Beyond Incrementalism:

Navigating Organizational Change, Behavioral Inertia, and Human Fear in AI-Driven Innovation

K. Russell[1,4], J. Wilck[2,4], S. Gerard[4], J. Hoag[3,4], B. Scott[1,4], J. Harrison[4], A. Sohn[4], J. May[4,5]

[1]Barton College School of Business & Innovation, Barton College, Wilson, NC 27893, USA

[2]Bucknell University, Lewisburg, PA 17837, USA

[3]Ohio State University College of Engineering, Columbus, OH 43210, USA

[4]Barton College Data Analytics Consortium, Barton College, Wilson, NC 27893, USA

[5]Richard Bland College, South Prince George, VA 23805, USA

Correspondence: innovate@barton.edu

Abstract

Organizations now have access to unprecedented compute, generative AI tools, multi-agent workflows, and data analytics capabilities, yet most AI initiatives fail to produce measurable transformation. This paper draws on sixteen transcribed sessions of the Barton College Data Analytics Consortium (October 2025–June 2026), a group of educators and practitioners in higher education, data science consulting, and industrial engineering, read alongside change-management, behavioral, and AI-productivity literature. We argue that the primary bottleneck to AI-driven transformation is neither financial nor technical but organizational: the human behavior that resists change. We identify eight distinct sources of resistance grouped into three families (threat-based, principled, and experiential), describe a “fluency trap” in which polished AI output suppresses scrutiny, and show that data literacy is a frequently missing prerequisite for AI readiness. Case evidence contrasts the automation of existing tasks with genuine agentic redesign of work. We conclude with a three-step framework for adoption (fix the process first, commit to scaling rather than downsizing, and run low-stakes, high-visibility pilots) and argue that without

confronting behavioral inertia, organizations do not innovate; they simply produce, in one author's words, "a digital version of a crappy process."

Keywords: Organizational Change, AI Adoption, Behavioral Inertia, Psychological Safety, Data Literacy, Agentic Workflows

Introduction

The Paradox of AI Readiness

Controlled studies of new AI tools and processes report real gains: a 14% average productivity increase for customer-support agents using a generative AI assistant, rising to 34% for novice workers [1]; a 40% reduction in time spent on professional writing tasks [2]; and 55.8% faster completion of a programming task with an AI pair programmer [3]. Practitioners in the consortium describe developer productivity in still stronger terms, as "not 10% better, but 10x better." The tools are real and the results are real, for those who have successfully adopted them.

The paradox is the gap between availability and transformation. RAND reports estimates that more than 80% of AI projects fail, roughly twice the failure rate of other IT projects [4], and an MIT study of enterprise generative AI pilots found that about 95% produced no measurable impact on profit and loss [5]. Ken Russell's account of being brought in to accelerate AI adoption at a client organization, only to be told "whoa, whoa, whoa, we don't want you to go that fast," illustrates that the resistance is not about ignorance of the tools. It is organizational.

The Incrementalism Trap

Scott Gerard defines the failure mode precisely: "People either propose tiny incremental changes or moonshots—the middle ground is where most value is, and where most people can't even imagine things." Organizations default to either trivial improvements or unrealistic ambitions. The

productive middle, genuine transformation without fantasy, requires exactly the kind of organizational change capacity many institutions have not built. The warning is an old one: three decades ago, Hammer argued that organizations should not use technology to "pave the cow paths" of existing processes but should redesign the work itself [6].

> Don't blindly pave the cow paths. If you digitize a crappy process, you now have a digital crappy process.
>
> —Scott Gerard

The incrementalism trap is visible in successive technology waves. As Ken Russell observed in the May 2026 session, organizations have been moving from plateau to plateau, adopting LLMs as a wave, settling into a new plateau, then waiting for the next incline, without ever fundamentally restructuring how work is done.

Thesis and Scope

This paper argues that the primary bottleneck to AI-driven transformation is not financial or technical but organizational change and the human behavior that resists it. The paper draws on change-management theory, case evidence, and first-person accounts from educators and practitioners across higher education, data science consulting, and industrial engineering to offer a practical, human-centered path forward.

The paper is grounded in this journal's dual mandate to balance scientific reason with philosophical curiosity. The "P" in PhD, doctor of philosophy, is a reminder the authors return to throughout: technology without philosophical grounding produces efficiency without wisdom.

Approach and Sources

The primary source material is a set of sixteen recorded and transcribed sessions of the Barton College Data Analytics Consortium held between October 2025 and June 2026 [7]. Participants

included faculty in computer science, information systems, and industrial engineering, data science and AI consultants, and practitioners working with community and business organizations. Sessions were unstructured discussions of AI adoption in participants' own institutions and client organizations. One case (Case D) is drawn instead from a written account provided to the authors in October 2026, after the sessions concluded, by a senior enrollment administrator at a small college; the institution is not named.

The transcripts were analyzed in four stages. First, AI-assisted extraction identified candidate themes, cases, and quotations across all sessions. Second, the authors curated this material by hand, verifying each quotation against the transcript and discarding themes that did not recur. Third, the authors held a recorded writing session to agree on the argument and its structure. Fourth, section leads drafted the text from that session and the curated material, and the full manuscript was checked with an originality-detection service. Quotations attributed to named authors are their own words from the sessions, and quotations in Case D are from the written account described above; dates refer to the session in which a remark was made. The consortium's own production process is thus a small case study in the creator-to-curator shift discussed below: AI made the human work faster and richer, but at no point did AI write the paper.

The consortium material is practitioner testimony rather than a representative sample, and its cases are illustrative rather than generalizable. Where the argument rests on empirical claims beyond the participants' own experience, it is supported by the published literature cited throughout.

Literature Review and Theoretical Framework

Organizational Change Models in the GenAI Era

Standard change-management frameworks, such as Kotter's eight-step model for leading transformation [8] and the unfreeze–change–refreeze sequence attributed to Lewin [9], assume technology cycles measured in years. Even the three-step reading of Lewin is a later simplification of a more dynamic account of group equilibria [10]. Generative AI compresses those cycles to months. The frameworks remain conceptually valid but operationally inadequate for the current pace: an organization that has barely "refrozen" after one wave of tooling is already facing the next.

External support for this argument comes from Mark Adams, whose keynote at the June 2026 Info-Tech LIVE conference set out a series of laws of innovation, as reported to the consortium by Ken Russell [11]. One of these laws applies directly: technology is easy to buy, but a rigid culture is hard to change. Status quo bias leads an organization to reject change in order to protect itself, so leaders must build psychological safety and break down rigid "tech priesthoods" before innovation can take root.

Both mechanisms in this law are well established. Status quo bias, the disproportionate preference for the current state of affairs, has been documented in decision making for decades [12], and psychological safety, a shared belief that a team is safe for interpersonal risk taking, predicts the learning behavior on which experimentation depends [13].

The literature itself has a bias worth noting. As the group observed in November 2025, most frameworks address the technology perspective and ignore non-technical workers: "when all you have is a hammer, everything looks like a nail."

Distinct Psychologies of Resistance

The consortium's discussions identify eight distinct sources of resistance to AI. They fall into three families, each requiring a different organizational response.

Threat-based resistance arises when AI appears to endanger a person's standing, livelihood, or safety at work:

- **Identity threat:** Professional identity is tied to expertise that AI seems to commoditize. A 20-year veteran who is proud of their craft does not want to hear that a chatbot can approximate it.
- **Economic threat:** A concrete, rational concern about jobs, wages, bargaining power, or the entry-level career ladder. The question "will I lose my job?" is not irrational. The gains and the risks of AI are unevenly distributed: experimental studies find the largest productivity gains among less experienced workers [1,2], while employment of early-career workers in the most AI-exposed occupations has fallen relative to their less-exposed peers [14]. For workers who have invested years mastering a skill that AI now performs, the anger stage of the Kübler-Ross model [15] is not stubbornness; it is proportionate to the threat.
- **Fear of experimentation:** Cognitive resistance in structured business environments where failure carries consequences. If trying something new and failing is penalized, employees revert to safe, incremental behavior. This is a design problem, not a people problem.

Principled resistance rests on values rather than self-interest and cannot be addressed with demonstrations or assurances about job security:

- **Ethical and environmental hesitancy:** Concern about the environmental costs of AI, the conduct of the companies building it, surveillance, bias, or the loss of human judgment.

This is a principled position, not ignorance, and it is treated in more detail later in the paper.

- **Existential and safety concern:** Concern about loss of control, catastrophic misuse, or the long-term trajectory of the technology, a concern shared by many respected researchers and AI lab leaders.

Experiential and relational resistance comes from a person's history with technology and the norms of the people around them:

- **Trust and accuracy skepticism:** Distrust rooted in a specific bad experience, or in a high-stakes domain with low tolerance for error, such as law, medicine, or safety engineering.
- **Social signaling:** The peer group is skeptical, and skepticism is the local norm. The position is relational rather than deeply held.
- **Exhaustion:** People who have lived through multiple "revolutionary" technology rollouts that overpromised and underdelivered pattern-match this one as more of the same.

John Hoag introduced the Kübler-Ross stages of grief [15] (denial, anger, bargaining, depression, and acceptance) not as metaphor but as an observed pattern among his colleagues. He noted in March 2026 that he is the only instructor in his computer science and information systems domain who acknowledges AI; his colleagues remain in the first three stages. A small college that deployed AI agents across enrollment reports a similar split: staff embraced AI while faculty were reluctant and, in many cases, still are (Case D later in the paper).

> My wife teaches middle schoolers. They're natives to this. They have far better prospects than the high schoolers… the worst of which are the incoming assistant professors, who have found that absolutely all of the rules have changed.
>
> —John Hoag

The Fluency Trap

Joe Wilck introduced a specific epistemic hazard that cuts across all three families of resistance: the fluency trap. “The more coherent the AI output, the less likely people are to question it.” The observation is consistent with research showing that fluently processed information is judged more true and more credible [16], and that people over-rely on automated aids, attending less carefully to their outputs [17]. AI’s greatest capability, producing polished and fluent text, is simultaneously its greatest danger. It creates an accountability vacuum: when AI makes decisions that no one reviews, no one is held accountable for the outcome. The trap also cuts the other way. Those who already trust AI review its output too lightly, while skeptics in the trust and accuracy group may see a single fluent but wrong answer as confirmation of their distrust.

Scott Gerard adds a corrective to the standard governance narrative: AI governance has not outrun implementation. The causality runs the other direction; implementation is outrunning governance. Organizations that treat governance as a trailing concern are creating the exact accountability vacuum the fluency trap produces.

Data Literacy as the Hidden Prerequisite

Joe Wilck’s April 2026 counterpoint reframes what AI readiness actually requires. When Ken Russell asked whether the share of graduating students able to build an agent would rise toward 100%, Wilck said it was going down. His explanation: students whose first computer was a Chromebook during pandemic-era schooling often do not know basic file structure and cannot locate files they have just saved. AI readiness assumes a foundation many learners, and many workers, do not have.

This connects to the consortium’s broader argument: data analytics is not a “waning” discipline soon to be replaced by AI. It is the prerequisite. Ken Russell frames it with a taxonomy from a book he co-authored [18, p. 58] (data to information, to knowledge, to insight, to wisdom,

to discernment) that extends the classic data–information–knowledge–wisdom hierarchy [19,20] and maps the value chain AI can accelerate but cannot manufacture from nothing.

> The data are screaming, and it's time for us to listen. But do we have the right tools to hear?
>
> —Ken Russell

From Manual Work to Agentic Workflows

Why Simple Automation Fails

Scott Gerard's minimal definition of an agent, a system that perceives, decides, and acts, follows the standard textbook definition [21] and distinguishes genuine agentic transformation from marketing hype. By contrast, current usage of the term "agentic" implies much more, requiring agents to exploit large language models, learning, and other associated machinery. Measured against either definition, most "AI adoption" is neither agentic nor transformative: it is the automation of existing tasks within existing structures.

Scott Gerard's weekend experiment with spec-driven design using Claude Code illustrates the difference. He began with general ideas; the AI repeatedly asked clarifying questions ("what about this? what about that?"), requiring iterative co-design rather than simple instruction. The output was a running Python application with a front-end interface, something he had no prior experience building. That is using agentic tooling to effect transformation. Layering a chatbot on an unchanged workflow is not.

The "composite AI sandwich" (probabilistic AI, then a deterministic process core, then probabilistic AI) is Scott Gerard's architectural framing for systems that are genuinely robust. Neither purely generative nor purely classical AI systems are adequate on their own for enterprise deployment.

Case Studies in True Transformation

Case A: The lean agentic enterprise. At a Raleigh-Durham Startup Week session that Becky Scott attended, a company of three people (two founders and one employee) described operating almost entirely through hundreds of agents handling sales, marketing, lead generation, and training delivery, and reported nearly $2 million in revenue. The founders were non-technical. Their interface was Bland AI, a voice-first agent platform: they spoke to it, described what they wanted, and it executed.

Case B: Bespoke human–AI integration. Becky Scott's Asana and Claude API build began with friction: an existing plugin required manual approval of every action, too much for real use. In approximately one hour, using the Asana API directly, she built an integration in which Claude creates its own task section, writes and tracks tasks, flags items needing human review, and closes them when complete. The human becomes the overseer rather than the executor.

Case C: The plumber's bid. A friend of an AI startup founder, a plumber, needed to compile a complex project bid that would have taken a full weekend by hand. An agent transcribed what the plumber described, gathered market intelligence, and built the bid in hours—something he had no inkling it could do. The package was so thorough it included additional intelligence on the potential client that could help the plumber structure and win future business as well. The case anticipates the seeing-is-believing argument below: as Becky Scott put it, "Sometimes people can't understand what it can do for them, and how it can save them time until they see it." The demonstration changed his mental model in a way no description would have. However, the agents ran on a complex, multi-agent, multi-platform setup that the average person would struggle to build; the plumber could not easily have made this change on his own.

Case D: Efficiency without redesign. Not every agentic deployment is a transformation, and the most instructive case here is a partial one. A senior enrollment administrator at a small

college saw generative AI as a way to separate the institution from “the slow moving machine that is higher education” while gaining efficiencies that an underfunded, overextended staff badly needed. Staff buy-in was built from the bottom up: the administrator educated teams on what AI could do and worked with them to decide what needed to remain human and what could be delegated to agents. The college then deployed more than 17 AI agents across recruitment, admissions, and student success.

After the first recruitment cycle, the efficiencies were real, but the outcome the college ultimately cared about, students enrolling and paying tuition, showed virtually no improvement. There was no comparison group, so holding enrollment steady may itself have been the return; without the agents, the numbers might have declined. The administrator’s own diagnosis, however, was that something important had been overlooked: “When you delegate tasks and workflows to AI, you also have to redesign the roles of the employees who previously owned that work.” The agents gave staff substantial time back, but no one had defined what they should do with it. “Without new expectations, responsibilities, and KPIs, people naturally gravitate back toward what they know.”

The administrator’s second lesson concerned what must stay human. Choosing a college, like buying a house, is a major life decision, and “people make decisions emotionally and then justify them logically.” In the revised model, humans build the relationships while AI supports them by handling transactional work, identifying behavioral signals, and flagging when a student needs an intervention. As the administrator summarized it: “AI delivered the efficiencies we were looking for, but we failed to fully capitalize on them because we had not reimagined how our people should spend the time it gave back.”

The Philosophical Shift: From Presenter to Creator

A bold investment in AI capabilities can permit instructors to create and maintain better content, enable active learning, and foster better student outcomes.

AI capabilities can sever remaining ties to traditional textbooks, even their online immersive, auto-graded successors, by "adopting" and ingesting Open Educational Resource (OER) and authored content into an LLM, specifically Google's Gemini Notebook (formerly NotebookLM), whose subject-matter context is limited to the sources the instructor supplies [22]. That workflow can produce collateral materials including slides, problem sets, Jupyter notebooks, and even timed scripts that can be voiced by ElevenLabs or turned into podcast episodes.

Such instructional content is amenable to both asynchronous online and flipped classroom models, depends on both the quality of the OER and editing by the instructor, and can produce reusable, "evergreen" content. More importantly, subject taxonomy, ontology, student maturity and demographics need to be considered: the same content should not be expected to transfer unchanged, so instructors should plan for continuous improvement, revising (and correcting) materials each term.

The Philosophical Shift: From Creator to Curator

When the human moves from creator to editor, from the person who produces the artifact to the person who reviews, shapes, and approves it, professional identity is restructured. Ken Russell described this shift concretely: when the Descript editing tool fixed a mispronunciation in a recorded video segment (changing "pro-ballistic" to "probabilistic") without a reshoot, the shift became visible in real time to the 19-year-old student and originator of the clip. She was stunned (and thankful) at not having to attempt a reshoot.

Becky Scott described the authorship process for another consortium's paper in the same terms: "AI did not write this paper. AI just kind of took our transcription and gave us some talking

points…" The production of this paper, described under Approach and Sources, follows the same pattern.

The Human Element: Overcoming Fear of the Unknown

Creating a Culture of Safe Experimentation

If failure is penalized, employees revert to incremental change. The design challenge is creating environments in which trying and failing carries no lasting cost: internal sandboxes, pilot permissions, and no-blame retrospectives. These are the structural conditions for the psychological safety that team learning requires [13].

An underappreciated barrier is that the AI tools themselves can undermine safe experimentation by being too agreeable. In March 2026 the consortium discussed Claude's tendency to be "too affirming," and Scott Gerard raised the concern that AI is "too agreeable." Research on language models documents the same tendency, termed sycophancy, in which models trained on human feedback match their answers to the user's stated views [23]. An AI that validates everything is not a safe experimentation partner; it removes the honest feedback loops that learning requires. This connects directly to the fluency trap: polished output without pushback inhibits the experimentation mindset.

> Chaos and ambiguity and fear—they're always there… lurking in the corner with those three cousins is the fourth cousin: clarity. You just have to get to it, and that's the hard part, but it's always there.
>
> —Ken Russell

Seeing Is Believing

Demonstration beats description. When Ken Russell showed a colleague what AI could do with a course-review process ("Ron's mind was blown"), it changed his perception in a way no white

paper could. Becky Scott's own trajectory, from the inaugural All Things AI conference to building custom API integrations a year later, is itself a case study in how repeated exposure accumulates into fluency. The plumber's bid (Case C above) is the same principle in a commercial setting.

The consortium demonstrated the principle in May 2026, when Joe Wilck ran a live comparison of Claude and ChatGPT on a security analysis of his own public online footprint. The comparison was spontaneous, low-stakes, and immediately useful: the ideal conditions for a seeing-is-believing moment.

June 2026 produced the most organic version. Becky Scott presented AI-assisted data analysis to professional peers, and their response was to ask for her spreadsheet so they could replicate the approach. A peer demonstration led directly to adoption requests.

Ethical and Environmental Hesitancy: A Distinct Category

Not all resistance is fear-based. Becky Scott encountered community members who are "absolutely adamant" about not using AI because of its environmental costs and ethical concerns about the companies building it. This is not the same as fear of job loss or fear of experimentation, and treating it as irrational Luddism misses a legitimate concern. Organizations that respond to it as they would to fear, with demonstrations or job-security assurances, will not move these people, and should not necessarily try.

This paper therefore treats principled hesitancy as a category of its own rather than folding it into the fear framework. Principled hesitancy is a legitimate position, and organizations that ignore it create the kind of forced-compliance environment that breeds more resistance and, later, outright rejection of AI.

AI Literacy as the New Compliance Training

Ken Russell observed in April 2026 that large companies are beginning to require employees to build a basic agent as part of standard compliance training, alongside harassment and cybersecurity modules. The baseline expectation is shifting from "use AI" to "build with AI."

Higher education is already running the natural experiment that compliance training is about to require at scale. John Hoag's classroom experiments provide the educational parallel in two contrasting approaches:

- **Software courses:** AI is banned entirely, preserving foundational skill-building. The rationale is not anti-AI sentiment but the recognition that students who cannot write code without AI assistance have not built the debugging instinct that makes them useful when AI fails.
- **Cyber courses:** "Impossible tasks on tight deadlines" are assignments designed so that human capacity alone cannot complete them, and students must use AI. Students who lean into the tools produce strong work. In an AI degree program, in which Hoag teaches one of six sections, he judges the student output to be "nearly 5x better" than that of the prior year's cohort. This is an instructor's assessment rather than a controlled measurement.

A Framework for Meaningful Organizational AI Adoption

The following three-step framework draws on both organizational change theory and the consortium's direct experience. It is informed by Ken Russell's Transact–Transform–Transcend model [18], which Scott Gerard applied to AI adoption in the June 2026 session, and which maps organizational ambition against readiness:

- **Transact:** doing today's tasks faster; the incrementalism trap.

- **Transform:** rethinking how work is done; the goal of this framework.
- **Transcend:** breaking existing paradigms and starting from scratch; the horizon.

Step 1. Strategic Alignment: Fix the Process First

Before introducing AI, audit the process it will touch. Organizations that skip this step automate their dysfunction [6]. The test is simple: if you would not be proud of the manual process, you should not automate it.

Strategic alignment also requires matching tool access to genuine need. Becky Scott observed in June 2026 that organizations with universal AI access were discovering employees using enterprise LLMs to run ordinary web searches, a misalignment between tool capability and task requirement that inflates cost without producing value. The lesson is structure, not restriction.

Another of Mark Adams' laws provides an external anchor: organizations should practice pre-commitment, setting boundaries in advance and limiting unnecessary choices, to protect the speed of execution and avoid "innovation tourism," endless testing without shipping [11].

Step 2. Psychological Safety: Scale, Don't Downsize

An explicit organizational commitment that automation will expand capacity rather than reduce headcount is not just ethically appropriate; it is strategically necessary. Shadow AI is the symptom of insufficient psychological safety: employees using AI clandestinely because the official environment is too risky or restrictive. A 2024 survey found that 75% of knowledge workers already use AI at work and that 78% of those users bring their own tools [24].

> Ken Russell's approach to organizational change is instructive here:
>
> When people bring me in to make change, I could easily come in and bowl people over, but that's not how I do it… we take a pragmatic approach, a softer approach, a listening-first approach, and that's how you make a dent.

—Ken Russell

In higher education, Becky Scott observed, faculty are rarely terminated. This makes it an ideal environment in which to make the psychological safety commitment credibly: the threat of job loss is low, but the promise of capability expansion is real. The question to reframe for employees is not "will AI replace me?" but "what will you do with the time AI gives back?" Case D shows what happens when that question goes unanswered: the agents returned time to staff, but without new roles, responsibilities, and KPIs to absorb it, people drifted back to familiar work and enrollment did not move. Redesigning the role is part of the commitment, not an afterthought.

Step 3. Low-Stakes, High-Visibility Pilot Projects

The ideal pilot is visible enough to shift perception but low-risk enough to survive failure. The Descript mispronunciation fix described above is a model case: a real problem, a visible result, no meaningful risk of failure, and an emotional impact (the student was stunned, and thankful) proportionate to the surprise.

Becky Scott's June 2026 Excel and ChatGPT data project is the organizational equivalent. She used tools already available (enterprise ChatGPT access), data already collected (community platform reports), and skills already in place (Excel familiarity). In one afternoon the result was ten new formula columns, pivot tables, health scores, 30-day and six-month trend data, and a curated prospect list for the sales team. Nothing new was acquired; everything was recombined. The work now runs as a weekly report with automated error checking, surfacing insights that had been waiting to be found.

Joe Wilck's counterpoint on data literacy warrants explicit acknowledgment in the framework: some populations need foundational scaffolding before a pilot can begin. A staged progression (prompt, then workflow, then agent, then multi-agent system) acknowledges that not

all users enter at the same point. The goal is not a single on-ramp but a path with multiple entry points.

Conclusion

AI adoption is a human problem that presents as a technical one. Organizations that treat it as a technical problem will continue to produce digital versions of broken processes. Even a well-executed deployment can stall if the roles around it are left unchanged: efficiency gained is not the same as value captured. Organizations that treat it as a human problem, one of culture, fear, literacy, and trust, will unlock the transformation that the tools already make possible.

Data analytics is not a waning discipline being absorbed by AI. It is the epistemological foundation that AI requires and amplifies. Without data literacy, AI produces confusion at scale. The consortium's recurring frame, "the data are screaming, and it's time for us to listen," is not a call for more technology. It is a call for better listening, which requires both the right tools and the human judgment to interpret what the tools surface.

Joe Wilck's observation captures the moment: "We no longer know what to teach a 5-year-old for 20 years from now." This is not a counsel of despair. It recognizes that adaptive capacity, the ability to learn new tools, ask new questions, and tolerate productive uncertainty, is now the primary skill. As Wilck also put it, "We've been augmenting since the beginning." From Taylor's stopwatch [25] to the printing press [26] to the spreadsheet, each wave of augmentation produced fear proportionate to the disruption and was eventually absorbed into what we simply call work. The current wave is not unprecedented in kind, only in scale and speed. The frameworks in this paper are not designed to produce AI experts but organizations capable of continuous adaptation.

That's a good reminder that everything that's wonderful is not necessarily technology.

—Ken Russell

The philosophical invitation, the “P” in PhD, is to remain curious about what we are doing and why, even as the tools change beneath us: to ask not only what AI can do but what we want it to do, and what we want to keep for ourselves. This paper does not celebrate AI adoption. It offers organizations a human-centered path through a transition that will happen with or without their intentional participation, and argues that intentional participation, grounded in organizational honesty about fear and culture, produces better outcomes for everyone.

Acknowledgments

The authors thank the participants in the Barton College Data Analytics Consortium sessions, October 2025–June 2026.

Use of AI Tools

AI tools were used to transcribe consortium sessions, to extract candidate themes and quotations from the transcripts, and to check formatting and references. All quotations were verified against the transcripts, and the argument and text were written by the authors, who take full responsibility for the content.

References

[1] Brynjolfsson, E.; Li, D.; Raymond, L. Generative AI at Work. *Quarterly Journal of Economics* **2025**, 140, 889–942. https://doi.org/10.1093/qje/qjae044

[2] Noy, S.; Zhang, W. Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science* **2023**, 381, 187–192. https://doi.org/10.1126/science.adh2586

[3] Peng, S.; Kalliamvakou, E.; Cihon, P.; Demirer, M. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. *arXiv* **2023**, 2302.06590. https://arxiv.org/abs/2302.06590

[4] Ryseff, J.; De Bruhl, B. F.; Newberry, S. J. *The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed: Avoiding the Anti-Patterns of AI*; Report RR-A2680-1; RAND Corporation: Santa Monica, CA, **2024**. https://doi.org/10.7249/RRA2680-1

[5] Challapally, A.; et al. *The GenAI Divide: State of AI in Business **2025***; MIT NANDA, Massachusetts Institute of Technology: Cambridge, MA, **2025**. Summarized in *Fortune*, Aug. 18, 2025. https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/

[6] Hammer, M. Reengineering Work: Don't Automate, Obliterate. *Harvard Business Review* **1990**, 68, 104–112.

[7] Barton College Data Analytics Consortium. *Consortium Session Transcripts, October **2025**–June 2026*; Unpublished records, Barton College: Wilson, NC, **2026**.

[8] Kotter, J. P. Leading Change: Why Transformation Efforts Fail. *Harvard Business Review* **1995**, 73, 59–67.

[9] Lewin, K. Frontiers in Group Dynamics: Concept, Method and Reality in Social Science; Social Equilibria and Social Change. *Human Relations* **1947**, 1, 5–41. https://doi.org/10.1177/001872674700100103

[10] Cummings, S.; Bridgman, T.; Brown, K. G. Unfreezing Change as Three Steps: Rethinking Kurt Lewin's Legacy for Change Management. *Human Relations* **2016**, 69, 33–60. https://doi.org/10.1177/0018726715577707

[11] Adams, M. Keynote Address on AI and Innovation. Presented at *Info-Tech LIVE **2026***, Las Vegas, NV, June 9–11, **2026**. As reported to the consortium by K. Russell.

[12] Samuelson, W.; Zeckhauser, R. Status Quo Bias in Decision Making. *Journal of Risk and Uncertainty* **1988**, 1, 7–59. https://doi.org/10.1007/BF00055564

[13] Edmondson, A. Psychological Safety and Learning Behavior in Work Teams. *Administrative Science Quarterly* **1999**, 44, 350–383. https://doi.org/10.2307/2666999

[14] Brynjolfsson, E.; Chandar, B.; Chen, R. *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*; Working Paper; Stanford Digital Economy Lab: Stanford, CA, **2025** (revised 2026).

https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/

[15] Kübler-Ross, E. *On Death and Dying*; Macmillan: New York, **1969**.

[16] Alter, A. L.; Oppenheimer, D. M. Uniting the Tribes of Fluency to Form a Metacognitive Nation. *Personality and Social Psychology Review* **2009**, 13, 219–235. https://doi.org/10.1177/1088868309341564

[17] Parasuraman, R.; Manzey, D. H. Complacency and Bias in Human Use of Automation: An Attentional Integration. *Human Factors* **2010**, 52, 381–410. https://doi.org/10.1177/0018720810376055

[18] Russell, K. E.; Ruppel, J. *Transact Transform Transcend: Becoming a Thoughtful Leader: Stories from the Journey*; Spencer House Publishing, **2024**.

[19] Ackoff, R. L. From Data to Wisdom. *Journal of Applied Systems Analysis* **1989**, 16, 3–9.

[20] Rowley, J. The Wisdom Hierarchy: Representations of the DIKW Hierarchy. *Journal of Information Science* **2007**, 33, 163–180. https://doi.org/10.1177/0165551506070706

[21] Russell, S.; Norvig, P. *Artificial Intelligence: A Modern Approach*, 4th ed.; Pearson: Hoboken, NJ, **2021**.

[22] Thompson, N.; Johnson, S. Writing in the Age of AI: Gemini Notebook's Steven Johnson with Nicholas Thompson. *The Most Interesting Thing in AI* [Podcast], August 12, **2026**. https://open.spotify.com/episode/0yUEZQKDqeDJBJGWrCHOUW

[23] Sharma, M.; et al. Towards Understanding Sycophancy in Language Models. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*, Vienna, Austria, **2024**. https://arxiv.org/abs/2310.13548

[24] Microsoft; LinkedIn. ***2024*** *Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part*; Microsoft Corporation: Redmond, WA, **2024**. https://news.microsoft.com/2024/05/08/microsoft-and-linkedin-release-the-2024-work-trend-index-on-the-state-of-ai-at-work/

[25] Taylor, F. W. *The Principles of Scientific Management*; Harper & Brothers: New York, **1911**.

[26] Eisenstein, E. L. *The Printing Press as an Agent of Change*; Cambridge University Press: Cambridge, UK, **1979**.