

Early AI adoption: How can you navigate time and trust tensions?

Is AI a sprint you're late for, or a long hike you haven't packed for?

ANNIKA ENGSTRÖM · DANIEL PITTINO · ALICE MOHLIN · NINA EDH · ANETTE JOHANSSON

Published on October 13, 2025

Leaders everywhere face a puzzling mix of signals. On one hand, competitors are touting AI wins, vendors promise quick ROI, and boards ask for timelines. On the other, data is messy, processes are fragile, and teams are unsure where AI will help (or harm) most. The result is a familiar organizational stall: pilots that never scale, sporadic automation, and a creeping sense of “we’re doing *something*, but are we actually moving?”

Recent evidence highlights why this stall happens. Early AI adoption reliably generates *two intertwined sets of tensions*: temporal (today’s reality vs. tomorrow’s ambition) and relational (human judgment vs. algorithmic autonomy). If you don’t name, and work with, these tensions, they silently reshape priorities, slow momentum, and drain trust. The good news: once visible, they can be steered.

WHAT WE STUDIED

This article distills insights from a qualitative study in large industrial settings undergoing early AI adoption. Researchers ran 23 focus groups with more than a hundred white-collar employees across functions. Their analysis surfaced nine recurring themes that cluster into two overarching tension types: (1) temporal (current capabilities vs. imagined futures) and (2) relational (humans vs. machines, trust vs. control, augmentation vs. replacement). We translate those research-backed themes into an action playbook you can put to work immediately.

KEY INSIGHTS

1) *The “Now vs. Next” Paradox Is Normal, Plan for Emergence, Not Certainty*

Organizations depict AI as both *a long road* (data quality isn’t there; processes aren’t ready) and *a sink-or-swim imperative* (we must “jump aboard now”). Many also treat AI as *familiar tech* (just more digital tools) and as *uncharted territory* (a paradigm shift). These simultaneous truths create an execution trap: leaders oscillate between incremental tweaks and grand visions, doing neither well. The research calls this an *emergent trajectory*: the path is non-linear, negotiated, and unfinished. Your operating system for AI must assume change and learning as defaults, not as exceptions.

Practical read: Stop searching for a master plan that survives first contact. Build a cadence that iterates strategy from evidence—quarterly, not annually.

2) *“Trust vs. Control” Is the Load-Bearing Wall*

Teams want explainability and guardrails, yet they also want speed and scale. People are more willing to use a system when they understand *how* it reasons, *where* it is strong/weak, and *who* is accountable when it fails. The research shows adoption accelerates when organizations replace vague assurance (“the model is good”) with concrete oversight rituals (e.g., human-in-the-loop checkpoints tied to risk levels) and when they clarify *when humans overrule vs. when machines lead*.

Practical read: Trust is not a feeling; it’s an architecture—policies, dashboards, thresholds, and escalation paths that everyone can see.

3) People See AI as Both “Helpful Assistant” and “Superhuman Rival”

Participants described AI as a way to offload drudgery and as a force that could outpace them. That ambivalence isn't resistance—it's information. If the only story is “AI frees you to be strategic,” but workflows still measure output like before, people will understandably keep their distance. Real adoption follows when you change the job to *fit the tool* (e.g., reweight performance metrics, redesign decision rights, and upskill toward the analysis/exception layers of the process).

Practical read: Role redesign must land **with** the tool, not a year later. Otherwise, augmentation looks like extra work.

4) Data Readiness Is Necessary—but It's Not the Bottleneck You Think

Yes, poor data quality slows progress. But the deeper blockers are *sensemaking and governance*: when to use AI, what risks are acceptable, and how to arbitrate disagreements between model outputs and human judgment. Teams need lightweight but explicit rules-of-engagement that can evolve. Without them, you get “pilot sprawl” and passive resistance (“we'll grow into it...someday”).

Practical read: Treat governance as a product: versioned, visible, and progressively hardened as you learn.

TAKEAWAYS

1) Build a Two-Speed AI Operating Rhythm

Why: You need breathing room for exploration and pressure to deliver.

How:

- **Quarterly “Now/Next” Reviews.** Every unit maintains a one-page map: *Now* (running AI use cases with owners, SLAs, and value realized) and *Next* (2–3 candidate experiments gated by a risk/impact rubric). Decisions: promote, pivot, or pause—no “someday” list.
- **Risk-Tiering by Use Case.** Classify by impact of error (e.g., Tier 1: financial or safety critical; Tier 3: advisory/low impact). Tie each tier to mandated controls (model cards, human checkpoint frequency, rollback plans).
- **Metric Duality.** Track *time-to-first-value* (speed) and *defect escape rate* (safety) so neither dominates.

2) Make Trust Operational, Not Aspirational

Why: Adoption rises when oversight is tangible.

How:

- **Decision Journals for AI.** For each automated decision type, log *input range, confidence, exceptions,*

and *final decision owner*. This creates an audit trail people can learn from.

- **Red Team Rituals.** Quarterly sessions where a cross-functional group stress-tests top models by probing edge cases and calibration drift. Publish findings and fixes.
- **Explainability by Audience.** Short “why it said that” summaries for front-line users; deeper model cards for specialists; board-level heatmaps of residual risk.

3) Redesign Roles Where the Work Actually Changes

Why: If work design lags model deployment, you get shadow processes and rework.

How:

- **Exception Leadership.** Shift skilled staff from routine processing to handling flagged edge cases, continuous prompt/process tuning, and customer-facing judgment calls.
- **Metric Realignment.** Move performance metrics from volume-based (tickets closed) to value-based (exceptions resolved, outcomes improved, defects avoided).
- **Career Ladders for AI-Infused Work.** Define paths like “AI Workflow Lead” or “Model Steward,” with competencies in data literacy, risk framing, and stakeholder communication.

4) Bring Ethical Reflexivity Into Regular Meetings

Why: Values discussions often stay abstract; product teams need concrete prompts.

How:

- **10-Question Ethics Checklist** (use at kick-off and at each major model update): Who is affected? Failure modes? Detectable bias? Human override? Appeal path? Data provenance? Retention? Monitoring plan? Energy/cost footprint? Communication plan?
- **Open Decision Memos.** When you choose between accuracy vs. transparency, or cost vs. control, publish a short memo of tradeoffs. Visibility builds credibility.
- **Community Review Windows.** Before scaling a high-impact use case, invite stakeholders to a time-boxed comment period with clear criteria for what feedback will change.

5) Start Where Paradox Is Manageable

Why: Early wins should teach you how to balance tensions, not just rack up ROI.

How:

- Pick use cases that are valuable but bounded (e.g., demand-forecast assist, invoice triage,

maintenance prioritization).

- Insist on “human-in-the-loop until proven otherwise,” and set *graduation criteria* for when the loop can loosen (e.g., sustained calibration performance, stable drift metrics, satisfied exception rates).
- Publish *retrospectives* that focus on how trust/controls evolved, not just the numbers.

IMPACT

Handled well, early AI adoption becomes a training ground for an organization's next decade: faster learning cycles, clearer accountability, and roles that emphasize creativity, empathy, and high-judgment choices. Handled poorly, it becomes a treadmill—ever more pilots, dashboard clutter, and quiet cynicism.

The research record suggests this isn't about choosing *speed* or *safety*, *automation* or *augmentation*. The durable advantage goes to organizations that can *hold the tension*: they treat AI as essential yet experimental; powerful yet bounded; helpful yet accountable to human values. That mindset, paired with the concrete practices in this article, reduces strategic drift and replaces it with compounding capability.

RECOMMENDATIONS

1. **Adopt an Emergent Strategy Cadence.** Establish quarterly “Now/Next” reviews, risk-tiering, and

dual metrics across all AI work. Make promotion/retirement decisions visible, and kill zombie pilots ruthlessly.

2. **Productize Governance.** Ship a living AI policy with version numbers, model cards, approval gates by risk tier, and rollback playbooks. Track policy adoption rates just like product usage.
3. **Engineer Trust.** Implement decision journals, red-team rituals, and explainability artifacts tailored to each audience. Tie funding for scale-up to evidence of these controls in action.
4. **Redesign Work with the Tool.** Update decision rights, performance metrics, and job architectures when models go live—not months later. Create pathways for “AI Workflow Lead” and “Model Steward” roles to retain and grow critical talent.
5. **Institutionalize Ethical Reflexivity.** Use a short ethics checklist at every milestone, publish tradeoff memos, and run community review windows before scaling high-impact use cases. Normalize the conversation about values as a standard engineering activity.
6. **Sequence for Learning.** Start with mid-stakes, high-signal use cases that teach the organization how to balance trust and control. Require “graduation criteria” before reducing human oversight.

Spotlight by CeFEO, *Early AI adoption: How can you navigate time and trust tensions?* Is AI a sprint you're late for, or a long hike you haven't packed for?. Downloaded on 28 July 2026 from <https://spotlight.cefeo.se>

CEFEO AUTHORS



Daniel Pittino

Professor

Jönköping International Business School

daniel.pittino@ju.se

PUBLISHED IN



Engström, A., Pittino, D., Mohlin, A., Edh, N., & Johansson, A. (2025). A paradox perspective on early AI adoption: Understanding temporal and relational tensions. *Journal of Organizational Change Management*.

<https://doi.org/10.1108/JOCM-02-2025-0086>

Spotlight is an online magazine that translates research from the Centre for Family Entrepreneurship and Ownership (CeFEO) at Jönköping International Business School, Jönköping University, into accessible insights for family business owners, practitioners, and policymakers. *Spotlight* is supported by the WIFU Foundation. This partnership advances dialogue and education in responsible family entrepreneurship and ownership.

