process and clarified what participation would involve.

Engagement 2: Reviewing Baseline Findings and Exploring Governance Maturity

The second engagement lasted approximately 60 minutes. An approximately 30-day follow-up interval was planned thereafter.

The session reviewed baseline assessment outputs, discussed conditions associated with different governance maturity phases, addressed questions, and established next steps concerning governance processes or protocols.

At this stage, the study remained oriented toward the original assessment model. The intention was to understand participants' governance arrangements in relation to the available information and the proposed maturity-assessment approach.

This engagement also exposed a limitation in the planned pathway: participants' governance arrangements were not consistently developed or documented enough to support the intended progression as originally conceived.

Engagement 3: Investigating Foundational Governance Needs and Testing Prototypes

The third engagement lasted less than 60 minutes and was redesigned in response to the preceding engagement.

The revised session focused on current AI use, gaps in organizational visibility and oversight, barriers to responsible adoption, and practical support participants believed would be useful.

Three prototypes were introduced to explore potential responses to these needs. Participants were invited to provide feedback on their understandability, anticipated usefulness, organizational fit, missing features, and likely adoption barriers.

The engagement concluded with an invitation to a community webinar, providing an opportunity for continued engagement beyond the immediate pilot activities.

The revised sequence reflects the study's shift from a primarily maturity-oriented assessment toward investigation of the conditions that may support more systematic AI oversight.

Methodological Pivot: From Governance Maturity to Foundational Capabilities

The methodological pivot was a refinement of the research approach in response to field observations.

Rather than treating the absence of formal governance arrangements primarily as a maturity-assessment outcome, the study redirected its attention toward four areas:

AI visibility. How AI-enabled tools and applications were being used in business operations, including potential gaps between employee use and management awareness.

Governance practices and gaps. How participants described their existing approaches to oversight, documentation, accountability, transparency, and risk identification.

Organizational barriers. What constrained the development of governance practices, including limited time, resources, expertise, or clarity about responsibility.

Practical intervention needs. What forms of guidance or lightweight processes might help participants make more informed decisions without imposing disproportionate administrative requirements.

The pivot does not establish that governance maturity frameworks are unsuitable for SMBs. Nor does it demonstrate that every organization should follow the same sequence of governance development.

It identifies a limitation in applying the original assessment pathway to this particular cohort and underscores the need to understand organizational circumstances before assuming that a predefined assessment sequence will be appropriate.

The revised approach also distinguished between documented governance and governance as practiced. The absence of a written policy does not, by itself, establish that an organization has no oversight practices. Equally, the existence of a policy does not demonstrate that employees understand it or that it is consistently implemented.

Investigating these distinctions required attention to everyday practices, decision-making, and organizational responsibilities, rather than relying exclusively on the presence or absence of formal governance artifacts.

The methodological pivot therefore shifted the emphasis from classifying organizations against a maturity-oriented framework to investigating the foundational conditions and practical interventions that might support more systematic oversight.

Interview Design and Research Traceability

The revised third engagement used a semi-structured interview approach to maintain consistency across participants while allowing questions to be explored in relation to each organization's circumstances.

Research objectives were mapped to representative interview questions and anticipated research or toolkit outputs. This alignment was intended to establish traceability between the problem being investigated, the evidence sought, and the intervention being explored.

The interview framework addressed three overarching areas:

1. **Organizational AI practices:** How is AI currently used, and what does the organization know about that use?
2. **Governance conditions:** What gaps, risks, uncertainties, and implementation barriers affect AI oversight?

3. **Potential interventions:** What practical resources could support better visibility, more informed decisions, and more consistent oversight?

Prototype feedback formed a related but distinct component of the interviews. Participants were asked to consider the understandability, anticipated usefulness, organizational fit, missing features, simplicity, and adoption conditions associated with each prototype.

This structure was intended to help distinguish between identifying a governance need and determining whether a particular design might address it.

The interview matrix should be included as supplementary material where appropriate, showing the relationship between research objectives, representative questions, and anticipated outputs.

The questions presented in such a matrix should be identified as representative samples rather than the complete interview protocol unless the full protocol is provided.

Prototypes as Research Instruments

Three early-stage prototypes were developed to investigate complementary dimensions of organizational AI governance: employee decision-making, visibility into AI tools, and understanding of operational AI use.

Each prototype served as a research instrument for examining a design hypothesis about organizational needs and the perceived feasibility of potential governance support. None was treated as a validated intervention.

Prototype A: Organization Decision Aid

This prototype examined the perceived value of accessible support for navigating everyday AI-related decisions. The inquiry focused on whether such support could be understandable, relevant to workplace conditions, and feasible within organizations with limited governance resources.

The prototype addressed a practical question: can employees be given usable guidance to help them interpret governance expectations when deciding how to use AI? Its effectiveness in improving decisions was not evaluated.

Prototype B: AI Tool Disclosure

This prototype explored organizational approaches to increasing visibility into AI-enabled tools used across business functions. The inquiry considered the practical challenges of disclosure, including participation, administrative burden, and organizational responsibility.

The underlying premise was that employees may have information about AI use that is not captured in formal inventories or approved-tool lists. However, disclosure is a discovery signal rather than necessarily an authoritative record. The pilot did not establish the completeness or reliability of any resulting inventory.

Prototype C: AI Use-Case Inventory

This prototype examined the potential value of documenting how AI is incorporated into business operations. It addressed the distinction between knowing which tools are present and understanding their organizational applications.

For example, identifying a particular AI-enabled application does not necessarily reveal whether it is used for internal brainstorming, customer communications, processing sensitive information, or supporting consequential business decisions. These applications may warrant different levels of oversight even when they involve the same tool.

The inquiry considered the perceived usefulness and feasibility of structured visibility without assuming that documentation alone improves governance.

Common Feedback Protocol

Participants reviewed the three prototypes using a shared set of questions concerning comprehensibility, anticipated use, organizational value, perceived gaps, simplicity, and adoption conditions. This common approach supported comparative assessment while allowing participants to articulate their own perspectives.

Feedback informed the refinement of design assumptions; it was not treated as evidence of effectiveness. Collectively, the prototypes explored three interdependent dimensions of foundational AI governance capacity:

- Visibility into AI tools.
- Understanding of AI use in business operations.
- Informed employee and leadership decision-making.

Their contribution at this stage was to structure inquiry into potential governance needs and implementation conditions, establishing a basis for further research and evaluation.

Evidence, Interpretation, and Methodological Limitations

The analysis distinguishes between observations documented during the pilot, interpretations of those observations, and propositions requiring further investigation.

This distinction is important because the study was exploratory and was designed to investigate organizational conditions and potential interventions rather than establish definitive conclusions about SMB AI governance.

The documented research design establishes the characteristics of the participating cohort, the original and revised engagement structure, and the introduction of three early-stage prototypes. Participant accounts and assessment records provide additional evidence about reported AI practices, governance arrangements, implementation barriers, and prototype responses to the extent captured in the underlying records.

These materials may inform interpretations about the importance of visibility into actual AI use, business context, and operational dependency. The strength of each interpretation depends on the supporting evidence, including the consistency of participant accounts and the extent to which alternative explanations were considered.

The study does not establish the prevalence of particular governance practices across SMBs, the completeness of any resulting AI inventory, or a causal relationship between proposed interventions and organizational outcomes.

Further analysis should distinguish recurring observations from isolated comments, preserve differences across participants, and document how contradictory or negative feedback informs conclusions. Where systematic coding or other formal analytical procedures were performed, these should be described explicitly. Where they were not performed, the report should avoid implying otherwise.

The small cohort and formative design limit the kinds of conclusions that can be drawn. The study can help identify problems encountered by participants, illuminate possible mechanisms, and inform subsequent research. It cannot establish population-wide prevalence or demonstrate that a particular

governance intervention produces better outcomes than alternative approaches.

Pilot Findings and Emerging Insights

The pilot's contribution should be assessed through the evidence collected during the engagements, the limitations encountered in the original assessment approach, and the questions raised by the three prototypes.

The following findings structure the principal insights that the research was designed to investigate. Before publication, each should be linked to the relevant interview records, assessment materials, or documented research decisions. The available draft does not provide sufficient participant-level evidence to establish the prevalence or consistency of these observations across the nine organizations.

Governance Documentation and Operational Practice Are Distinct

The original assessment approach relied in part on governance artifacts to establish organizational expectations, recognized risks, and assigned responsibilities.

The pilot encountered difficulty obtaining and interpreting these materials in ways that supported the planned assessment sequence. This prompted us to broaden the inquiry beyond formal documentation.

The methodological implication is that documentary evidence and operational evidence answer different questions. Policies and procedures can establish what an organization intends to happen, but they cannot independently demonstrate that employees understand or follow those expectations.

Conversely, the absence of formal documentation does not establish that no oversight occurs. Informal decision-making, management review, and employee judgment may also influence how AI is used.

For SMB governance assessments, documentary evidence should therefore be considered alongside accounts of actual practices and the organization's ability to identify and respond to consequential AI use.

Implication: Assess governance documentation in relation to operational practice rather than treating the existence or absence of formal artifacts as a complete measure of readiness.

AI Adoption Visibility Is a Distinct Governance Need

The revised research approach focused on whether organizations could identify AI use across their operations and understand the activities that use supported.

This focus reflects the possibility that AI enters an organization through multiple channels, including employee experimentation, departmental tools, and embedded software capabilities. A centrally maintained list of approved tools may not capture all these activities.

The tool disclosure prototype explored one potential means of improving visibility. The use-case inventory explored a complementary approach that connects AI tools to business activities.

Neither approach was established as complete or authoritative through the pilot. Their value lies in the questions they help investigate: what AI is being used, where it is used, who relies on it, and what the organization knows about its implications.

Implication: AI discovery should be treated as an ongoing organizational activity rather than a one-time exercise that produces a permanently complete inventory.

Business Context Can Make AI Discovery More Actionable

Knowing which AI tools are present provides only part of the information needed for governance decisions.

A tool's significance depends partly on what it does in practice, the information it processes, the decisions it influences, and the business activities that rely on it.

An application used for low-consequence internal experimentation may warrant a different level of oversight from the same application used to process sensitive information or support consequential decisions.

The use-case inventory was designed to explore this distinction by connecting tools with their operational applications.

The pilot did not establish that a business-oriented inventory produces better outcomes than a tool list alone. It did, however, provide a basis for investigating whether business context can help leaders make more informed prioritization decisions.

Implication: Discovery should capture enough context to explain why an AI use matters, not simply whether a tool exists.

Operational Dependency May Be an Important Prioritization Factor

AI governance often begins with questions about which tools are used and how frequently they are used. These are useful considerations, but they may not fully reflect operational significance.

A business process can become dependent on an AI-enabled service even when the underlying tool is used infrequently. Conversely, frequent experimentation may have limited consequences if it remains isolated from sensitive information and important business decisions.

Operational dependency therefore merits attention alongside the nature of the task, the information involved, and the potential consequences of an error or disruption.

This is a proposition for further investigation, not a demonstrated comparative finding of the pilot. Additional research is needed to determine how SMBs can assess dependency consistently and how it should influence the prioritization of AI-related risks.

Implication: Assess what the business relies on AI to accomplish and what would happen if the capability became unavailable, unreliable, or inappropriate for use.

Practical Decision Support May Help Bridge Governance Intent and Practice

The organization decision aid explored the perceived value of accessible guidance for everyday AI-related decisions.

The underlying design question was whether employees could benefit from practical support that helps them interpret expectations in situations where a policy alone may not provide a sufficiently specific answer.

Such support could potentially help employees recognize sensitive information, identify situations requiring human review, or determine when to seek guidance. However, these are possible functions of a decision aid, not demonstrated outcomes of the prototype.

The pilot did not establish whether participants used the decision aid in practice, whether it improved decision quality, or whether it reduced AI-related risks.

Implication: Employee-facing guidance merits further evaluation as a complement to formal policies, with attention to understandability, practical use, and the conditions under which employees seek assistance.

Cross-Cutting Interpretation

These areas of inquiry support a practical hypothesis: SMB AI governance may benefit from a progression that begins with visibility, connects AI use to business context, prioritizes consequential applications, and applies controls proportionate to the identified exposure.

The pilot does not establish that this progression is superior to alternative approaches. Its value is that it offers a coherent direction for further research and practical experimentation.

Validation should examine whether organizations can use lightweight discovery methods to produce sufficiently reliable information, whether leaders can prioritize uses consistently, and whether the resulting controls improve oversight without imposing disproportionate administrative burdens.

Implications and Recommendations

A Practical Progression for SMB AI Governance

SMBs can begin investigating AI governance readiness without implementing a comprehensive governance program. The proposed approach consists of four activities: establishing visibility, understanding business context and dependency, prioritizing consequential uses and applying appropriate controls.

We suggest the starting point be identifying AI tools, embedded capabilities, experimental uses and the business activities they support. This is intended to engage relevant employees, including tech or operations leaders. The result becomes an initial baseline that serves as an independent accounting inventory of adopted tools and other AI-enabled software.

For consequential uses, SMBs can then document the capability's purpose, information inputs, affected stakeholders and operational dependencies. This examination gives better introspection into the information sensitivity, human

oversight, potential consequences of unreliable outputs and relevant vendor practices and limitations.

Next, SMBs should strongly consider how they evaluate AI uses according to information sensitivity, potential harm, affected stakeholders, operational dependency and their organization's ability to detect and correct errors. A documented leadership assessment provides a practical approach to prioritizing the consequential uses.

And lastly, it's important to address identified exposures early through targeted measures, including employee guidance, vendor due diligence, information-sharing restrictions, human review, technical safeguards, accountability, escalation procedures and policy updates. An effective method is to match controls to the use case and its consequences rather than applying uniform requirements across all AI activity.

Priority Recommendations for SMB Leaders

Recommendation 1: Conduct Focused AI Discovery. Establish an initial inventory of AI tools, embedded capabilities, and experimental uses through a short, leadership-led exercise. Document uncertainties and areas requiring further investigation.

Intended outcome: A practical baseline of organizational AI use and unresolved questions.

Recommendation 2: Assess Material Risks and Dependencies. Examine consequential uses in relation to information sensitivity, decisions supported, affected stakeholders, potential failure consequences, and operational reliance. Prioritize by potential harm and business significance, not usage frequency alone.

Intended outcome: A defensible basis for determining where additional oversight is warranted.

Recommendation 3: Implement Targeted Controls. Address identified gaps through proportionate employee guidance, vendor review, information restrictions, human oversight, technical safeguards, and defined accountability or escalation procedures.

Intended outcome: Stronger oversight of consequential AI use without imposing unnecessary burdens on low-consequence experimentation.

Concluding Remarks

SMBs need a practical way to understand where AI is entering their business, which uses create meaningful risk or operational dependency, and where leadership intervention is warranted.

This pilot explored the limitations of relying on governance artifacts alone and investigated three lightweight approaches to addressing the resulting visibility challenge. Its methodological evolution highlighted the importance of examining actual practices and organizational circumstances rather than assuming that formal governance documentation provides a complete picture of readiness.

The study does not establish that the proposed prototypes improve governance outcomes or that a visibility-first approach is universally superior to alternative frameworks. It does, however, provide a basis for further investigation into a practical sequence for AI oversight: identify AI use, understand its business significance, prioritize consequential applications, and implement controls proportionate to the risks and dependencies involved.

The central proposition is straightforward: effective SMB AI governance should begin with enough understanding to make informed decisions. The next step is to test whether lightweight discovery and targeted interventions can provide that understanding reliably, sustainably, and without imposing disproportionate administrative burdens.

Source Materials

1. Andrea Bettoni, Davide Matteri, Elias Montini, Bartłomiej Gładysz, Emanuele Carpanzano (2021). An AI adoption model for SMEs: a conceptual framework, *IFAC-PapersOnLine, 54*(1), 702–708. https://doi.org/10.1016/j.ifacol.2021.08.082
2. Jan Jöhnk, Malte Weißert, & Katrin Wyrtki. (2021). *Ready or Not, AI Comes—An Interview Study of Organizational AI Readiness Factors*. Business & Information Systems Engineering, 63(1), 5–20. https://doi.org/10.1007/s12599-020-00676-7
3. Jeanne McClure & Gregg Gerdau. (2026). Why AI Readiness Is an Organizational Learning Problem, Not a Technology Purchase. Preprint. arXiv:2604.16369. Available at https://arxiv.org/abs/2604.16369
4. Esther Sánchez, Reyes Calderón, Francisco Herrera (2025). Artificial intelligence adoption in SMEs: Survey based on TOE–DOI framework, primary methodology and challenges. Applied Sciences, 15(12), 6465. https://doi.org/10.3390/app15126465.
5. Emmanouil Papagiannidis, Ida Merete Enholm, Christian Dremel, Patrick Mikalef, & John Krogstie. (2023). *Toward AI Governance: Identifying Best Practices and Potential Barriers and Outcomes*. Information Systems Frontiers, 25, 123–141. https://doi.org/10.1007/s10796-022-10251-y.

6. Sukanlaya Sawang and Virach Sornlertlamvanich (2026). Artificial intelligence (AI) maturity in small and medium-sized enterprises: A framework of internalized and ecosystem-embedded capabilities (arXiv:2603.08728). arXiv. https://arxiv.org/abs/2603.08728
7. Thomas Übellacker (2025). Making sense of AI limitations: How individual perceptions shape organizational readiness for AI adoption (arXiv:2502.15870). arXiv. https://doi.org/10.48550/arXiv.2502.15870

Acknowledgements

This material is based on work supported in part by the Institute for Trustworthy AI in Law and Society (TRAILS), which is supported by the National Science Foundation under Award No. 2229885. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the National Institute of Standards and Technology.

Authors

Brandeis Marshall, Ph.D. is founder and CEO of DataedX Group®, a data and AI governance consulting agency. Under her "move slower and build better people-first solutions" leadership, cross-functional teams are guided and supported in improving their data strategy, responsible AI practices and overall AI orchestration efforts. She is also the author of Data Conscience: Algorithmic Siege on our Humanity (Wiley, 2022), co-editor of Mitigating Bias in Machine Learning (McGraw-Hill, 2024) and contributing author in The Black Agenda (Macmillan, 2022).

About This Publication

DataedX Group® is a for-profit organization working across data strategy, responsible AI, governance, and organizational capability. This publication is part of its public scholarship work examining how data and AI practices affect institutions, communities, and public decision-making.

The report's purpose is to contribute evidence and practical recommendations to public discussions about AI governance. Its findings reflect the scope and limitations of the pilot and do not constitute a comprehensive assessment of the entire ecosystem.

Contact: DataedX Group® · dataedx.com