

THE \$2.3 TRILLION QUESTION

Why Digital Transformation Waste Persists

*A White Paper
Ukemi Way Paper Series
Publication # 001*

April 10, 2026

Version 1.0

By

Christian Stambouli
Founder, Ukemi Way

Copyright Notice

This document contains proprietary intellectual property belonging to Ukemi Way.

No part of this publication may be reproduced, distributed, transmitted, stored in a retrieval system, or translated into any language, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the author, except for brief quotations used in reviews or scholarly commentary.

Copyright © Ukemi Way 2026. All rights reserved.

EXECUTIVE SUMMARY

Why Knowing Doesn't Fix the Problem

The digital transformation industry has documented its failure problem thoroughly for more than three decades. Failure rates between 70% and 88%. Trillions of dollars in cumulative waste. A persistent gap between what organizations spend on transformation and what they sustain.

What the industry has not adequately answered is the more important question: if the failure data is this widely known, why hasn't anything changed?

This paper argues that digital transformation waste persists not because organizations lack knowledge, methodologies, or skilled practitioners, but because four structural forces actively work against the knowledge they already have: the psychology of decision-makers under pressure, flawed approval processes, cognitive model misalignment in leadership, and incentive misalignment in the delivery ecosystem.

Until those forces are confronted directly, investments in better methodology, stronger change management, and improved technology selection will continue to produce marginal gains against a problem that shows no signs of resolution.

A fifth and accelerating factor has emerged: AI integration is surfacing and amplifying pre-existing structural dysfunction in ways that traditional implementation could absorb and conceal.

This paper draws on published research from organizational psychology, behavioral economics, infrastructure economics, digital transformation literature, and current enterprise AI research. It concludes with structural interventions that address root causes rather than symptoms.

The Four Structural Forces

- 1: Escalation of Commitment. The organizational and psychological tendency to continue investing in failing programs because stopping carries higher short-term costs than continuing.
- 2: Strategic Misrepresentation in the Approval Process. Structural pressure to make projects look better on paper than honest analysis supports, producing business cases designed to win approval rather than accurately representing what success requires.
- 3: The Alignment Gap. The mismatch between the cognitive models that experienced leaders use to make decisions and the decision-making logic (iterative, data-driven, and dependent on continuous adaptation) that digital transformation requires.

4: Vendor Incentive Misalignment. The delivery ecosystem is compensated for completion, not sustained outcomes. Accountability ends when value erosion most commonly begins.

ABOUT THIS RESEARCH

This paper is published as a standalone research contribution. The findings and structural analysis are drawn entirely from peer-reviewed literature, major industry research, and longitudinal project performance data.

SECTION I

The Scale of the Problem

The familiar numbers, and the one that isn't.

McKinsey reports failure rates of roughly 70 percent [14]; Gartner finds only 48 percent of projects fully meet their targets [2]; Bain's 2024 research found 88 percent of transformations fail to achieve their original ambitions [3]. Industry estimates place the cumulative global waste from failed digital transformation at roughly \$2.3 trillion to date [16], while worldwide digital transformation spending is projected to reach \$3.4 trillion in 2026 [15].

The AI-specific picture is no better. 74 percent of companies struggle to achieve and scale AI value despite widespread adoption. Organizations average 4.3 pilots, but only 21 percent reach production scale with measurable returns [13]. A recent MIT study found 95 percent of enterprise generative AI pilots fail to deliver measurable P&L impact [4], while 42 percent of companies scrapped most AI initiatives in 2025, up from 17 percent in 2024 [17]. The abandonment rate more than doubled in a single year.

But the most significant and least discussed data point in the transformation literature is not a failure rate. It is what happens to projects counted as successes.

A 2023 McKinsey survey found that while 56 percent of respondents reported achieving most, or all, of their transformation goals, only 12 percent sustained those goals for more than three years. During the later stages of a large-scale change effort, an average of 42 percent of financial benefits is lost [1].

More than half of organizations believe they have succeeded. Less than one in eight sustains that success for more than three years. A significant portion of what the industry counts as successful transformation is, by any honest measure, temporary.

Spending on digital transformation has continued to grow. Failure rates have not meaningfully declined. This report examines the gap between widely documented knowledge and persistent behavior.

SECTION 2

The Measurement Problem Nobody Talks About

The data describing the problem is weak, and that weakness matters.

The Numbers Reported Don't Measure the Same Thing

The spread between reported figures (70 percent failure in some studies, 88 percent in others) reflects a fundamental inconsistency in how the field defines failure. Some studies measure outcomes at the level of the whole transformation program; others count each project within a transformation separately. These are not comparable measurements, yet their figures are routinely cited alongside each other as if they were.

A project delivered on schedule and within budget counts as a success in some studies, regardless of whether the organization effected a transformation. A multi-year enterprise-wide initiative is a failure in others if it falls short of any stated objective. Combining these different definitions into a single headline figure produces a number that feels empirically solid but rests on conceptually inconsistent foundations.

This is not a minor methodological footnote. It is a structural problem with the entire evidence base. After three decades of research, there is no agreement on what constitutes success or failure.

The Go-Live Illusion

The problem is compounded by what is called the go-live illusion: declaring success at technical completion long before sustained business value can be assessed. These are cases where systems go live on schedule, data migration is successful, and integrations function as designed, but business outcomes fall short, revenue targets are missed, and efficiency gains never materialize.

Because most research captures outcomes at or near implementation completion, technical success combined with business failure is systematically undercounted. The 12 percent sustained success figure [1] is the rare data point that tracks what happens after go-live. The real failure rate, measured by sustained business outcomes rather than deployment milestones, is higher than the headline numbers suggest.

What This Tells Us

When organizations benchmark against an incomplete and optimistic picture of industry performance, they systematically underestimate the difficulty of their initiative. When the research base counts temporary achievement as success, it obscures the most important question: not whether transformation can be achieved, but whether it can be sustained.

SECTION 3

Why Knowing Doesn't Change Behavior

The behavioral science explanation for a pattern that three decades of documentation have failed to change.

McKinsey's own senior partner acknowledged: "No one sets out to fail, but research shows that 70 percent of the time, companies do just that. In our experience, it isn't a lack of knowledge that leads to unsuccessful outcomes." [14]

The problem is not information. Organizations know the failure rates. They hire experienced practitioners. They invest in methodology and governance frameworks. And they fail anyway at roughly the same rates they always have. Behavioral science has a well-documented explanation. It operates through a mechanism called escalation of commitment. It describes a pattern most experienced DX practitioners will recognize immediately.

Escalation of Commitment: Staying in a Burning Building

A program begins to show warning signs. Milestones slip, adoption lags, but benefit projections look increasingly optimistic. The rational response to such a pattern is to stop, reassess, and either course-correct or exit. But instead, the organization invests more. Resources are added. Timelines are extended. Projected benefits revised upward. The program continues.

This is not irrationality. It is a well-documented behavioral pattern. Escalating commitment refers to the tendency of decision-makers to persist with failing courses of action, due in part to their unwillingness to admit that their prior resource allocations were in vain. It is what researchers describe as the self-justification explanation [6].

Stopping means admitting the previous investment was wasted. Continuing means that the investment might yet be vindicated. Under organizational pressure, with careers and reputations attached to program outcomes, the psychology of self-justification consistently overrides rational calculus.

"Escalation of commitment is not a sign of incompetence. It is a form of self-protection. People who are invested in a failing project believe that the time and resources already expended will be wasted if they walk away. They may also fear that backtracking will negatively impact their reputation. [6]"

The practitioners and executives leading these programs are aware that their programs are in trouble. They continue anyway because the organizational and psychological costs of stopping exceed the perceived cost of continuing, at least in the short term.

Why Colleagues and Peers Also Escalate

The escalation pattern is not limited to those directly leading a program. Others in the organization, colleagues, peers, and senior leaders who are aware of the program but not running it, also tend to support continuation rather than correction. Research found that multiple forms of psychological connectedness, including perspective-taking, shared team identity, and interdependent working relationships, lead people to vicariously justify others' initial decisions and escalate their own commitment, even in the face of personal financial costs [7].

The implication for DX programs is sobering. Bringing in fresh eyes from outside the project team does not guarantee an independent assessment. If those individuals are organizationally connected to the original decision-makers, they will justify the decisions rather than challenge them. The larger and more visible the transformation program, the wider that circle of connectedness becomes, and the stronger the collective pressure toward continuation rather than correction.

Institutional Inertia as Amplifier

Organizations have weak sensory systems, which make them vulnerable to dangerous signals in their environment [6]. Compounding this, individual and team incentives are often misaligned with the organization's interests. When those incentives reward program continuation over honest assessment, escalation becomes self-reinforcing across the entire organization, not just among those directly leading it.

In practice, signals that indicate a transformation program is failing, adoption metrics, benefit realization gaps, and user resistance data travel through organizational layers designed to absorb and neutralize rather than act on them. By the time failure is visible enough to force a decision, the waste has already accumulated to a scale that makes honest accounting politically untenable.

This is why publishing better failure data, conducting post-mortems, and disseminating better frameworks have not made a difference over three decades. The mechanisms driving waste are structural and psychological. You cannot solve such a problem with information.

SECTION 4

The Approval Process as a Waste Generator

A significant portion of DX waste is designed into programs before execution begins.

Most analysis of DX failure focuses on execution. What went wrong after a program was funded and underway? This section moves the analysis upstream to the approval process itself. A significant and underexamined portion of DX waste is not due to poor execution. It is designed into programs before execution begins. The research basis comes primarily from the work of Bent Flyvbjerg, whose decades of study on large-scale project performance produced findings that map directly onto the DX context [5].

Strategic Misrepresentation: Designing Projects to Get Funded

In environments where projects must compete with other investment opportunities, the approval process creates structural pressure to make a project look as attractive as possible on paper, regardless of whether that picture is accurate.

Strategic project planners and managers sometimes underestimate cost and overestimate benefits to achieve approval. Optimistic planners do this unintentionally; strategic actors do it deliberately. The result is the same: cost overruns and benefit shortfalls. It is called “strategic misrepresentation.” In both cases, unintentional or strategic, project execution is based on a business case designed to win approval rather than accurately represent what the implementation requires and delivers [5].

The competitive dynamics of the approval process make it self-reinforcing. When honest projections compete with optimistic ones for the same budget, optimistic projections always win. Over time, this creates a selection effect. Flyvbjerg describes it as inverted Darwinism: it is not the best projects that get implemented, but the projects that look best on paper. For very large projects, the most significant behavioral bias is political. Strategic misrepresentation is driven by power relations, office politics, salesmanship, and jockeying for position and funds, with millions and sometimes billions of dollars at stake.

Optimism Bias: The Unintentional Version

Not all distortion is deliberate. Optimism bias, the well-documented human tendency to overestimate positive outcomes and underestimate costs and risks, particularly for initiatives we are personally invested in, produces the same result without bad intent. The specific failure mode is called “reference class forecasting failure”: decision-makers consistently treat their project as unique rather than benchmarking against a historical performance of comparable projects [5].

The combination of deliberate strategic misrepresentation at the institutional level and unintentional optimism bias at the individual level means that most DX business cases are structurally unreliable before a single dollar of execution budget is spent.

What the Solution Requires

Better forecasting techniques and appeals to ethics won't address strategic misrepresentation. What's required is institutional change. It requires organizations to redesign the approval process itself by building structural safeguards that don't depend on individuals acting against their own career interests to produce honest assessments [5]. Transparency, independent peer review, and incentive structures that reward honesty and penalize deliberate misrepresentation are some of the key components of such a change.

SECTION 5

The Alignment Gap: Where Projects Are Lost Before They Begin

The most widely experienced structural force, and the one least frequently named as structural.

Misalignment during program inception is almost universally treated as a communication problem. A failure of stakeholder engagement, requirements definition, or executive clarity. This framing leads to communication-level solutions: better workshops, clearer documentation, and more rigorous governance processes. These interventions help at the margins. They do not address the root cause.

The Cognitive Model Gap

The gap is the misalignment between management intuition, developed through years of formal education and practical experience based on traditional economic, strategic, and operating models, and the new logic of a reality shaped by digital technologies. Managers have built their careers on such intuition and make day-to-day decisions accordingly, even subconsciously, while understanding the stated imperatives of successful digital transformation [8].

The misalignment lies between the cognitive models that experienced leaders use to make decisions and the decision-making logic required for successful digital transformation. Leaders do not abandon those models when a transformation program is initiated. They apply them, often unconsciously, in contexts that yield predictably poor results. The result is that transformation programs can have clear executive sponsorship, well-documented objectives, and well-resourced program structures, and still produce misaligned execution, because the mental models guiding daily decisions are not aligned with the transformation's actual requirements.

This is a structural problem. It cannot be solved with a better project charter.

The Budget Allocation Signal

One of the clearest observable signals of this cognitive model gap is transformation budget allocation. Historically, 80 to 90 percent of digital transformation budgets go to technology, leaving insufficient funds for organizational change management, business process optimization, and strategic planning. A useful benchmark: if more than 50 percent of the budget is spent on technology, the strategy is likely off-balance [8].

This allocation pattern reflects the cognitive model of leaders who understand transformation primarily as a technology acquisition, which is precisely the model that produces success declared at go-live and sustained outcome failure.

The Multi-Level Misalignment Problem

The cognitive model gap is compounded by misalignment operating simultaneously at multiple organizational levels. Across teams, business units, and partners; between those running pilot initiatives and those in charge of operations; and between internal stakeholders and external vendors [8].

No single intervention closes it. Aligning executive sponsors does not align business units. Aligning business units does not align operational teams. Each layer of misalignment compounds the others. All rooted in the same problem: different groups operating from different cognitive models of what transformation means and what it requires.

The Financial Consequence

The alignment gap leaves its financial mark in cost overruns and missed objectives, even when it is rarely named as the cause. The cognitive model misalignment operating at inception is the upstream cause that most post-mortems never identify.

SECTION 6

The Delivery Ecosystem's Structural Misalignment

Not a capability problem. A structural compensation problem.

The three forces examined so far operate primarily inside organizations. This section examines a fourth structural force in the external ecosystem of consultants, system integrators, and technology vendors that delivers most large-scale transformation work. The argument is not about the competence or integrity of individual firms. It is about a structural misalignment between how the delivery ecosystem is compensated and what organizations need.

How the Delivery Ecosystem Is Paid

The dominant business model in consulting is built around time and deliverables, not sustained business outcomes. At its core, consulting sells time, access to expert knowledge, and a skilled workforce for a specific period. Project-based fixed-fee pricing is used by 36 percent of consultants, 26 percent employ value-based pricing, and 23 percent charge hourly fees. Combined, most consulting engagements are priced on inputs or fixed deliverables rather than on outcomes sustained over months and years [12].

The structural consequence is direct: the vendor's financial obligation ends at go-live. The period when value erosion most commonly begins is post-implementation, when adoption is supposed to take hold and benefits materialize. It is precisely when the delivery relationship has formally concluded, and accountability has dissolved.

Two Categories of Vendor

Not all vendors have the same line of sight into a program's trajectory, and accountability obligations should reflect that distinction.

System vendors, the platform providers whose technology sits at the center of the transformation, have direct access to platform telemetry. They can see whether the system is being used as implemented, where adoption falls short of baseline expectations, and where technical friction is generating failure signals. What they cannot reliably see is the organizational context behind those signals. A system vendor knows the platform is underused. It cannot know why.

Consulting and implementation vendors occupy a different position. They are embedded in the organization. They sit in steering committees. They observe leadership alignment problems, political friction, capability gaps, and readiness failures that surface during implementation. They can accumulate organizational intelligence that is directly relevant to whether the transformation

will sustain. The commercial structure, however, creates no incentive for them to do so. Several disincentives make silence the path of least resistance. Raising concerns about organizational readiness risks being perceived as scope expansion, or triggering client defensiveness, and putting the contract relationship under strain. The problem is not vendor intent. It is that the commercial model was never designed to reward the kind of honest reporting that would most benefit the client at this stage.

Accountability as Transparency, Not Outcome Ownership

The argument is not that vendors should own outcomes that depend on what client organizations do or fail to do. Leadership alignment, adoption behavior, process redesign, and capability building are the responsibility of the client organization. That ownership cannot and should not be transferred.

For system vendors, the accountability obligation is transparency: reporting what their platform instrumentation already captures during a defined post-go-live window. Not the interpretation of organizational causes, that is beyond their line of sight, but the adoption signals, utilization gaps, and configuration friction that their telemetry already measures.

For consulting and implementation vendors, the obligation is disclosure: a structured reporting requirement at the close of engagement that documents organizational readiness gaps, alignment risks, and warning signals observed during implementation. Not outcome ownership, but honest reporting of what they saw, on the record, before they leave. If that is not specified in the contract, vendors will take the path of least resistance: fulfilling their contract obligations without volunteering what isn't their mandate.

The Industry Is Being Forced to Change

The consulting industry is now under pressure from clients to shift toward outcome-based pricing, quantified ROI, and sustained capability transfer. Clients are increasingly skeptical of broad transformation promises. They are demanding implementation rigor and measurable outcomes [3].

This shift is happening because clients have absorbed enough failure to make the old model politically untenable. The fact that it is happening now, under pressure, is itself evidence that the misalignment existed all along. Section 7 explains why the pressure intensified at this specific moment.

SECTION 7

AI as Diagnostic: Why the Dysfunction Is Now Visible

AI does not create new structural dysfunction. It makes pre-existing dysfunction impossible to conceal.

Every structural force examined in this paper has been present throughout the entire history of enterprise technology implementation. ERP programs in the 1990s produced these patterns. CRM implementations in the 2000s reproduced them. Each successive wave of transformation technology arrived with the promise of breaking the failure cycle and departed having largely repeated it. What is different now is not the dysfunction. It is that the dysfunction is no longer concealable.

AI Exposes What Traditional Implementation Could Absorb

“AI is revealing the weaknesses in many existing transformation projects by exposing cracks in their cloud and data foundations. These revelations are compelling innovation and transformation leaders to reevaluate their playbooks. KPMG, 2024 [9]”

When organizations implemented traditional ERP or CRM systems, a significant level of organizational dysfunction could be absorbed. Shadow processes emerged. Workarounds became embedded in daily operations. Data quality problems were manually rectified by experienced staff. The system went live. Dysfunction continued around it. The organization declared success.

AI embedded in those same systems does not tolerate dysfunction the same way. It requires clean data, coherent process logic, and genuine behavioral alignment. Otherwise, it produces outputs that are visibly wrong, unreliable, or actively harmful. The dysfunction that organizations have sidestepped for decades is now producing observable, attributable failures that cannot be explained away as implementation growing pains.

The Deterministic-Probabilistic Collision

There is a technically precise reason why AI specifically surfaces organizational dysfunction. Organizations are running two incompatible operating systems simultaneously. A deterministic backbone of ERP, financial controls, and compliance, where every decision must be explainable, repeatable, and governed. Alongside is a probabilistic AI layer that produces likelihoods rather

than certainties. Most companies are attempting to integrate probabilistic AI into deterministic, rule-bound systems or replace the deterministic core altogether. But neither works [10].

This is a fundamental architectural tension that forces organizations to confront the quality, coherence, and governance of their existing systems and data before AI integration can proceed. Legacy data and infrastructure architectures cannot power real-time, autonomous AI. Organizations need to evaluate whether their technology foundation is ready to support AI deployments [10].

The Scaling Gap Confirms the Diagnosis

Nearly two-thirds of organizations have not yet begun scaling AI across the enterprise. While 39 percent report EBIT impact, just 6 percent qualify as genuine AI high performers, attributing more than 5 percent of EBIT to AI [13].

The stalling at the pilot phase is the “escalation of commitment” pattern operating in real time. Organizations invest in AI pilots, visible, manageable, and politically safe. Pilots succeed on their own terms. The structural problems that will prevent scaling are present but not yet visible at that scale. When scaling is attempted, those structural problems surface. Rather than confronting them, organizations continue piloting. The waste accumulates.

Why This Explains the Vendor Accountability Shift

Nearly nine out of ten employees now use AI at work, yet only 28 percent of organizations are positioned to turn AI deployment into high-value outcomes. When new technology lands on fragile talent foundations, absent learning infrastructure, misaligned incentive structures, and insufficient capability pipelines, productivity benefits lag by over 40 percent [11].

Clients are demanding outcome-based accountability, not because their standards have risen, but because AI has made the consequences of misaligned delivery impossible to conceal. The 42 percent abandonment rate for AI initiatives in 2025 [17] is not primarily a story about AI. It is a story about organizational foundations that were never adequate, and a delivery ecosystem that was not accountable for sustaining outcomes.

SECTION 8

Implications and Solution Directions

Structural problems require structural solutions. What follows are intervention directions, not prescriptions.

On Escalation of Commitment

The core design principle is to create decision points defined before a project begins, not after it is in trouble. When off-ramp criteria are established prospectively, before organizational and psychological investment has accumulated, they carry a different weight than assessments made mid-program when careers are attached to outcomes. Approaches worth examining include pre-defined “go/no-go” criteria at key milestones, independent challenge mechanisms outside the project sponsor’s authority, and explicit separation between those who approved the investment and those who evaluate its continuation.

On Strategic Misrepresentation in the Approval Process

The approval process itself requires structural redesign. The most evidence-backed intervention is “reference class” benchmarking: requiring that business case projections be anchored to the actual historical performance of comparable projects, not to optimistic internal estimates. This shifts the burden of proof; sponsors must demonstrate specifically why this initiative will outperform the historical record of similar programs. Independent review of business cases that exceed defined investment thresholds, conducted by parties with no stake in the approval outcome, provides a structural check on both deliberate misrepresentation and unintentional optimism bias.

On the Alignment Gap

The alignment gap requires intervention at the level of cognitive models. The starting point is a structured assumption-surfacing exercise at program inception: explicit mapping of the mental models and operating assumptions that key stakeholders bring to the program before those assumptions have shaped design decisions.

On Vendor Incentive Misalignment

The intervention is not outcome-based compensation, which is difficult to structure fairly when a significant share of the outcome depends on the client organization after delivery. The intervention is transparency obligations calibrated to what each vendor observes and can report.

For system vendors, contracts should include a defined post-go-live reporting period during which platform telemetry, adoption rates, utilization data, and configuration health indicators are formally delivered to the client. The data already exists. What is absent is a contractual obligation to define it in a structured way during the period when it is most actionable.

For consulting and implementation vendors, engagement close-out should include a documented readiness assessment: an on-the-record account of organizational risks, alignment gaps, and warning signals observed during delivery. This shifts the incentive structure; vendors who have been honest throughout the engagement have nothing to fear from this requirement. Those who have not are exposed.

Both obligations are forms of transparency, not outcome ownership. For AI-specific engagements, both obligations extend to include AI adoption signals, data quality findings, and pilot-to-production readiness indicators observed during implementation.

On AI Readiness as a Structural Prerequisite

An AI readiness assessment is not a technical exercise. It is an organizational assessment. Before embedding AI in core business systems, the structural forces documented in this paper need to be evaluated honestly, because AI will surface dysfunction, and the cost of discovering that dysfunction after AI deployment is substantially higher than discovering it before. Organizations that treat AI implementation as a technology project layered onto existing foundations will reproduce the go-live illusion at greater speed and greater cost than any previous technology wave.

The Common Thread

All four intervention directions point to the same underlying governance requirement: structures that operate with meaningful independence from the people and organizations most invested in a program's continuation and its declaration of success. Individuals placed in contexts where their psychological and financial interests are aligned with program continuation will tend to act in ways that serve those interests. That is a predictable human response to structural incentives, not a character failure. Governance that accounts for this reality will consistently outperform governance that assumes individuals will act against their own interests when evidence requires it.

CONCLUSION

Are Organizations Prepared to Address the Root of the Problem?

The \$2.3 trillion cumulative cost of digital transformation failure is not primarily a competence problem. The practitioners who lead these programs are experienced. The organizations that fund them are sophisticated. The methodology available to guide them has never been more developed.

The waste persists because it is structurally generated by behavioral dynamics that knowledge interventions cannot reach, by approval processes that reward optimistic misrepresentation, by cognitive model gaps that shape decision-making even when leaders are genuinely committed, and by a delivery ecosystem compensated for completion rather than sustained outcomes.

Three decades of documenting the failure rate have not made meaningful change. This is the clearest possible signal that documentation and awareness are not the solution. The forces generating waste are structural. The interventions required to address them are also structural.

AI integration has not created new structural forces. It has made the existing ones impossible to conceal. The organizations that interpret this as a new technology challenge to be managed are likely to reproduce familiar failure patterns at greater speed and cost. The organizations that interpret it as a diagnostic signal, a rare opportunity to diagnose and address structural dysfunction that has persisted for decades, are positioned to do something that three decades of transformation effort have not produced: break the pattern.

The question is not whether the \$2.3 trillion problem can be addressed. The question is whether organizations are prepared to address it at its root.

Recommended Citation

Stambouli, C. (2026). The \$2.3 Trillion Question: Why Digital Transformation Waste Persists. Ukemi Way Paper Series. <https://www.ukemiway.com/publications/>

METHODOLOGY NOTE

This paper was produced through systematic curation and synthesis of published research from peer-reviewed academic literature, major industry research firms, and longitudinal studies on project performance. Sources were selected based on methodological rigor, recency, and relevance to the digital transformation practitioner. Where findings from adjacent fields, behavioral economics, infrastructure economics, organizational psychology, are applied to DX, the application logic is made explicit in the text. The paper does not advocate for any specific vendor, platform, or proprietary methodology.

REFERENCES

References

- [1] McKinsey & Company (2023). How to implement transformations for long-term impact. McKinsey Insights.
- [2] Gartner (2024). CIO and Technology Executive Survey: Digital Initiative Success Rates. Gartner.
- [3] Bain & Company (2024). Transformation: A Company's Most Important Bet. Bain Insights.
- [4] MIT Project NANDA (2025). The GenAI Divide: State of AI in Business 2025. MIT Media Lab.
- [5] Flyvbjerg, B. (2014). What You Should Know About Megaprojects and Why: An Overview. *Project Management Journal*, 45(2), 6–19.
- [6] Staw, B.M. (1976). Knee-deep in the Big Muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance*, 16(1), 27–44.
- [7] Brockner, J., Rubin, J.Z., Fine, J. (1983). Factors affecting entrapment in escalating conflicts. *Journal of Research in Personality*, 17(2), 195–207.
- [8] Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *The Journal of Strategic Information Systems*, 28(2), 118–144.
- [9] KPMG (2024). Technology Industry Innovation Survey: AI Readiness and Transformation. KPMG Insights.
- [10] Deloitte (2026). Enterprise AI Report: Data Readiness and Infrastructure for AI Deployment. Deloitte Insights.
- [11] EY (2025). Work Reimagined Survey: AI Adoption and Productivity Outcomes. EY Insights.
- [12] Bain & Company / SaaS Capital (2024). Outcome-Based Pricing in Enterprise SaaS. Bain Insights.
- [13] McKinsey Global Institute (2024). AI High Performers: Scaling AI Value Across the Enterprise. McKinsey Insights.
- [14] McKinsey & Company (2022). Common pitfalls in transformations: A conversation with Jon Garcia. McKinsey Insights.
- [15] IDC (2024). Worldwide Digital Transformation Spending Guide. IDC.
- [16] Harkin, B. (2024). Evolving from Digital Transformation to Digital Acceleration Using the Galapagos Framework. Taylor & Francis.
- [17] S&P Global Market Intelligence (2025). Voice of the Enterprise: AI & Machine Learning. S&P Global.