LINKS GLOBAL
Communicating with Purpose. Powered by Insight.

THOUGHT LEADERSHIP SERIES

The AI Adoption Problem Nobody Is Talking About

Three quiet failures, why they keep happening, and the question to ask first

A Links Global thought leadership article by M. Sanchez, founder of Links Global, a communications firm working at the intersection of trust science, narrative design, and civic engagement. Part of the Links Global article series: AI and Technology. Drawn from 25 years of field practice and evidence-based research in strategic communication, public engagement, and institutional trust.

In Brief

The AI debate is stuck between evangelists and skeptics while a quieter problem does the real damage: tools adopted without assessment, without integration, and without governance. This article describes the three adoption failures that are actually happening, the pressures that drive them, and what responsible adoption looks like for organizations that cannot afford to get it wrong.

The conversation about AI adoption in organizations has a characteristic shape. On one side: the enthusiasts, who see transformative efficiency gains, competitive advantage, and the democratization of capabilities previously available only to resource-rich institutions. On the other: the skeptics, who catalog the risks, bias, job displacement, data privacy, algorithmic opacity, and the unaccountable concentration of power in the hands of a small number of technology companies.

Both conversations are real. Both are necessary. And both are systematically missing the adoption problem that is doing the most damage in the most organizations right now.

That problem is not primarily about whether to adopt AI. It is about the gap between the pace of AI adoption and the development of the organizational, cultural, and governance conditions that determine whether AI adoption produces its promised benefits or reproduces and amplifies the institutional failures it was supposed to solve.

Organizations are deploying AI tools into workflows, decision processes, and public-facing services without the foundational conditions that would make those deployments work. And the consequences (not the dramatic, headline-generating failures of AI systems gone catastrophically wrong, but the quieter, systematic patterns of AI tools that underperform, entrench existing inequities, erode professional judgment, and create institutional dependencies on tools that cannot be governed or held accountable) are accumulating faster than the field is able to recognize them.

The Three Adoption Failures That Are Actually Happening

Adoption without assessment. The most pervasive AI adoption failure is the deployment of AI tools without adequate prior assessment of whether the tool is appropriate for the specific organizational context, population, and decision domain in which it will operate.

AI tool vendors make claims about their products that are almost always true in aggregate and in the populations on which the tools were validated, and almost always incompletely true in the specific contexts where the tools are being deployed. A predictive analytics tool that accurately identifies high-risk individuals in a health system serving primarily insured, English-speaking, urban populations will not perform identically in a health system serving primarily uninsured, multilingual, rural populations. The performance degradation is often significant. It is also often invisible: the tool produces confident outputs regardless of the adequacy of its underlying data, and the organizations deploying it may have neither the technical capacity nor the evaluation framework to detect the gap.

Assessment before deployment means asking, in sufficient technical depth to get honest answers: what was the training data? What populations were represented in it, and in what proportions? Where has this tool been independently validated? What are the documented performance differences across demographic subgroups? What happens when the tool encounters inputs outside its training distribution? Organizations that cannot get adequate answers to these questions from vendors are not in a position to evaluate the tools being sold to them. And the organizations least equipped to conduct this assessment are typically the ones serving populations most at risk of harm when the tools fail. The failure is not rare. RAND researchers cite estimates that more than 80 percent of AI projects fail, twice the rate of technology projects that do not involve AI, and a 2025 MIT NANDA report found that about 95 percent of enterprise generative AI pilots produced no measurable profit-and-loss impact, though it is a non-peer-reviewed study that defines success narrowly. In both cases the explanation was largely organizational, not technical.

Adoption without integration. The second adoption failure is the deployment of AI tools without genuine integration into the organizational workflows, professional cultures, and accountability structures they are meant to support.

This failure takes two characteristic forms. The first is AI as overlay: the tool is deployed on top of existing workflows without changing the underlying processes, creating a system in which the AI output is one more input into an already-complex human decision process. In this configuration, the AI's actual effect on decisions is opaque: it is unclear whether, or how much, the AI output is influencing the final decision, making it impossible to hold either the AI or the human decision-maker accountable for specific outcomes.

The second form is AI as replacement: the tool takes over a decision domain previously managed by human judgment, without adequate audit mechanisms or human review processes. This configuration is efficient when the AI performs well. When it performs poorly (when it encounters the edge cases,

the novel situations, the populations inadequately represented in its training data) the failure is systemic rather than individual, and the professional judgment capacity that would have caught the error has been atrophied by disuse.

Neither overlay nor replacement is inherently wrong. Both can be appropriate in specific contexts with specific governance. What makes them failures is the absence of deliberate design: the failure to think through, before deployment, what the appropriate role of AI judgment is relative to human judgment, how the outputs will be audited, who is accountable when the system produces harmful outputs, and how the organization will detect and respond to performance degradation over time.

Adoption without governance. The third and most consequential adoption failure is the deployment of AI tools without adequate governance infrastructure: the policies, accountability mechanisms, audit processes, and oversight structures that determine whether AI systems operate within acceptable boundaries and can be corrected or removed when they don't.

The NIST AI Risk Management Framework describes what adequate governance looks like, and it is voluntary, which is part of the problem. AI governance in most organizations is currently either absent or nominal: a set of policies that exist on paper without the institutional capacity to enforce them, a committee that meets quarterly to discuss AI use without the technical expertise to evaluate what they're reviewing, or a vendor contract that specifies the tool's performance without adequate mechanisms for the organization to verify that performance independently.

The absence of governance does not reduce the organization's accountability for the outcomes its AI tools produce. It increases it, because the organization can no longer plausibly claim that the tool's failures were beyond its ability to anticipate or prevent. In our experience, courts, regulators, and the public are increasingly sophisticated about the difference between organizations that adopted AI thoughtfully and organizations that adopted AI opportunistically. The organizations in the second category are accumulating liability faster than they recognize.

Why Adoption Failure Keeps Happening

The three adoption failures above are well-documented. They have been described in academic literature, investigated by civil society organizations, litigated in courts, and examined in congressional hearings. And they keep happening, at scale, in organization after organization, sector after sector. Understanding why requires looking at the structural incentives that drive AI adoption decisions.

The vendor incentive. AI vendors are incentivized to sell tools, not to ensure appropriate adoption. Their business model rewards rapid deployment, not careful assessment. Vendor representatives are almost always technically knowledgeable about their tools' capabilities and much less forthcoming about their limitations. The organizations best equipped to negotiate appropriate assessment, integration, and governance requirements into AI procurement are organizations with substantial

internal technical expertise, which are typically the largest, best-resourced organizations. Smaller organizations, community-serving institutions, and public-sector entities, precisely the organizations whose AI adoption decisions most affect vulnerable populations, are typically least equipped to conduct this negotiation.

The efficiency pressure. Organizations facing resource constraints (which includes virtually every public health agency, community health center, municipal government, and nonprofit operating in the current funding environment) face genuine pressure to find efficiency gains wherever they exist. AI tools are marketed as efficiency tools, and they often deliver real efficiency gains in specific, measurable ways. The efficiency gains are visible, quantifiable, and immediate. The adoption risks are diffuse, probabilistic, and delayed. In resource-constrained organizations operating under short-term performance pressure, the rational adoption decision often diverges significantly from the responsible one.

The FOMO dynamic. The fear of being left behind (of competitors, peer organizations, or more technologically sophisticated actors gaining advantages that are difficult to close) is a powerful driver of AI adoption decisions that bypass adequate assessment and governance. This dynamic is most visible in competitive industries but is present, in adapted form, in public-sector and nonprofit contexts: municipalities that feel pressure to demonstrate technological sophistication to funders and constituents, public health agencies that see peer institutions deploying AI tools and feel the urgency to keep pace, organizations whose leadership has been exposed to compelling AI demonstrations and is impatient with the assessment and governance processes that responsible adoption requires.

The expertise gap. Most organizations that are adopting AI lack the internal expertise to assess what they are adopting. This is not a criticism: AI literacy at the level required for responsible organizational adoption is genuinely rare, genuinely specialized, and genuinely difficult to develop quickly. The expertise gap is structural, and it creates a systematic information asymmetry between AI vendors and the organizations procuring their tools that favors rapid adoption over responsible adoption.

What Responsible Adoption Actually Looks Like

Responsible AI adoption is not the absence of AI adoption. It is adoption that is preceded by adequate assessment, designed with genuine integration, and governed by accountability structures adequate to the stakes of the deployment.

The assessment phase should be non-negotiable before any AI deployment that affects decisions about people's lives. It should include independent evaluation of the tool's performance in populations comparable to those the organization serves, equity analysis of performance differences across demographic subgroups, and honest documentation of the conditions under which the tool is

expected to perform adequately and the conditions under which it is not. This assessment should be conducted with genuine technical expertise: either developed internally or contracted from sources that are genuinely independent from the vendor.

The integration design should specify, before deployment, what the appropriate relationship between AI outputs and human judgment is in the specific decision domain. Which decisions can AI outputs inform without human review? Which require human review of every AI output? Which are categorically inappropriate for AI involvement regardless of the tool's performance? These are not technical questions. They are governance and ethics questions that require professional judgment informed by understanding of the AI system's actual capabilities and limitations.

The governance infrastructure should include: a defined owner of AI system accountability within the organization; regular audit processes that test system performance against deployment specifications, with particular attention to performance for underserved or high-risk subpopulations; a clear escalation process for when AI system performance degrades or produces harmful outputs; and a defined process for modifying or removing AI systems that are not performing within acceptable parameters.

None of this is technically complex. All of it requires organizational commitment that most AI adoption conversations currently bypass in the rush to deployment.

The Adoption Question Organizations Need to Ask

Behind every AI adoption decision is a question that is rarely asked explicitly: what is this tool for? Not what can it do, vendors will tell you that, but what specific problem, in this specific organizational context, for these specific populations, does this tool address better than the alternatives?

When that question is asked seriously, with honest assessment of the tool's actual capabilities and limitations in context, many AI adoption decisions look different than they do when the question is "how do we implement this tool that is being widely adopted in our field?"

Some adoption decisions still proceed. The tool genuinely addresses a specific problem, its performance in the relevant population is adequately validated, the integration design is thoughtful, and the governance is adequate. These are the AI adoptions that produce their promised benefits.

Others do not proceed, or proceed with significant modifications. The tool's performance in the relevant population is inadequately validated. The decision domain is one where the consequences of AI error are too consequential to accept without human review that the adoption design did not include. The governance infrastructure does not exist and cannot be built within the deployment timeline. These are the adoptions that, when they proceed anyway, produce the failure patterns the field is increasingly documenting.

The organizations that ask the question seriously (that make AI adoption decisions based on honest assessment of fit rather than on enthusiasm, peer pressure, or the efficiency gains visible in the vendor's demonstration) are the ones whose AI deployments produce benefits rather than reproduce and amplify the institutional failures they were supposed to solve.

That is not a high bar. It is the minimum standard of responsible organizational leadership in the AI era.

Where This Fits

Adoption without assessment is the same mistake, in a different field, as communicating without a trust reading. The Links Global Trust Audit™ shows what a condition-first discipline looks like; the companion article on AI literacy describes the understanding that responsible adoption requires.

Schedule a conversation with Links Global: www.linksglobal.org

Sources

- Ryseff, J., De Bruhl, B. F., & Newberry, S. J. (2024). *The root causes of failure for artificial intelligence projects and how they can succeed* (RR-A2680-1). RAND Corporation. rand.org/pubs/research_reports/RRA2680-1.html
- Challapally, A., et al. (2025). *The GenAI divide: State of AI in business 2025*. MIT NANDA. Reported in Fortune, August 18, 2025. fortune.com
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). nist.gov

Ownership and use. © 2026 Links Global. All rights reserved. The Links Global Trust Audit™, Narrative Stewardship Infrastructure™, NSI™, Epistemic Re-Regulation Architecture™, and ERA™ are trademarks of Links Global. You are welcome to read this article, link to it, share it, and quote brief passages with attribution to Links Global and a link to this page. Reproducing it in whole or in substantial part, adapting or translating it, or using the three adoption failures and the adoption assessment approach in training, consulting, publications, or derivative works requires prior written permission from Links Global. Please cite as: Sanchez, M. (2026). *The AI Adoption Problem Nobody Is Talking About*. Links Global. www.linksglobal.org.