

Flow Like the River: A Learning-Led Framework for Organizational Change

AI Learning Systems Architect | claredygart.com

Clare Dygart

Abstract

The dominant change management frameworks (Lewin's Unfreeze-Change-Refreeze, Kotter's 8-Step Process, and Prosci's ADKAR Model) share three structural flaws that help explain why organizational change efforts are so widely experienced as costly failures: they treat change as a bounded event, locate design authority at the top of the hierarchy, and apply undifferentiated strategies to fundamentally different types of change. This paper proposes a learning-led alternative grounded in continuous change theory, distributed governance, and instructional content-type analysis. Rather than managing change as an episodic disruption requiring top-down orchestration, this framework positions change as the natural state of a healthy organization. It positions learning strategy, not project management, as the discipline best equipped to sustain it.

The Problem with Existing Models

A 70% failure rate is frequently attributed to organizational change initiatives. The statistic has been cited by McKinsey, Kotter, Gartner, and Deloitte, among others, and has become axiomatic in the field. However, as Mark Hughes demonstrated in a critical review published in the *Journal of Change Management* (2011), there is no valid and reliable empirical evidence supporting this figure. The number traces to Hammer and Champy's 1993 acknowledgment that their estimate was "unscientific," and it was subsequently repeated by Kotter and others as though it were research-backed, with each source citing the previous in a self-reinforcing loop.

We do not need to defend the precision of this statistic to accept its signal. The fact that the narrative of change failure is so persistent and so widely believed tells us something real about how people experience organizational change. The question is not whether the number is exactly right, but why the experience is so consistently negative. This paper argues that the answer lies in three structural flaws shared by every dominant change management framework.

Flaw 1: Change as Bounded Event

All three major frameworks treat change as a discrete initiative with a beginning, middle, and end. Lewin's model unfreezes a current state, transitions to a new state,

and refreezes. Kotter's eight steps move from building urgency to anchoring new approaches in culture. ADKAR walks individuals through five psychological gates from awareness to reinforcement.

In each case, the assumption is that the organization exists in a default state of stability that is periodically interrupted by change. But organizations are in continuous flux. Workflow drift, incremental adaptation, unmanaged degradation of systems and practices, emergent innovations at the team level: the most consequential changes often happen outside any formal change effort. The dominant models have no mechanism for addressing this ambient, ongoing change.

Karl Weick and Robert Quinn identified this distinction in a landmark 1999 paper in the *Annual Review of Psychology*, contrasting episodic change, described as "infrequent, discontinuous, and intentional," with continuous change, which is "ongoing, evolving, and cumulative." They argued that the two types follow fundamentally different sequences: episodic change follows unfreeze-transition-refreeze, while continuous change follows the inverse, freeze-rebalance-unfreeze. In the continuous model, the intervention is to pause and make visible what is already happening, rebalance it, and then resume improvisation and learning.

The dominant frameworks are built entirely on the episodic model. They assume a default state of frozen. This paper assumes a default state of flow.

A note on readiness. The ACMP Standard for Change Management and similar professional frameworks treat organizational readiness as a pre-change assessment, a discrete gate that must be passed before a change effort begins. This framing is a direct consequence of the episodic model: if change is an event, then readiness is the condition you achieve before the event starts. But in a learning-led organization operating in continuous flow, readiness is not something you assess before change begins. It is the ongoing condition of a team that has built adaptive capacity into its regular functioning. The question is not "are we ready for this change?" It is "have we built the kind of organization where change is already something we do well?" The former question belongs to episodic change management. The latter belongs to learning strategy.

Flaw 2: Top-Down Locus of Design

Even where these models acknowledge the need for buy-in or individual psychological readiness, the design authority remains with leadership. People experiencing the change are positioned as subjects to be moved through stages, not as agents capable of generating and navigating change themselves. This frames resistance as a problem to overcome rather than a signal to read.

These are fundamentally patriarchal models. Authority flows downward. Compliance is the measure of success. The people closest to the work have the least agency in shaping the change. Even Kotter's language of "coalition-building" and ADKAR's focus on individual readiness maintain the assumption that the locus of design stays with leadership. The people experiencing change are still subjects of someone else's plan, not co-designers of adaptive capacity.

This top-down orientation generates the very resistance it then seeks to manage.

When people feel acted upon rather than empowered, they push back, not because they oppose the change, but because they have been denied agency in shaping it.

Flaw 3: Undifferentiated Change Types

The dominant frameworks apply a single methodology regardless of the nature of the change. But consider the range of changes organizations regularly face:

- A knowledge management system migrates to a new platform.
- A new procedure is established for requesting PTO.
- The organization fundamentally reimagines how it relates to its customers.
- An AI automation initiative transforms core workflows.

These are not the same kind of change. Borrowing from instructional content-type analysis, a taxonomy distinguishing among facts, concepts, processes, procedures, and decision-making frameworks, we can see that some changes are factual or procedural (the knowledge base moved; here is how you request PTO now) while others require deep mental model shifts (rethinking the customer relationship; integrating AI into creative work).

Applying the same eight-step process to a knowledge base migration and a cultural transformation is both wasteful and counterproductive. It over-engineers simple transitions and under-supports complex ones. An unnuanced approach to change types is nearly a guarantee of failure, or at least of wasted resources and unnecessary friction.

Core Principles of a Learning-Led Alternative

Change Is Continuous and Synonymous with Innovation

If we accept that the irreducibly human contribution in an AI-augmented workplace is creative and innovative work (what this paper terms the 15% thesis, referring to the proportion of work that AI cannot perform and that represents uniquely human judgment, creativity, and relational intelligence), then a healthy organization is constantly producing change from within.

Change is not a disruption to be managed. It is evidence that people are doing the work that matters.

The organizational question shifts from “How do we get people through this change?” to “How do we build the conditions where people generate and navigate change as a normal part of their work?” An organization oriented around innovation is an organization that has made change its operating rhythm rather than its occasional crisis.

Unlearn, Innovate, Integrate

Rather than a one-time unfreeze-refreeze sequence, teams in this framework operate in a continuous cycle of unlearning outdated models, innovating new approaches, and integrating new understanding into practice.

This cycle is metabolic, not episodic. And the concept of unlearning here departs significantly from how the change management literature typically frames it. Un-

learning is not a discrete, painful cognitive act: stop believing X, start believing Y. It is a disposition: holding knowledge lightly, treating the space between one's ears as liminal space where reframing, re-examining, and shifting are always possible.

Research on organizational unlearning supports the importance of this capacity. Sinkula (2002) suggests that organizational unlearning begins with changing cognitive structures, mental models, and dominant logics. Empirical studies have found that daily routines and risk aversion are barriers to unlearning, while providing temporal and spatial freedom and facilitating an error-forgiving climate support teams in breaking free from obsolete patterns (Klammer et al., 2020; Amaya et al., 2022). Research on technology implementation specifically identifies prior knowledge and established mental models as critical barriers to change and positions unlearning as a way to address this resistance (Becker, 2010).

But these studies still tend to treat unlearning as an event, something triggered by a crisis or mandated by leadership. This framework reframes it as a practiced skill and an ongoing state. When the Eastern sages say that nothing is unchangeable except the reality of change, there is something profoundly stabilizing about internalizing that message. You no longer have to worry about change because it is simply a given. The anxiety that traditional change management spends enormous energy trying to manage, through communications campaigns, urgency-building, and reinforcement, dissolves when people live in a state where transition is normal rather than exceptional.

Integration, the final phase of the cycle, replaces Lewin's refreeze. There is no new frozen state to arrive at. Integration means absorbing new understanding into the way work gets done, and continuing to move. The river does not stop flowing because it navigated around a rock.

Distributed Governance and Matriarchal Leadership

This framework rejects the patriarchal assumption that change must be directed from above. Instead, governance is distributed. The locus of control for change sits at the point of change, with the teams and individuals closest to the work.

This is not leaderlessness. It is a matriarchal model where leadership means creating the conditions for others to act, not commanding action. Authority is relational and situational rather than positional.

The research base for distributed leadership is substantial and growing. Harris, Leithwood, Day, Sammons, and Hopkins (2007) found a positive relationship between distributed leadership and organizational change, with coordinated patterns of distribution yielding the strongest results. MIT Sloan's Deborah Ancona has argued that the future of management requires a shift from command-and-control to distributed leadership, emphasizing autonomy to innovate and noncoercive alignment around common goals. Research published in *Frontiers in Psychology* (2021) found that distributed leadership positively influences proactive behavior and meaningfulness of work among employees, and that during periods of organizational change, it can help overcome the limitations of pyramidal leadership structures.

Distributed governance also serves as a retention strategy. When people feel autonomous and empowered, when they are agents of change rather than subjects of it, they stay. Traditional change management creates turnover by making people feel acted upon. Distributed governance prevents it by making people feel like co-creators of the organization's direction.

The Manager as Change Doula

A doula does not deliver the baby. The mother does. A doula creates the conditions, reduces friction, provides support, reads the situation, and adapts in real time. The doula is not in charge of the process; the process is natural. The doula's job is to make sure nothing gets in the way and to know when something needs a different kind of intervention.

The manager or team lead in this framework operates the same way. They do not drive, sell, or orchestrate change. They support a process that is already happening. And their first move, the skill that makes everything else possible, is diagnosis.

The Doula's First Move Is Diagnosis

Not all changes are the same, and they do not warrant the same response. The doula's essential competency is the ability to read the environment and diagnose what kind of change is actually in play. This is the foundational act. Before any strategy is selected, before any support is provided, the doula asks: what are we actually dealing with here?

The diagnosis distinguishes among fundamentally different types of change:

Factual and procedural changes (a system migration, a new policy, an updated process) call for updated job aids, adjusted links, and frictionless access to information at the point of need. These are embodied cognition tools, not communications campaigns. If the knowledge base moved from one platform to another, the doula's job is to analyze where people currently look for the link and make sure it is updated. No town hall required. No urgency to build.

Mental model shifts (new ways of thinking about customers, markets, collaboration, or the role of AI in work) require space for unlearning, dialogue, practice, and reflection. These changes cannot be accomplished through information delivery. They require the kind of supported cognitive and social process that learning design is specifically built to facilitate.

Getting the diagnosis right is what makes proportionate response possible. Applying a procedural response to a mental model shift will fail. Applying a mental model intervention to a procedural change wastes resources and generates cynicism. The doula who misdiagnoses the type of change will provide the wrong kind of support, and the change will stall, not because people resisted it, but because the support did not match the need.

This diagnostic capability is trainable, and developing it at scale is one of the most high-leverage investments an organization can make. It is also where learning strategy stops being a support function and starts leading. The doula-as-diagnostician is

the primary mechanism through which learning design becomes the engine of organizational adaptability.

Friction as the Primary Diagnostic

Where traditional models add layers of process, communications plans, training roll-outs, reinforcement campaigns, stakeholder management frameworks, this framework treats friction as the enemy. The question is always: what is getting in the way of people doing what they are already trying to do?

Transaction cost economics, originating with Coase (1937) and developed by Williamson (1975, 1985), provides the theoretical language here. Transaction costs are the costs of dissipation in resource exchange, what Stigler (1972) explicitly compared to friction in physics. In organizational terms, friction shows up as coordination overhead, information latency, approval bottlenecks, and redundant process.

Research on organizational coordination suggests that companies waste as much as 80% of their time on coordination tasks: meetings, approvals, handoffs, and rework. Research on information management in organizations finds that fragmented knowledge, inconsistent taxonomies, and disordered data are direct sources of friction that impede transformation.

Much of what passes for change management actually introduces friction. Every communications campaign, every training rollout, every reinforcement checkpoint is a process layer that costs time and attention. The goal of this framework is to identify and remove barriers, not to build elaborate scaffolding around them.

A note on coherence across distributed teams: one legitimate concern with distributed change governance is how teams executing change locally remain aligned with the broader organization. The answer lies not in centralized approval but in shared standards: common diagnostic vocabulary, transparent decision logs, and lightweight peer visibility. Coherence in a flow model is achieved through shared language and trust, not through hierarchical checkpoints.

Reduced Information Costs

When the locus of control for change sits at the point of change, information costs drop dramatically. The traditional model requires information to travel up the hierarchy, get processed into a change plan, and be communicated back down. Every step adds latency, distortion, and cost.

When the people closest to the change have the authority and capability to respond, the signal-to-action distance collapses. The team that notices a workflow problem can address it directly rather than escalating it through three levels of management, waiting for a change initiative to be designed, and then receiving instructions on how to change.

This is one of the core economic arguments for distributed governance: it is not just more humane, it is more efficient. The cost of change drops when you stop routing every adaptation through a centralized change management apparatus.

The Flow Metaphor

Unfreeze-change-refreeze assumes a default state of frozen. This framework assumes a default state of flow.

A river does not stop moving when it encounters an obstacle. It flows around it. The people in the current do not need to be unfrozen, motivated, or driven through stages of psychological readiness. They need a clear channel and the absence of unnecessary obstruction.

When a person is walking through a crowded city and encounters someone standing still in the middle of the sidewalk, the traditional change management approach would be to assess the blockage, build a coalition of fellow pedestrians, communicate the urgency of moving, develop a plan to convince the person to step aside, execute the plan, and celebrate the short-term win. The alternative, flow like the river, means the current was already moving. Just keep moving.

The organizational implications are profound. An organization in flow does not need to spend money boiling water to unfreeze its people. It invests instead in maintaining the conditions that keep people moving: clear information at the point of need, distributed authority to act, managers who read the environment and reduce friction, and a culture where holding knowledge lightly is a practiced skill rather than a terrifying demand.

Connections to Existing Theory

This framework does not emerge from a vacuum. It builds on and extends several significant bodies of work.

Weick and Quinn (1999) provide the foundational distinction between episodic and continuous change, and their description of the continuous change process, freeze-rebalance-unfreeze, aligns closely with this paper's unlearn-innovate-integrate cycle. Their observation that continuous change is characterized by "recurrent interactions, shifting task authority, response repertoires, emergent patterns, improvisation, translation, and learning" essentially describes the organizational culture this framework seeks to create.

Peter Senge's learning organization (1990) provides essential groundwork, particularly the five disciplines of systems thinking, personal mastery, mental models, shared vision, and team learning. Senge argued that the only sustainable competitive advantage is an organization's ability to learn faster than the competition. This framework extends Senge's vision by distributing the learning function to the point of change rather than requiring organization-wide transformation initiated from the top, addressing a known limitation of Senge's model, which critics have noted is difficult to implement in bureaucratic organizations through learning initiatives alone (Finger and Brand, 1999).

The organizational unlearning literature, including work by Hedberg (1981), Nystrom and Starbuck (1984), Tsang and Zahra (2008), and the recent comprehensive review by Grisold et al. (2024), supports the centrality of discarding outdated mental models as a precondition for adaptive capacity. This framework's contribution is to

reframe unlearning as a continuous disposition rather than an event-driven intervention.

Distributed leadership research, notably Harris et al. (2007), Spillane (2006), and Ancona at MIT Sloan, provides empirical grounding for the claim that distributing leadership authority to the point of change produces superior organizational outcomes, increased proactive behavior, and greater meaningfulness of work.

Transaction cost economics, Coase (1937), Williamson (1975, 1985), Stigler (1972), provides the theoretical vocabulary for the friction diagnostic and the information-cost argument for distributed governance.

Implications for Practice

For L&D Leaders

Learning strategy is not a support function in this framework. It is the primary discipline through which organizational change capacity is built. The chief learning officer, or whatever the role is called, becomes the architect of the organization's adaptive infrastructure: the systems, tools, and cultural conditions that allow teams to unlearn, innovate, and integrate continuously.

This means L&D must move beyond course design and content delivery into organizational design, team cognition analysis, and the development of manager learning-design literacy.

For Managers and Team Leads

The manager's role shifts from change driver to change doula, and the doula's first move is always diagnosis. What kind of change is this? Does it require a job aid or a conversation? An updated link or a team dialogue about mental models? The ability to diagnose correctly and then select a proportionate response is the core competency. It is trainable, and it replaces the need for expensive, centralized change management infrastructure.

In practical terms, planning in a flow model looks different from traditional change planning. There is no lengthy change plan document traveling up the hierarchy for approval. Instead, the plan is local, iterative, and held by the team closest to the change. It is lightweight by design: a shared diagnosis, a proportionate response, and a feedback loop. Managers develop the habit of asking: what do we know, what do we need, and how will we know if it's working? The answers stay at the level where the work is happening.

For Executives

The economic case is straightforward. Traditional change management is expensive, not because the consultants cost too much, but because the model itself generates friction, latency, and resistance as structural features. Distributed governance reduces information costs by collapsing the signal-to-action distance. Investing in adaptive

capacity at the team level is cheaper and more sustainable than funding episodic change initiatives.

It is also a retention strategy. People who feel like agents of change stay. People who feel like subjects of change leave.

Measuring Flow: Indicators of Organizational Adaptive Capacity

Traditional change management metrics, adoption rates, milestone completion, stakeholder satisfaction surveys administered at project close, are artifacts of the episodic model. They measure whether a bounded change event “landed.” In a continuous flow model, there is no project close, no adoption deadline, and no post-rollout survey that captures what matters. The measurement paradigm itself must shift.

This does not mean measurement becomes impossible. It means the field needs to stop reaching for the easy metrics and start asking what organizational adaptive capacity actually looks like when it is working, and then finding ways to observe it.

A Portfolio of Proxies

No single indicator tells the full story of an organization in flow. What is needed is a small portfolio of proxy indicators, observable data points that stand in for the underlying condition, reviewed periodically and interpreted together. Triangulation matters: an escalation rate that drops could mean teams have gained agency, or it could mean people gave up escalating. A single data point is ambiguous. Two or three, read in combination, become meaningful.

The measurement approach should also distinguish between **leading indicators**, data that predicts the health of adaptive capacity before a change event occurs, and **lagging indicators**, data that describes trend over time after the fact. Both are useful; neither alone is sufficient.

The following indicators are proposed as candidates for a portfolio approach. They are not exhaustive, and organizations should expect to calibrate the mix based on their context and data availability.

Escalation rate. In most organizations, some version of escalation data already exists in ticketing systems, project logs, or approval workflows. As distributed decision-making capacity develops, teams should escalate fewer decisions upward, not because they are avoiding accountability, but because more decisions fall within their diagnosed scope of authority. Tracking escalation rate over time provides a lagging indicator of whether distributed governance is actually taking hold.

Time-to-resolution on workflow problems. Traceable in project management tools, this measures how quickly teams move from noticing a problem to acting on it, the signal-to-action distance described in the transaction cost economics section. Shorter loops over time suggest teams are acting rather than waiting. Extended loops suggest friction, unclear authority, or insufficient diagnostic capability.

Voluntary knowledge-sharing activity. Wiki edits, internal documentation contributions, forum posts, and other behavioral traces of knowledge sharing are imperfect

but meaningful proxies for a learning culture. They are behavioral, not self-reported, and they trend over time. An organization in flow generates and shares knowledge as a byproduct of regular work.

Turnover patterns. Exit interview data is messy, but it is a real signal, particularly whether people leave citing a sense of being acted upon versus empowered. This connects directly to the retention argument in the distributed governance section. Parsed carefully, turnover data can indicate whether the culture is moving toward agency or away from it.

Escalation Literacy as a Measurable Construct

One of the most important mental model shifts this framework requires is a changed relationship to escalation itself. In a traditional top-down model, escalation is either compliance (I followed the process) or failure (I couldn't handle it alone). In a flow model, escalation is a skill: knowing when a decision genuinely exceeds your scope and moving it appropriately, versus knowing when you have the authority and capability to act locally. Getting this diagnosis right is part of what makes distributed governance work.

This construct, escalation literacy, is measurable, but not through direct behavioral questions. Asking “did you escalate when you thought it was appropriate?” invites social-desirability bias and tells you almost nothing about the underlying mental model. A more rigorous approach uses **vignette-based or situational judgment measures**: scenario prompts that reveal how respondents reason about escalation rather than how they rate their own behavior.

A vignette might read: *Your team identified a workflow problem affecting two other departments. Walk me through what you did.* The response reveals both the mental model (when is this mine to solve?) and the behavioral tendency (did they act, wait, or escalate?). Administered periodically, not as a post-change survey but as a regular organizational pulse, vignette-based measures provide a leading indicator of whether distributed diagnostic capability is developing.

Used alongside the behavioral proxy data described above, escalation literacy surveys provide the qualitative texture that quantitative proxies alone cannot. Together they form a measurement approach that is honest about its limitations, triangulated enough to be meaningful, and consistent with the continuous rather than episodic nature of the framework itself.

Areas for Further Development

This framework opens two significant lines of inquiry that warrant deeper exploration:

Team cognition as the unit of change capacity. If change happens at the team level, then the team's collective cognitive capacity, not individual readiness, is the relevant unit of analysis. How do teams develop shared mental models, and how do those models adapt under pressure? The growing literature on team cognition, trans-active memory systems, and collective intelligence provides a foundation; connecting it explicitly to change capacity is a meaningful research direction.

Transactional vs. transformational learning. Not all learning is created equal. How does the distinction between transactional learning (information transfer) and transformational learning (mental model shifts) map onto the change-typing framework proposed here? Clarifying this relationship would sharpen both the diagnostic framework and the practice implications for learning designers and managers.

Conclusion

The dominant change management frameworks have been justified for decades by a failure statistic that nobody can substantiate, applied through models that treat change as an event to be managed rather than a condition to be embraced, and implemented through top-down processes that generate the very resistance they seek to overcome.

A learning-led alternative starts from different premises: change is continuous; it emerges from within; different types of change require different responses; the people closest to the change are best positioned to navigate it; and the leader's role is to create conditions, not command compliance.

Flow like the river. The current is already moving. The work is not to push the water. It is to clear the channel.

References

- Amaya, A., et al. (2022). Team unlearning and new product development success. *NPD Research*.
- Ancona, D. (2022). Distributed leadership. MIT Sloan School of Management.
- Becker, K. (2010). Facilitating unlearning during implementation of new technology. *Journal of Organizational Change Management*, 23(3), 251-268.
- Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16), 386-405.
- Finger, M., & Brand, S. B. (1999). The concept of the learning organization applied to the transformation of the public sector. In M. Easterby-Smith, L. Araujo, & J. Burgoyne (Eds.), *Organizational Learning and the Learning Organization*.
- Grisold, T., et al. (2024). Organizational unlearning as a process: What we know, what we don't know, what we should know. *Management Review Quarterly*.
- Hammer, M., & Champy, J. (1993). *Reengineering the Corporation*. Harper Business.
- Harris, A., Leithwood, K., Day, C., Sammons, P., & Hopkins, D. (2007). Distributed leadership and organizational change: Reviewing the evidence. *Journal of Educational Change*, 8, 337-347.
- Hedberg, B. (1981). How organizations learn and unlearn. In P. C. Nystrom & W. H. Starbuck (Eds.), *Handbook of Organizational Design* (Vol. 1, pp. 3-27). Oxford University Press.
- Hughes, M. (2011). Do 70 per cent of all organizational change initiatives really fail? *Journal of Change Management*, 11(4), 451-464.

- Klammer, A., et al. (2020). Two forms of organizational unlearning. *Management Learning*, 51, 598-619.
- Kotter, J. P. (1996). *Leading Change*. Harvard Business School Press.
- Nystrom, P. C., & Starbuck, W. H. (1984). To avoid organizational crises, unlearn. *Organizational Dynamics*, 12(4), 53-65.
- Senge, P. M. (1990). *The Fifth Discipline: The Art and Practice of the Learning Organization*. Doubleday.
- Spillane, J. P. (2006). *Distributed Leadership*. Jossey-Bass.
- Stigler, G. J. (1972). The law and economics of public policy: A plea to the scholars. *Journal of Legal Studies*, 1(1), 1-12.
- Tsang, E. W. K., & Zahra, S. A. (2008). Organizational unlearning. *Human Relations*, 61(10), 1435-1462.
- Weick, K. E., & Quinn, R. E. (1999). Organizational change and development. *Annual Review of Psychology*, 50, 361-386.
- Williamson, O. E. (1975). *Markets and Hierarchies*. Free Press.
- Williamson, O. E. (1985). *The Economic Institutions of Capitalism*. Free Press.