

Enterprise AI Adoption as an Organisational State Transition

ENTERPRISE AI

ORGANISATION THEORY

STATE TRANSITIONS

AI ADOPTION

ABSTRACT

Organisations with access to the same frontier AI capability end up in very different places: most remain in pilots, a minority delegate consequential work at scale, and some retreat after apparent success. The prevailing explanation treats adoption as accumulated readiness, scored across capability, data, governance and skills, in which any progress counts. This paper proposes a different account. Adoption is an organisational state, defined as sustained delegation of consequential work to autonomous systems, and it is governed by two mechanisms. First, organisational controls are complements: the weakest control binds, so an organisation cannot average its way in. Second, delegation generates the operating evidence that sustains institutional confidence in delegation, a feedback that makes the adoption state self-maintaining once entered and out of reach until the loop closes. The feedback separates the conditions for entering the state from the conditions for remaining in it, so movement between states is abrupt, asymmetric and history-dependent: a severe incident can eject an organisation whose audited controls never changed, and recovery depends on whether the accumulated evidence record survives. The entry and exit thresholds are characterised exactly in a minimal model, and the account is stated as falsifiable predictions measurable from ordinary workflow and incident records.

KEYWORDS enterprise AI adoption; organisational state; delegation; hysteresis; complementarities; state transition.

1 The adoption puzzle and the claim

The controlled evidence on generative AI is task-level: a randomised deployment in customer support and a field experiment in consulting both find substantial gains [11, 19]. There is no comparable evidence on sustained organisational adoption. What exists is survey material from consultancies and industry bodies [51, 55, 56]: self-reported, drawn largely from AI-active respondents, and not controlled in any respect. Read as description rather than measurement, it is consistent in direction: widespread experimentation, much less scaled operational deployment, and little attributed earnings impact. None of it establishes the distribution or dynamics claimed below. It does motivate a sharper puzzle than slow diffusion: organisations with access to the same capability appear to end up in different places and stay there, some stalled at pilots, some in retreat after apparent success.

The standard readiness account treats deployment success as an additive function of capability, data, governance, engineering, skills, and change management. Its operational instruments are maturity ladders and averaged scorecards: progress on any dimension counts as movement towards the adoption state. Read as a theory of dynamics, that form predicts approximately continuous organisational variation, incremental movement, and no large downward transition caused by one event. We instead study the conjecture that consequential delegation is approximately bimodal over relevant populations and that movement between modes is abrupt, asymmetric, and history-dependent.

Stated plainly, the explanation offered here is this. Enterprise AI adoption resists point solutions because it is a systems problem twice over. Organisational controls (verification, permissioning, observability, process design, staff capability) are complements: the weakest one binds, so progress on the others cannot compensate, and no averaged readiness score reveals whether an organisation can cross. And even with adequate controls across the board, adoption only establishes itself once delegation begins generating the evidence that sustains institutional confidence in delegation; once established, the same loop holds the organisation in place through routine failures. The weakest link explains why an organisation cannot average its way in. The evidence loop explains why crossing is a genuine state transition, abrupt and asymmetric, rather than the last increment of a checklist. If this account is right, the practical consequences follow directly: maturity scorecards measure the wrong thing; investment must be coordinated across controls rather than concentrated where progress is easiest; and the institutional record of delegated work (audited outcomes, decision histories, post-mortems) is part of the adoption state itself, worth protecting through vendor changes and reorganisations, because its survival determines whether an organisation recovers after a serious incident.

The intended contribution, stated in one sentence: complementary controls in which the weakest binds, combined with a single institutional feedback in which delegation generates the evidence that sustains confidence in delegation, are sufficient to make enterprise AI adoption a two-state system with exactly characterised entry and exit thresholds, so that where an organisation sits depends on its history and not only on its capability or controls. This is a firm-level synthesis of familiar ingredients rather than a claim that complementarities, trust, path dependence, or hysteresis are individually new. The paper combines: (i) an observable, consequence-weighted delegation measure; (ii) a latent organisational state with an adoption region; (iii) complementary controls; and (iv) an institutional feedback linking delegation, evidence, and confidence. The resulting transition boundaries are characterised exactly, and the mechanism yields predictions about threshold entry,

incident regression, and conditional recovery that can be checked in workflow and incident records. Appendix A gives the full evidence base and literature comparison; Appendices B–G provide definitions, proofs, numerical illustrations, empirical designs, limitations, and reproducibility details.

2 Adoption as an organisational state

Fix a class of consequential work: tasks whose outputs bind the organisation through commitments, transactions, altered records, external communications, or decisions that take effect without per-instance human ratification. The measure of adoption is the consequence-weighted share of that work that takes effect without prior human approval, computed over an accounting window from ordinary workflow and approval records; Appendix B states the definition formally. An organisation is in the adoption state when this delegation share stays above a pre-specified level throughout a window long enough to contain its normal incident flow. The definition excludes advisory use in which every output is ratified, and it excludes a deployment that operates only until its first routine failure. Adoption means sustained delegation maintained through routine failures; a model in place, or one successful demonstration, does not qualify.

The delegation share indicates state membership; it is not the complete state. Two organisations can show the same share today yet respond differently to an incident because they differ in accumulated evidence, institutional confidence, or control configuration. The underlying state is partially hidden, and the research question separates into two parts: where an organisation currently sits, and how its state moves near the boundary of the adoption region.

Controls (verification coverage, permissioning and identity for computational agents, observability of agent actions, process redesign around delegation, staff capability) are aggregated so that weak components can dominate. The additive scorecard, in which progress on any dimension counts equally, is the special case in which components substitute perfectly for one another; the aggregate used here nests it and extends to a weakest-link limit in which the weakest control binds. Capability improvements lower the base incident rate and thereby move the boundary of the adoption region. They do not relocate the organisation in its control configuration, and they do not erase its history.

3 The delegation–confidence–evidence mechanism

The minimal dynamic state has three coordinates: institutional confidence, the delegation share, and an institutional evidence stock. Confidence governs a target level of delegation that rises steeply around an authorisation threshold. Delegation generates successful operating experience; confidence is updated from that experience and reverts towards a sceptical prior when evidence is not renewed. Incidents produce multiplicative confidence losses, and retrenchment can be faster than expansion. Sustained delegation also builds the evidence stock: audited outcomes, calibrated expectations, decision records, post-mortems, and evidence about which failures are ordinary rather than disqualifying.

The stock enters the mechanism in one place: it dampens the effect of new incidents on confidence. A deep record lets an organisation contextualise a failure that would terminate a pilot. The result is a loop, from delegation to evidence to confidence and back to delegation, operating in an incident environment set by control adequacy. The loop makes the adoption state self-maintaining. Low delegation produces too little evidence to support high confidence; high delegation produces the evidence that keeps confidence above the authorisation threshold.

This mechanism yields the conjectured structure exactly. Every equilibrium of the deterministic model lies on a one-dimensional locus, and the folds of that locus are the entry and exit thresholds: between them, a stable low-delegation state and a stable adoption state coexist, separated by an unstable boundary (Theorem C.1). The band exists because of the evidence stock. Strengthen its buffering and the band widens, with the larger displacement on the adoption branch; remove it and entry and exit coincide, so history ceases to matter (Corollary C.2). How fast an organisation expands or retrenches changes how quickly collapse and recovery occur, and nothing about where the thresholds sit. At the baseline parameters the analytic folds are $q^\uparrow = 0.8788$ and $q^\downarrow = 0.7906$, and forward simulation of the full model recovers both to the third decimal. Figure 1 shows the mechanism and the fixed-point geometry. The full model is in Appendix B, exact statements and proofs in Appendix C, and numerical illustrations in Appendix D.




Figure 1. Mechanism and fixed-point geometry. Solid branches are stable in the reduced flow; the dashed branch is unstable. Numerical thresholds use the baseline parameterisation.

4 Empirical implications, failure conditions, and scope

The geometry yields three primary tests. First, *interactions should dominate additivity*: control configurations and their weakest components should predict entry into the adoption state better out of sample than averaged maturity scores. Second, *adoption should exhibit state-transition signatures*: where the population spans the transition band, sustained delegation should be approximately bimodal, time to entry should rise sharply near $q^\ddagger$, and the same organisation can remain in the adoption state at a q that would not have taken it there from a pilot state. Third, *severe incidents should cause regression at unchanged audited controls*, with recovery strongly conditional on survival of the institutional evidence stock. Inside the band, a surviving record can keep re-entry feasible; destroying or discrediting the record restores the higher entry requirement.

These predictions fail cleanly. The account should be rejected if additive models match interaction models out of sample; if a population known to span the proposed boundary has a unimodal, continuously filled distribution of sustained delegation; if severe matched incidents do not alter delegation; or if recovery is independent of whether decision records, outcome histories, and post-mortem records survive. Secondary predictions are that materially different control configurations can sustain the same adoption state, provided no critical component falls below the relevant floor, and that variance of delegation fluctuations may rise before shock-driven collapse without the rising lag-1 autocorrelation associated with classical critical slowing down.

The dependent variable is measurable from records organisations already keep: agent action logs, pre-approval and human-touch rates, rollback events, transaction or sensitivity weights, and incident histories. A longitudinal platform panel can test interaction structure, bimodality, transition hazards, and configuration dispersion. Incident event studies can compare delegation trajectories with matched controls and interact the event with evidence-record survival. Shared capability releases and staggered regulatory obligations provide candidate event and natural-experiment designs, subject to the identification qualifications in Appendix E.

The model is an existence proof, not an estimated organisational system. Its numerical values are not claims about firms. The control taxonomy is illustrative, and the theory is limited to consequential delegation by systems autonomous enough to fail consequentially; it does not describe ordinary advisory use. Its value is narrower but testable: familiar organisational mechanisms are sufficient to generate lasting divergence, threshold entry, asymmetric regression, and conditional recovery without assuming any lasting difference in capability.

5 Conclusion

Enterprise AI adoption is better represented by a delegation share, a boundary, and alternative stable states than by a single maturity score. The central structural claim is that a durable institutional evidence stock converts delegation into a self-maintaining state and separates the conditions for entering the adoption state from those for remaining there. This makes governance records, audited outcome histories, and post-mortem records part of the adoption state itself rather than documentation around it. The next empirical task is state estimation from workflow and incident records.

A Evidence base, inherited theory, and novelty positioning

A.1 The adoption puzzle and its evidentiary status

Two bodies of evidence motivate the theory. The first concerns capability. In a randomised deployment across 5,172 customer-support agents, access to a generative assistant increased issues resolved per hour by about 15% on average and by roughly a third for novice workers [11]. In a field experiment with 758 consultants, subjects using GPT-4 completed more tasks and achieved substantially higher assessed quality on tasks inside the model's competence [19]. These are controlled task-level estimates and are large by the standards of workplace interventions.

The second concerns organisational adoption. McKinsey's late-2025 survey reported AI use in at least one function among most respondents, but much lower rates of scaling beyond pilots or attributing material earnings impact to AI [55]. Deloitte reported that only a minority of AI-active firms had converted a substantial share of pilots into scaled operational deployment [51]. S&P Global reported a rising share of companies abandoning most AI initiatives [56]. Gartner separately forecast high cancellation rates for generative and agentic AI projects [27, 28]. A widely circulated MIT report described a very high share of pilots without measurable profit-and-loss impact; the source manuscript treats it as evidence of the practitioner narrative, not as a reliable measurement, because the methodology was criticised [13]. Representative public statistics also show that firm-level AI use remains much lower than individual use and is sensitive to survey wording [1, 4, 5].

These sources have important limits, and none of them is a controlled study. Consultancy surveys sample AI-active leaders and measure perceptions. Project-abandonment statistics do not distinguish termination of pilots from regression after sustained consequential delegation. Binary self-reports do not distinguish an employee drafting an email from an institution allowing an autonomous system to execute a transaction. None of the sources measures an organisation-wide delegation distribution, identifies entry and exit thresholds, times a transition, or attributes regression to an incident. The evidence therefore motivates but does not establish the theory. The empirical programme in Appendix E is part of the contribution because the required object is not available in existing public data.

Historical analogies explain delay but not the full conjecture. Measured productivity lagged electrification because factories had to be reorganised around the technology [18]; the information-technology productivity puzzle similarly resolved in favour of organisational co-investment and intangible complements [6, 7, 10]. These precedents predict slow and heterogeneous diffusion. The sharper proposition here is that matched organisations may occupy alternative stable delegation states, cross the boundary abruptly, and regress after apparent success.

A.2 What the model inherits

The paper uses established mechanisms and makes the synthesis, unit of analysis, measurement, and exact transition geometry the contribution.

Complementarities and weakest links. Complementarity theory shows that mutually reinforcing organisational practices produce non-additive returns, multiple coherent configurations, and failures of piecemeal imitation [8, 45, 46]. Empirically, returns to information technology concentrated in firms making complementary organisational investments [9]. Related weakest-link and rugged-landscape models show why a mean score can conceal a binding component [38, 41, 49]. Most directly, Gans and Goldfarb [25] derive discrete bundled automation from O-ring complementarity in a static task-level model. The present paper uses a CES aggregate so that the additive account is an explicit nested special case and then couples that aggregate to organisational state dynamics.

Punctuated change, thresholds, and hysteresis. Punctuated-equilibrium theories describe long periods of convergence interrupted by reorientation [29, 50, 60]. Formal diffusion and catastrophe models produce metastability and adoption hysteresis [33, 43]. Critical-transition theory supplies the language of alternative stable states, folds, and warning signals [52–54]. Investment theory shows how durable or sunk stocks can separate entry and exit thresholds [21, 22]. The present paper identifies a specific organisational stock that performs this role and derives its comparative static exactly.

Confidence, reliance, and organisational memory. Reliance on automation depends on trust and observed performance [40, 48]. Gao and Lee [26] provide a dynamical operator-level loop in which reliance gates the evidence from which trust is updated. Algorithm aversion and threat rigidity motivate asymmetric responses to failure [20, 57], while organisation-level trust repair motivates the importance of institutional context after failure [30]. Absorptive capacity and path dependence explain how accumulated organisational knowledge can sustain or lock out future action [3, 14, 17]. The present model aggregates these ideas at the organisation level and makes the stored evidence stock observable through its consequences for delegation and recovery.

Equifinality and sociotechnical organisation. Equifinality permits different configurations and paths to support a functionally equivalent state [31, 37, 61]. Configurational studies have documented multiple pathways to AI adoption in cross-section [39, 44]. Sociotechnical and dynamic-capabilities traditions emphasise heterogeneous organisational responses to the same technology [58, 59]. The present contribution is a dynamical derivation of equifinality with a threshold and a weakest-component floor, rather than another static typology.

Complexity concepts can become unfalsifiable metaphors in management research [12, 34], and tipping-point language is easily overused [47]. The discipline adopted here is to specify observable dynamical signatures in advance [23]:

bimodality under a population condition, within-organisation hysteresis, incident regression at unchanged controls, and evidence-dependent recovery.

A.3 Nearest neighbours and the remaining gap

Several recent works are close but differ in the conjunction of unit, mechanism, and predictions. Ercole [24] model health-AI adoption as an evolutionary game among clinicians, obtaining bistability through coordination externalities and cultural lock-in. Huang and Vishnoi [35] obtain path dependence in individual delegation through human-skill erosion. Gans and Goldfarb [25] derive discrete bundled automation in a static task-level model. Cua and Bernardo [16] simulate population-level generative-AI adoption with a trust stock and asymmetric recovery but do not derive firm-level delegation bistability or exact transition boundaries. Liu et al. [42] estimate polarised, state-like delegation dynamics among individual store managers. Configurational studies establish cross-sectional equifinality rather than the transition mechanism studied here.

Against those neighbours, the source manuscript claims the conjunction of five commitments:

1. the unit is an individual organisation rather than a population of adopters or an individual worker;
2. adoption is sustained, consequence-weighted delegation computable from workflow records;
3. the mechanism is a delegation–confidence–evidence feedback in which hysteresis is generated by an institutional evidence stock rather than coordination externalities or skill erosion;
4. threshold entry, hysteretic exit, incident regression, conditional recovery, and configurational equifinality arise jointly from that mechanism, with transition boundaries characterised analytically; and
5. the account is stated as falsifiable hypotheses with named disconfirmers.

The literature search reported in the source manuscript covered arXiv, SSRN, and scholarly indexes in July 2026 under the paper’s terminology and older synonyms. That search is not exhaustive. The defensible novelty claim is therefore that this specific synthesis and dynamical form were not found in the searched corpus, not that no related framework exists.

B

Definitions, state space, and full model

B.1 Definitions

Definition B.1 (Organisational state). An organisational state is the latent configuration of organisational properties that determines the organisation’s propensity to delegate consequential work to autonomous computational systems under uncertainty. It is partially observed through routing, approval, incident, rollback, audit, and recovery behaviour.

Definition B.2 (Delegation order parameter). Fix consequential work W and a weight function w specified ex ante from a pre-registered sector taxonomy, with w_i proportional to a stated consequence measure such as transaction value or record sensitivity. Over an accounting window W_t , the delegation share is

$$x(t) = \frac{\sum_{i \in W_t} w_i \mathbf{1}[\text{instance } i \text{ took effect without prior human approval}]}{\sum_{i \in W_t} w_i} \in [0, 1]. \quad (1)$$

Definition B.3 (Adoption state). An organisation is in the adoption state with respect to W over a window τ if $x(t) \geq x_A$ throughout τ , where τ is fixed ex ante and contains the organisation’s normal incident flow. Adoption is sustained delegation maintained through routine failures.

Definition B.4 (Transition and regression). Let $\mathcal{A}$ be the region whose delegation coordinate satisfies $x \geq x_A$. An adoption transition is movement from outside $\mathcal{A}$ to inside it, sustained over the prescribed window. An adoption regression is the reverse movement, similarly sustained.

The order parameter is behavioural and independent of self-report. Consequence weighting excludes advisory use in which a human approves every instance. The sustained criterion ensures that a system that delegates only until its first failure has not yet established adoption. Familiar labels such as experimentation, pilot, controlled deployment, and autonomous deployment may be represented as regions of state space, but the regions need not be totally ordered and movement need not be monotone. The numerical model introduces a pilot floor $x_p > 0$; this narrows its domain relative to the general definition.

B.2 State representation and complementary controls

Write the state coordinates as

$$S(t) = (x(t), T(t), E(t), s_1(t), \dots, s_K(t), \dots), \quad S(t+1) = F(S(t), \varepsilon(t)).$$

The delegation share is directly measurable. Control components are auditable in principle. Confidence and evidence are latent but produce behavioural signatures and, under the conditions of Theorem C.3, are exactly recoverable in the noiseless known-parameter model.

Controls are represented by $s = (s_1, \dots, s_K)$ and aggregated by

$$q(s; \rho) = \left(\frac{1}{K} \sum_{j=1}^K s_j^\rho \right)^{1/\rho}, \quad \rho \leq 1, \quad (2)$$

with the geometric mean at $\rho = 0$. At $\rho = 1$, components are perfect substitutes and adequacy is their arithmetic mean. For $\rho < 1$, weak components increasingly dominate; as $\rho \rightarrow -\infty$, q approaches $\min_j s_j$. Illustrative control components include verification coverage, permissioning and identity for computational agents, observability of agent actions, process redesign around delegation, and staff capability. The theory requires interacting components but does not depend on that particular taxonomy.

Conditional on delegation, incidents occur at a rate declining in q . The baseline intensity is $\lambda_0(1-q)x_tM$, where M is consequential-work volume and λ_0 is a technology-dependent base rate. A capability improvement lowers λ_0 and moves the boundary of the adoption region. It does not directly alter the organisation's position s or evidence stock E .

B.3 Stochastic model

The numerical model uses state $(T, x, E) \in [0, 1] \times [x_p, 1] \times [0, \infty)$ with s held fixed within each experiment except where stated. Let $\sigma(z) = (1 + e^{-z})^{-1}$ and let $\text{clip}(z, a, b) = \min\{b, \max\{a, z\}\}$. Each period,

$$N_t \sim \text{Poisson}(\lambda_0(1-q)x_tM), \quad V_t = \sum_{j=1}^{N_t} v_j + V_t^{\text{shock}}, \quad v_j \sim \text{Exp}(\text{mean } \bar{\sigma}), \quad (3)$$

$$D_t = 1 + c_E E_t, \quad (4)$$

$$T_{t+1} = \text{clip}\left(T_t + \alpha x_t(1 - T_t) - \frac{\beta}{D_t} T_t V_t - \omega(T_t - T_0), 0, 1\right), \quad (5)$$

$$E_{t+1} = \max\{E_t + x_t^2 - \epsilon_E E_t, 0\}, \quad (6)$$

$$x_t^* = x_p + (1 - x_p)\sigma(\kappa(T_t - \theta)), \quad (7)$$

$$r_t = \delta_{\text{down}} + (\delta_{\text{up}} - \delta_{\text{down}})\sigma(k_s(x_t^* - x_t)), \quad (8)$$

$$x_{t+1} = \text{clip}(x_t + r_t(x_t^* - x_t), x_p, 1). \quad (9)$$

The interpolation r_t is smooth on the interior and approximates fast retrenchment when $x_t^* < x_t$ and slow expansion when $x_t^* > x_t$. The deterministic skeleton replaces V_t by its expectation $\lambda_0(1-q)x_tM\bar{\sigma}$. An alternative accrual specification replaces x_t^2 in Equation (6) by $x_t^2/(1 + \phi V_t)$ so that incident-heavy periods contribute less to the evidence stock.

B.4 Baseline parameterisation

Table 1. Baseline parameters used in all numerical experiments in the source manuscript.

Parameter	Value	Parameter	Value
K	5	λ_0	0.1
M	1000	$\bar{\sigma}$	0.1
α	0.2	β	0.3
ω	0.02	T_0	0.15
κ	12	θ	0.5
δ_{up}	0.15	δ_{down}	0.9
k_s	200	x_p	0.03
ϵ_E	0.02	c_E	0.15
ρ	-2	adoption criterion	$x > 0.5$ for 12 periods

Periods are nominally months. The parameter values are chosen to exhibit the mechanism, not estimated from organisations.

C

Exact transition geometry and proofs

C.1 Fixed-point locus

For $T \in (0, 1)$ define

$$\begin{aligned} A(T) &= \alpha x^*(T)(1 - T) - \omega(T - T_0), \\ B(T) &= \beta T \lambda_0 x^*(T) M \bar{\sigma}, \\ D(T) &= 1 + c_E \frac{x^*(T)^2}{\epsilon_E}, \\ Q(T) &= 1 - \frac{D(T)A(T)}{B(T)}. \end{aligned}$$

Here A is net confidence gain absent incidents, B scales expected incident loss, and D is the damping supplied by the equilibrium evidence stock.

Theorem C.1 (Exact fixed-point geometry). *Consider the deterministic skeleton of Equations (3)–(9) at fixed control adequacy q .*

(i) *A point (T, x, E) with $x \in (x_p, 1)$ is a fixed point if and only if*

$$x = x^*(T), \quad E = \frac{x^*(T)^2}{\epsilon_E}, \quad Q(T) = q.$$

(ii) *Suppose Q is monotone on $(0, T_a)$, (T_a, T_b) , and $(T_b, 1)$, with a strict local maximum $q^\uparrow = Q(T_a)$ and strict local minimum $q^\downarrow = Q(T_b)$, and suppose both outer branches span $(q^\downarrow, q^\uparrow)$. Then every $q \in (q^\downarrow, q^\uparrow)$ has at least three fixed points.*

(iii) *In the reduced flow*

$$\dot{T} = A(T) - (1 - q) \frac{B(T)}{D(T)},$$

obtained by placing x and E on their equilibrium relations, a fixed point with $Q'(T) > 0$ is asymptotically stable and one with $Q'(T) < 0$ is unstable. If $Q''(T) \neq 0$ at a fold, the fold is a saddle-node bifurcation. For the reduced discrete map, local asymptotic stability holds when

$$0 < Q'(T)C(T) < 2, \quad C(T) = \frac{B(T)}{D(T)}.$$

Proof. At a fixed point, Equation (6) gives $E = x^2/\epsilon_E$. Because the adjustment rate in Equation (8) is strictly positive, Equation (9) forces $x = x^*(T)$. Substituting these relations into Equation (5), with incident volume replaced by its expectation $\lambda_0(1 - q)xM\bar{\sigma}$, gives

$$A(T) - (1 - q) \frac{B(T)}{D(T)} = 0,$$

which is equivalent to $Q(T) = q$ because $B(T) > 0$ on the interior. This proves (i).

For (ii), Q is continuous. Each of the three monotone intervals contributes a solution of $Q(T) = q$ for q strictly between the local minimum and maximum, provided the two outer branches span that interval. The corresponding values of x and E follow from (i).

For (iii), write $g(T) = A(T) - (1 - q)C(T)$ with $C = B/D > 0$. At equilibrium, $A = (1 - q)C$. Since $Q = 1 - A/C$,

$$g'(T) = -Q'(T)C(T).$$

Thus the reduced flow is stable exactly where $g' < 0$, equivalently $Q' > 0$, and unstable where $Q' < 0$. At a fold, $g = g' = 0$, while $g'' = -Q''C \neq 0$ and $\partial g/\partial q = C > 0$, giving the saddle-node normal form. For the reduced discrete map $T_{t+1} = T_t + g(T_t)$, local asymptotic stability requires $|1 + g'| < 1$, or $-2 < g' < 0$, which becomes $0 < Q'C < 2$. ■

The fixed-point characterisation in part (i) is exact for the full three-dimensional deterministic skeleton; no time-scale approximation is used. The reduction enters only in the scalar stability statement. At the baseline parameters, the two folds are

$$q^\uparrow = 0.8788, \quad q^\downarrow = 0.7906.$$

Forward simulation of the full skeleton locates thresholds at 0.8789 and 0.7903, respectively. The adjustment speeds δ_{up} , δ_{down} , and k_s do not appear in Q , so they affect transients but not the fixed-point thresholds.

Corollary C.2 (The evidence stock generates the band). *Wherever $A(T) > 0$,*

$$\frac{\partial Q(T)}{\partial c_E} = -\frac{x^*(T)^2}{\epsilon_E} \frac{A(T)}{B(T)} < 0.$$

The displacement is proportional to $x^(T)^2$, so stronger evidence buffering moves the high-delegation branch more than the pilot branch. In the baseline model, the width $q^\uparrow - q^\downarrow$ increases monotonically with c_E and is zero at $c_E = 0$.*

Proof. Only $D(T) = 1 + c_E x^*(T)^2 / \epsilon_E$ depends on c_E . Differentiating $Q = 1 - DA/B$ at fixed T gives the displayed derivative. Since $x^*(T)$ is much larger on the adoption branch, the same increase in c_E produces a larger vertical displacement there. At $c_E = 0$ the baseline locus has no interior folds, as confirmed both analytically and by deterministic simulation. ■

C.2 Full-map stability at the baseline fixed points

At $q = 0.835$, the deterministic skeleton has fixed points

$$(T, x, E) = (0.258, 0.080, 0.32) \quad \text{and} \quad (0.704, 0.923, 42.6),$$

corresponding to low-delegation and adoption states. At a fixed point, $x = x^*$, so derivatives of r_t with respect to x vanish in the product $r_t(x_t^* - x_t)$ and the mid-rate $\bar{r} = (\delta_{\text{up}} + \delta_{\text{down}})/2$ applies. Let $v = \lambda_0(1 - q)M\bar{\sigma}$ and $D = 1 + c_E E$. The Jacobian in coordinates (T, x, E) is

$$J = \begin{pmatrix} 1 - \alpha x - \frac{\beta}{D}vx - \omega & \alpha(1 - T) - \frac{\beta}{D}Tv & \frac{\beta c_E}{D^2}Tv x \\ \bar{r}(1 - x_p)\kappa\sigma'(\kappa(T - \theta)) & 1 - \bar{r} & 0 \\ 0 & 2x & 1 - \epsilon_E \end{pmatrix}.$$

For a characteristic polynomial $z^3 + a_2z^2 + a_1z + a_0$, the Jury criteria evaluated in the source manuscript are:

Criterion	Low-delegation	Adoption
$P(1) > 0$	5.5×10^{-4}	2.0×10^{-3}
$P(-1) < 0$	-5.61	-5.05
$ a_0 < 1$	0.423	0.337
$ a_1 - a_0a_2 < 1 - a_0^2$	$0.797 < 0.821$	$0.790 < 0.886$

All four hold. The spectral radii are 0.983 and 0.986. The dominant eigenvalue is close to $1 - \epsilon_E = 0.98$, reflecting slow evidence-stock decay rather than proximity to a bifurcation.

C.3 Exact observability of the latent state

Theorem C.3 (Exact observability). *Let*

$$g^* = \max_{u>0} [\sigma(u) + u\sigma'(u)] \approx 1.0998.$$

If

$$\frac{\delta_{\text{down}}}{\delta_{\text{up}}} < \frac{g^*}{g^* - 1} \approx 11.0,$$

then the one-period delegation update is strictly increasing in the target delegation x_t^ . Along any trajectory on which clipping is inactive, $T_t \in (0, 1)$, $x_t \in (x_p, 1)$, and $V_t > 0$, three consecutive observations of x together with the incident record identify (T_t, E_t) exactly.*

Proof. For fixed x , write $u = k_s(x^* - x)$ and

$$h(x^*) = x + [\delta_{\text{down}} + (\delta_{\text{up}} - \delta_{\text{down}})\sigma(u)](x^* - x).$$

Then

$$h'(x^*) = \delta_{\text{down}} + (\delta_{\text{up}} - \delta_{\text{down}})[\sigma(u) + u\sigma'(u)].$$

For $u \leq 0$, the bracket is less than $1/2$, so $h' > (\delta_{\text{up}} + \delta_{\text{down}})/2 > 0$. For $u > 0$,

$$h'(x^*) \geq \delta_{\text{down}} - (\delta_{\text{down}} - \delta_{\text{up}})g^*,$$

which is positive under the stated ratio bound. Strict monotonicity makes x_t^* recoverable from (x_t, x_{t+1}) . Equation (7) then gives

$$T_t = \theta + \kappa^{-1} \text{logit} \left(\frac{x_t^* - x_p}{1 - x_p} \right).$$

Applying the same inversion to (x_{t+1}, x_{t+2}) yields T_{t+1} . From Equation (5),

$$D_t = \frac{\beta T_t V_t}{T_t + \alpha x_t (1 - T_t) - \omega (T_t - T_0) - T_{t+1}},$$

and Equation (4) gives $E_t = (D_t - 1)/c_E$. ■

At the baseline retrenchment-to-expansion ratio of six, the condition holds with margin. At ratio twelve, it fails and delegation behaviour need not be invertible. Thus the asymmetry that makes regression abrupt can, if too extreme, make the latent state unidentifiable from behaviour alone. The empirical filtering problem is the noisy, parameter-uncertain version of this exact inversion.

C.4 Why collapse is not classical critical slowing down

At the high-delegation adoption fixed point, the dominant eigenvalue is governed mainly by slow evidence decay and changes little as q approaches $q^\downarrow$. A severe incident can push the state across the basin boundary while local restoring forces remain finite. The model therefore predicts rising variance as the basin shrinks but not necessarily the rising lag-1 autocorrelation characteristic of quasi-static passage through a fold [53]. This distinction narrows the early-warning claim in Appendix E.

D Numerical illustrations and robustness

D.1 Experiment designs

The source manuscript reports five experiments under one parameter set and no per-experiment tuning:

E1: deterministic hysteresis sweep.

Sweep $q \in [0.5, 1]$ from low-delegation initial conditions $(T_0, x_p, 0)$ and from adoption-state initial conditions $(0.9, 0.9, 0.81/\epsilon_E)$ until convergence.

E2: matched-mean configurations.

Simulate 1,200 organisations whose control vectors have arithmetic mean exactly 0.90 but differ in balance, for 360 periods from low-delegation initial conditions.

E3: incident regression.

Simulate 200 organisations in the adoption state at $q = 0.835$ for 480 periods, injecting at period 240 a severity-30 shock, 300 times the mean routine incident. Compare a surviving-stock arm, a destroyed-stock arm, and 200 no-shock controls.

E4: equifinality.

Simulate 300 organisations with heterogeneous initial controls, heterogeneous aspiration caps, and randomised sequential investment paths over 360 periods.

E5: early warning.

Simulate 200 organisations in the adoption state while ramping q from 0.95 to 0.70 over 600 periods, comparing detrended forty-period windows far from and immediately before exit.

D.2 Reported results

E2 is deliberately designed so that an additive aggregate cannot discriminate: every organisation has the same arithmetic mean. It therefore demonstrates a mechanism-level separation, not a statistical comparison in observational data. The flat profile likelihood cautions against claiming precise recovery of the CES elasticity. E3 contains both an analytic and a simulated component. If the stock is reset to zero at a q inside the baseline band, re-entry requires crossing $q^\uparrow$ and is therefore impossible at unchanged q ; the recovery rate in the surviving-stock arm is the stochastic output.

D.3 Band-width comparative static

The agreement is to the third decimal across the sweep. At $c_E = 0$, entry and exit coincide at approximately $q = 0.941$ in simulation.

Table 2. Numerical results reported in the source manuscript. They illustrate sufficiency of the mechanism; they are not estimates of organisational magnitudes.

Experiment	Result
E1	Analytic folds $q^\uparrow = 0.8788$ and $q^\downarrow = 0.7906$; simulated thresholds 0.8789 and 0.7903. Near the entry threshold, time to transition rises sharply: 45 periods at $q = 0.99$, 111 at 0.92, 200 at 0.90, 332 at 0.89, and 507 at 0.885.
E2	Despite identical arithmetic means, 41.9% reach the adoption state. Rates rise from 27% in the weakest-component tercile to 56% in the strongest. An additive classifier is at chance by construction (AUC 0.510); the data-generating CES aggregate reaches AUC 0.671. The additive value $\rho = 1$ is 55.5 log-likelihood points below the maximum, but the likelihood is nearly flat over $\rho \in [-8, 0.5]$, so the direction is identified more strongly than the elasticity.
E3	All 200 shocked organisations regress and all 200 unshocked controls remain in the adoption state. When the evidence stock is destroyed, none recover inside the horizon. When it survives, 93% recover, with a median of 63 periods below the threshold.
E4	112 of 300 organisations reach the adoption state. Adopter control vectors are about as dispersed as non-adopter vectors (mean pairwise dispersion 0.121 versus 0.126), while the weakest component is higher among adopters (0.887 versus 0.869).
E5	Variance of detrended delegation fluctuations rises by a factor of 105 in the forty periods before collapse, while lag-1 autocorrelation is close to unchanged (0.79 versus 0.80).

Table 3. Hysteresis-band width against evidence buffering. Analytic values are fold ordinates of Q ; simulated values come from the deterministic skeleton.

c_E	0.00	0.05	0.10	0.15	0.20	0.30
Analytic width	0.000	0.004	0.042	0.088	0.137	0.239
Simulated width	0.000	0.005	0.043	0.089	0.138	0.239

D.4 Robustness variants

When evidence accrual is damped by incident flow with $\phi = 0.3$, the band narrows to $[0.836, 0.881]$. Repeating E3 at $q = 0.858$, 200 of 200 organisations regress; 58% recover when the stock survives, with median 71 periods below threshold, and none recover when it is destroyed. Thus the qualitative structure survives while magnitudes shift.

Sweeping the retrenchment-to-expansion ratio over 2, 4, 6, 9, and 12 leaves both thresholds unchanged, as Theorem C.1 requires. Transients vary: recovery within the horizon declines from 100% at ratio two to 66% at ratio twelve, and a pure confidence excursion fails to become a full regression at ratio two because delegation retreats slowly enough for confidence to recover first. Ratio twelve lies just beyond the observability bound in Theorem C.3.

E Hypotheses, measurement, and empirical identification

E.1 Full hypothesis set

Each hypothesis is conditional on a population and measurement design capable of observing the relevant state signature.

H1: Capability events do not relocate organisations.

Shared capability improvements, such as frontier-model releases or price reductions, should predict transitions weakly relative to the organisation's prior control configuration. Transition hazard should load on configuration rather than cluster tightly around the capability calendar. The hypothesis is disconfirmed if heterogeneous organisations transition together around shared releases after conditioning on controls.

H2: Interactions dominate.

Models with interaction structure, whether the CES form with $\rho < 1$ or flexible equivalents, should outperform additive maturity scores out of sample. The weakest component should carry predictive information not captured by the mean. The claim is directional: the theory requires rejection of the additive case, not precise recovery of one elasticity. It is disconfirmed if additive models match interaction models or if $\hat{\rho}$ is indistinguishable from one in adequately powered data.

H3: Threshold behaviour.

In a population whose control-adequacy distribution places meaningful mass across the transition band, the distribution of sustained delegation should be bimodal, with mass near a pilot floor and in the adoption region. Time to transition should increase sharply near the entry threshold. The population condition matters: a sample entirely above both thresholds may be unimodal-high without contradicting the mechanism. The hypothesis is disconfirmed by a

continuously filled, unimodal delegation distribution in a population known to span the boundary.

H4: Incident-triggered regression and conditional recovery.

Severe incidents should trigger sustained regression at unchanged audited controls; regression should be faster than the expansion that preceded it; and recovery should depend on survival of the institutional evidence stock. The source manuscript pre-commits to treating the mechanism as disconfirmed if destroyed-record organisations recover at unchanged controls at least half as often as surviving-record organisations, because the baseline model places the former near zero and the latter in a large majority.

H5: Configurational equifinality.

Organisations in the adoption state should not share a unique control configuration. Adopter configurations may be about as dispersed as non-adopter configurations, while adopters share a higher floor on their weakest component or an equivalent complementary aggregate. The hypothesis is disconfirmed if adoption concentrates in one narrow configuration or if adopters' weakest components are indistinguishable from those of non-adopters at matched means.

A secondary early-warning prediction follows from the shock-driven exit mechanism: variance of detrended delegation fluctuations should rise before regression, but lag-1 autocorrelation need not. Because eroding q also raises incident forcing in the baseline model, the variance increase combines a shrinking basin with stronger shocks and should not be interpreted as a pure critical-slowness signature.

E.2 Measurement protocol

The dependent variable is Equation (1). A workable data construction requires the following fields: Enterprise orchestration

Table 4. Illustrative data dictionary for measuring enterprise delegation.

Construct	Observable fields	Pre-registration requirement
Task universe W	Task class, workflow identifier, timestamp, organisational unit	Define consequential work before outcome analysis
Human ratification	Approval event, pre-approval policy, human touch before effect, exception override	Specify what counts as prior ratification for each task class
Consequence weight w_i	Transaction value, record sensitivity, legal commitment, customer or operational exposure	Publish weighting rule and sensitivity analyses
Delegation share x	Weighted autonomous instances divided by total weighted instances	Fix accounting window and missing-data rule
Adoption	$x \geq x_A$ over a sustained window	Fix x_A , window length, and normal incident-flow criterion
Controls s_j	Permissioning, observability, verification, process design, staff capability, audit coverage	Fix taxonomy and scoring protocol
Incidents	Severity, timing, affected task class, rollback, financial or operational consequence	Fix severity scale and event-inclusion rule
Evidence survival	Availability and organisational credibility of decision records, outcome histories, post-mortems, and audit trails after the incident	Define survival independently of recovery outcome

and governance platforms can in principle provide the relevant action, approval, rollback, and incident records; vendor usage data already distinguish automation-shaped from augmentation-shaped usage at the API level [2]. Survey instruments cannot substitute for the order parameter because binary AI-use questions are sensitive to wording and do not distinguish advisory use from consequential delegation [1]. Early firm-level work on AI delegation and oversight suggests the construct is becoming measurable [15, 62].

E.3 Identification designs

Longitudinal platform panel. Construct monthly delegation shares and quarterly control audits for organisations on a common platform. H2 can be tested by nested or flexible interaction models, profile-likelihood intervals for ρ , and held-out prediction. H3 can be tested through multimodality statistics, mixture diagnostics, and configuration-dependent transition hazards. H5 can be tested by comparing pairwise configuration dispersion and weakest-component distributions among adopters and non-adopters. The main challenge is distinguishing genuine bistability from persistent mover-stayer heterogeneity. Within-organisation evidence is decisive: the same organisation remaining stable at a q below its historical entry point is more diagnostic than a cross-sectional mixture.

Incident event studies. Use severity-graded events from internal incident registers, insurer claims, or public repositories. Compare delegation trajectories in affected organisations with matched controls and estimate the interaction between incident exposure and evidence-stock survival. The central prediction is the interaction itself; a negative incident effect alone does

not test the mechanism. Post-incident reorganisations create both substantive variation and measurement risk: they may destroy the evidence stock while also changing staff, process, or formal controls.

Shared capability events and natural experiments. Frontier-model releases and common price changes affect technology parameters broadly and can test whether transitions cluster around capability events or depend on prior organisational configuration. Regulatory timetables can shift the cost or timing of specific control investments across sectors and jurisdictions, providing candidate instruments for configuration. Any use of current legal timetables should be re-verified at the time of empirical execution; the source manuscript's 2026 policy statements are not themselves an identification result.

Prospective early-warning study. Track detrended delegation fluctuations in organisations near regression and compare them with matched organisations that remain in the adoption state. Test variance and lag-1 autocorrelation separately. A finding that autocorrelation, rather than variance, provides the reliable signal would revise the shock-driven mechanism even if regression remains state-like.

E.4 Feasibility and threats to validity

No public dataset currently contains the complete order parameter and control history. Three access paths are plausible: single-vendor research partnerships; multi-vendor consortia through industry bodies, insurers, or regulators; and population-scale business surveys containing behavioural delegation modules. The last cannot reproduce workflow-level x exactly but may detect bimodality through pre-approval requirements, autonomous transaction shares, and rollback behaviour.

The main threats are:

1. **Reverse causality.** Delegation can induce control investment as well as result from it. Longitudinal timing, regulatory shocks, and incident studies are preferable to cross-section alone.
2. **Selection.** Platform records sample organisations that adopted a particular platform. Population survey modules and multi-vendor designs are needed to assess transportability.
3. **Construct validity.** Without a pre-registered consequential-work taxonomy, the order parameter collapses into another measure of general AI use.
4. **Unobserved heterogeneity.** A mixture of intrinsically different organisation types can mimic cross-sectional bimodality. History dependence within the same organisation is the key discriminator.
5. **Auxiliary functional forms.** Logistic authorisation, superlinear accrual, and the chosen damping function are model conveniences. Failure of one form should revise the model without automatically rejecting the broader state-transition claim; failure of the core dynamical signatures should reject it.

Pilot abandonment must not be used as a direct proxy for regression from the adoption state. The model begins at a positive pilot floor and contains no decision to stop piloting; its object is exit after sustained consequential delegation.

F Implications, limitations, and future programme

F.1 Theoretical implications

If enterprise AI adoption is maintained by feedback, binary AI-use variables and single maturity scores omit the objects governing dynamics. The relevant quantities are the delegation order parameter, the transition boundary, basin membership, and the evidence stock. The historical analogy with electrification and information technology becomes more precise: coordinated complements are required to cross a boundary, so partial investment can appear unproductive until the transition occurs. The present transition may differ from capital-intensive predecessors because institutional confidence can be erased by an incident even when installed technical capability remains. The historical comparison is a conjecture, not a demonstrated result.

The framework also separates capability from organisational position. A model release can reduce the base failure rate and move the boundary without changing controls, evidence, or confidence. This gives a reason for shared technical progress to yield heterogeneous organisational outcomes without assuming that some organisations lack access to the technology.

F.2 Practical implications

The practical conclusions follow from the model geometry and are not yet empirical prescriptions.

1. At matched budgets, investment in a weak complementary control can dominate further improvement in an already strong component. Averaged scorecards may measure the wrong moment of the control distribution.
2. An organisation just above its entry threshold can take much longer to reach the adoption state than one well above it. A long flat period need not imply that coordinated investment has no value, though the model does not identify the optimal investment policy.

3. Decision records, audited outcome histories, calibrated expectations, and post-mortem records should persist independently of the vendor, system, team, or reorganisation that generated them. In the model, their durability determines the width of the hysteresis band and whether recovery after a severe incident remains feasible at unchanged controls.

The source model deliberately excludes organisational intervention. Once investment, process redesign, and incident response enter as control inputs, the problem changes from explaining state transitions to choosing policies that move or keep the organisation inside the adoption basin.

F.3 Limitations

The parameterisation is illustrative and no numerical magnitude in Appendix D is an estimate of a firm. The control taxonomy is not validated. The model concerns systems autonomous enough to fail consequentially and does not address ordinary drafting, summarisation, retrieval, or decision support under universal human ratification. The mechanism is sufficient, not claimed to be unique. Bistability in a stylised model does not establish bimodality in organisations.

The theory should be abandoned or substantially revised if appropriately measured data show that additive models predict as well as interaction models, sustained delegation is continuous rather than state-like in populations spanning the proposed threshold, within-organisation hysteresis is absent, incident regression does not occur at unchanged controls, or recovery does not interact with evidence-record survival. The explicit failure conditions prevent the latent-state language from insulating the account against evidence.

F.4 Future research programme

Four questions follow if the empirical signatures survive.

State estimation.

Estimate latent confidence and evidence from noisy histories of delegation, approvals, incidents, and recovery. Theorem C.3 solves only the noiseless known-parameter case; the practical problem is nonlinear filtering with parameter uncertainty, followed by validation through predicted dynamics [36].

Boundary estimation.

Estimate entry and exit thresholds and distinguish true hysteresis from unobserved types using panel threshold methods, transition hazards, and within-organisation history [32].

Organisational control.

Treat interventions as control inputs and study reachability of the adoption region, least-cost paths across the entry boundary, sequencing towards weak components, and incident-response policies that preserve basin membership.

Adaptive enterprise systems.

Compose state estimation, boundary location, intervention choice, and response observation into a closed loop. The open governance question is whether such a loop can operate safely, at what cadence, and with what division of authority between the observing system and the institution it models.

G

Reproducibility and restructuring record

The source manuscript states that code and a `results.json` file accompany the paper and that all numerical values are generated from one parameter set with fixed seeds. Those files were not included with the PDF supplied for this restructuring. This revision therefore does not claim to reproduce the original stochastic experiments independently.

The revision package includes:

- one consolidated LaTeX source containing the main body and technical appendices;
- the bibliography used by the restructured manuscript;
- a vector mechanism-and-geometry figure;
- the data file used to draw the baseline fixed-point locus; and
- a Python script that recomputes the analytic locus, fold locations, and band-width table from the equations and baseline parameters.

The stochastic experiment values in Table 2 are retained from the supplied source manuscript. A submission version should reconnect the original simulation artefact, confirm that section and figure identifiers match the new structure, and record its repository identifier, version, licence, environment, seeds, and checksums.

References

- [1] Jeffrey S. Allen. Monitoring AI Adoption in the U.S. Economy. *FEDS Notes*, April 2026. doi:10.17016/2380-7172.4032. Board of Governors of the Federal Reserve System, 3 April 2026.
- [2] Ruth Appel, Peter McCrory, Alex Tamkin, Michael Stern, Miles McCain, and Tyler Neylon. Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption. Technical report, Anthropic, September 2025. URL <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>. First-party API usage telemetry; vendor report, not peer reviewed.
- [3] W. Brian Arthur. Competing Technologies, Increasing Returns, and Lock-In by Historical Events. *The Economic Journal*, 99(394): 116–131, 1989. doi:10.2307/2234208.
- [4] Alexander Bick, Adam Blandin, and David J. Deming. The Rapid Adoption of Generative AI. *Management Science*, January 2026. doi:10.1287/mnsc.2025.02523. Published online 20 January 2026; earlier circulated as NBER Working Paper 32966.
- [5] Kathryn Bonney, Cory Breaux, Catherine Buffington, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, and Keith Savage. Tracking Firm Use of AI in Real Time: A Snapshot from the Business Trends and Outlook Survey. Technical Report 32319, National Bureau of Economic Research, April 2024. doi:10.3386/w32319.
- [6] Erik Brynjolfsson. The Productivity Paradox of Information Technology. *Communications of the ACM*, 36(12):66–77, 1993. doi:10.1145/163298.163309.
- [7] Erik Brynjolfsson and Lorin M. Hitt. Beyond Computation: Information Technology, Organizational Transformation and Business Performance. *Journal of Economic Perspectives*, 14(4):23–48, 2000. doi:10.1257/jep.14.4.23.
- [8] Erik Brynjolfsson and Paul Milgrom. Complementarity in Organizations. In Robert Gibbons and John Roberts, editors, *The Handbook of Organizational Economics*, pages 11–55. Princeton University Press, 2013. doi:10.2307/j.ctt1r2ggg.5.
- [9] Erik Brynjolfsson, Lorin M. Hitt, and Shinkyu Yang. Intangible Assets: Computers and Organizational Capital. *Brookings Papers on Economic Activity*, 2002(1):137–198, 2002. doi:10.1353/eca.2002.0003.
- [10] Erik Brynjolfsson, Daniel Rock, and Chad Syverson. The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372, 2021. doi:10.1257/mac.20180386.
- [11] Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. Generative AI at Work. *The Quarterly Journal of Economics*, 140(2): 889–942, 2025. doi:10.1093/qje/qjae044.
- [12] Bernard Burnes. Complexity Theories and Organizational Change. *International Journal of Management Reviews*, 7(2):73–90, 2005. doi:10.1111/j.1468-2370.2005.00107.x.
- [13] Aditya Challapally, Chris Pease, Ramesh Raskar, and Pradyumna Chari. The GenAI Divide: State of AI in Business 2025. Technical report, MIT NANDA, July 2025. URL https://web.archive.org/web/20250818145714id_/https://nanda.media.mit.edu/ai_report_2025.pdf. Preliminary, non-peer-reviewed report; archived copy of the original first-party PDF; methodology criticised.
- [14] Wesley M. Cohen and Daniel A. Levinthal. Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly*, 35(1):128–152, 1990. doi:10.2307/2393553.
- [15] Innessa Colaiacovo and Rembrand Koning. Delegating to AI: Beliefs and the Organization of Work at New Ventures. Working Paper 26-093, Harvard Business School, June 2026. Dated 18 June 2026; doi:10.2139/ssrn.6962978.
- [16] Kelsey Adrienne Cua and Ezekiel Bernardo. A System Dynamics Decision-Support Model for Governing Generative AI Adoption: Leverage-Point Analysis and Intervention Simulation. Preprint SSRN 6968358, De La Salle University, June 2026. Posted 19 June 2026; doi:10.2139/ssrn.6968358.
- [17] Paul A. David. Clio and the Economics of QWERTY. *American Economic Review, Papers and Proceedings*, 75(2):332–337, 1985. URL <https://www.jstor.org/stable/1805621>.
- [18] Paul A. David. The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *American Economic Review, Papers and Proceedings*, 80(2):355–361, 1990. URL <https://www.jstor.org/stable/2006600>.
- [19] Fabrizio Dell’Acqua, Edward III McFowland, Ethan Mollick, Hila Lifshitz, Katherine C. Kellogg, Saran Rajendran, Lisa Kraye, François Candelon, and Karim R. Lakhani. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science*, 37(2):403–423, March 2026. doi:10.1287/orsc.2025.21838.
- [20] Berkeley J. Dietvorst, Joseph P. Simmons, and Cade Massey. Algorithm Aversion: People Erroneously Avoid Algorithms after Seeing Them Err. *Journal of Experimental Psychology: General*, 144(1):114–126, 2015. doi:10.1037/xge0000033.
- [21] Avinash Dixit. Entry and Exit Decisions under Uncertainty. *Journal of Political Economy*, 97(3):620–638, 1989. doi:10.1086/261619.

- [22] Avinash Dixit. Investment and Hysteresis. *Journal of Economic Perspectives*, 6(1):107–132, 1992. doi:10.1257/jep.6.1.107.
- [23] Kevin J. Dooley and Andrew H. Van de Ven. Explaining Complex Organizational Dynamics. *Organization Science*, 10(3):358–372, 1999. doi:10.1287/orsc.10.3.358.
- [24] Ari Ercole. The Partial Adoption Trap: Coordination Failure, Trust, and Cultural Lock-in in Health AI Adoption. [arXiv:2605.17388](https://arxiv.org/abs/2605.17388), May 2026. Submitted 17 May 2026.
- [25] Joshua S. Gans and Avi Goldfarb. O-Ring Automation. Technical Report Working Paper 34639, National Bureau of Economic Research, January 2026. doi:10.3386/w34639.
- [26] Ji Gao and John D. Lee. Extending the Decision Field Theory to Model Operators’ Reliance on Automation in Supervisory Control Situations. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 36(5):943–959, 2006. doi:10.1109/TSMCA.2005.855783.
- [27] Gartner. Gartner Predicts 30% of Generative AI Projects Will Be Abandoned After Proof of Concept By End of 2025, July 2024. URL <https://www.gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025>. Press release, 29 July 2024; analyst forecast, methodology unpublished.
- [28] Gartner. Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027, June 2025. URL <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>. Press release, 25 June 2025; analyst forecast, methodology unpublished.
- [29] Connie J. G. Gersick. Revolutionary Change Theories: A Multilevel Exploration of the Punctuated Equilibrium Paradigm. *Academy of Management Review*, 16(1):10–36, 1991. doi:10.2307/258605.
- [30] Nicole Gillespie and Graham Dietz. Trust Repair After An Organization-Level Failure. *Academy of Management Review*, 34(1):127–145, 2009. doi:10.5465/amr.2009.35713319.
- [31] Christopher Gresov and Robert Drazin. Equifinality: Functional Equivalence in Organization Design. *Academy of Management Review*, 22(2):403–428, 1997. doi:10.2307/259328.
- [32] Bruce E. Hansen. Sample Splitting and Threshold Estimation. *Econometrica*, 68(3):575–603, 2000. doi:10.1111/1468-0262.00124.
- [33] Paul A. Herbig. A Cusp Catastrophe Model of the Adoption of an Industrial Innovation. *Journal of Product Innovation Management*, 8(2):127–137, 1991. doi:10.1111/1540-5885.820127.
- [34] Kath Houchin and Donald MacLean. Complexity Theory and Strategic Change: An Empirically Informed Critique. *British Journal of Management*, 16(2):149–166, 2005. doi:10.1111/j.1467-8551.2005.00427.x.
- [35] Lingxiao Huang and Nisheeth K. Vishnoi. Path Dependence under Adaptive AI Delegation. [arXiv:2603.02950](https://arxiv.org/abs/2603.02950), March 2026. Submitted 3 March 2026; revised 7 May 2026.
- [36] Rudolf E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35–45, 1960. doi:10.1115/1.3662552.
- [37] Daniel Katz and Robert L. Kahn. *The Social Psychology of Organizations*. John Wiley & Sons, Inc., New York, 1966. URL <https://search.worldcat.org/title/The-social-psychology-of-organizations/oclc/1470995928>.
- [38] Michael Kremer. The O-Ring Theory of Economic Development. *The Quarterly Journal of Economics*, 108(3):551–575, 1993. doi:10.2307/2118400.
- [39] Philipp Laut, Philipp Dumbach, and Bjoern M. Eskofier. Integration of Artificial Intelligence in the Organizational Adoption—A Configurational Perspective. In *ICIS 2021 Proceedings*, December 2021. URL https://aisel.aisnet.org/icis2021/is_implement/is_implement/2/. Paper 1165, 12 December 2021.
- [40] John D. Lee and Katrina A. See. Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, 46(1):50–80, 2004. doi:10.1518/hfes.46.1.50.30392.
- [41] Daniel A. Levinthal. Adaptation on Rugged Landscapes. *Management Science*, 43(7):934–950, 1997. doi:10.1287/mnsc.43.7.934.
- [42] Junming Liu, Wei Thoo Yue, Alvin Chung Man Leung, and Xin Zhang. Find the Good. Seek the Unity: A Hidden Markov Model of Human-AI Delegation Dynamics. *MIS Quarterly*, 49(3):1185–1204, 2025. doi:10.25300/MISQ/2024/18232.
- [43] Christoph H. Loch and Bernardo A. Huberman. A Punctuated-Equilibrium Model of Technology Diffusion. *Management Science*, 45(2):160–177, 1999. doi:10.1287/mnsc.45.2.160.
- [44] Lili Mi, Wei Liu, Yu-Hsi Yuan, Xuefeng Shao, and Yifan Zhong. Adoption of AI in Response to COVID-19—A Configurational Perspective. *Personal and Ubiquitous Computing*, 27(4):1455–1467, 2023. doi:10.1007/s00779-023-01711-6.
- [45] Paul Milgrom and John Roberts. The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *American Economic Review*, 80(3):511–528, 1990. URL <https://www.jstor.org/stable/2006681>.

- [46] Paul Milgrom and John Roberts. Complementarities and Fit: Strategy, Structure, and Organizational Change in Manufacturing. *Journal of Accounting and Economics*, 19(2–3):179–208, 1995. doi:[10.1016/0165-4101\(94\)00382-F](https://doi.org/10.1016/0165-4101(94)00382-F).
- [47] Manjana Milkoreit. Social Tipping Points Everywhere?—Patterns and Risks of Overuse. *WIREs Climate Change*, 14(2):e813, 2023. doi:[10.1002/wcc.813](https://doi.org/10.1002/wcc.813).
- [48] Raja Parasuraman and Victor Riley. Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors*, 39(2):230–253, 1997. doi:[10.1518/001872097778543886](https://doi.org/10.1518/001872097778543886).
- [49] Jan W. Rivkin. Imitation of Complex Strategies. *Management Science*, 46(6):824–844, 2000. doi:[10.1287/mnsc.46.6.824.11940](https://doi.org/10.1287/mnsc.46.6.824.11940).
- [50] Elaine Romanelli and Michael L. Tushman. Organizational Transformation as Punctuated Equilibrium: An Empirical Test. *Academy of Management Journal*, 37(5):1141–1166, 1994. doi:[10.2307/256669](https://doi.org/10.2307/256669).
- [51] Jim Rowan, Nitin Mittal, Beena Ammanath, and Costi Perricos. The State of AI in the Enterprise: The Untapped Edge. Technical report, Deloitte AI Institute, January 2026. URL <https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html>. Survey of 3,235 leaders in 24 countries; industry report, not peer reviewed.
- [52] Marten Scheffer, Steve Carpenter, Jonathan A. Foley, Carl Folke, and Brian Walker. Catastrophic Shifts in Ecosystems. *Nature*, 413(6856):591–596, 2001. doi:[10.1038/35098000](https://doi.org/10.1038/35098000).
- [53] Marten Scheffer, Jordi Bascompte, William A. Brock, Victor Brovkin, Stephen R. Carpenter, Vasilis Dakos, Hermann Held, Egbert H. van Nes, Max Rietkerk, and George Sugihara. Early-Warning Signals for Critical Transitions. *Nature*, 461(7260):53–59, 2009. doi:[10.1038/nature08227](https://doi.org/10.1038/nature08227).
- [54] Marten Scheffer, Stephen R. Carpenter, Timothy M. Lenton, Jordi Bascompte, William Brock, Vasilis Dakos, Johan van de Koppel, Ingrid A. van de Leemput, Simon A. Levin, Egbert H. van Nes, Mercedes Pascual, and John Vandermeer. Anticipating Critical Transitions. *Science*, 338(6105):344–348, 2012. doi:[10.1126/science.1225244](https://doi.org/10.1126/science.1225244).
- [55] Alex Singla, Alexander Sukharevsky, Bryce Hall, Lareina Yee, Michael Chui, and Tara Balakrishnan. The State of AI in 2025: Agents, Innovation, and Transformation. Technical report, QuantumBlack, AI by McKinsey, November 2025. URL <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>. Global survey published 5 November 2025; industry report, not peer reviewed.
- [56] S&P Global Market Intelligence. Generative AI Experiences Rapid Adoption, but with Mixed Outcomes—Highlights from VotE: AI & Machine Learning. Technical report, S&P Global Market Intelligence, May 2025. URL <https://www.spglobal.com/market-intelligence/en/news-insights/research/ai-experiences-rapid-adoption-but-with-mixed-outcomes-highlights-from-vote-ai-machine-learning>. Research article published 30 May 2025; industry source, not peer reviewed.
- [57] Barry M. Staw, Lance E. Sandelands, and Jane E. Dutton. Threat Rigidity Effects in Organizational Behavior: A Multilevel Analysis. *Administrative Science Quarterly*, 26(4):501–524, 1981. doi:[10.2307/2392337](https://doi.org/10.2307/2392337).
- [58] David J. Teece, Gary Pisano, and Amy Shuen. Dynamic Capabilities and Strategic Management. *Strategic Management Journal*, 18(7):509–533, 1997. doi:[10.1002/\(SICI\)1097-0266\(199708\)18:7<509::AID-SMJ882>3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z).
- [59] Eric L. Trist and K. W. Bamforth. Some Social and Psychological Consequences of the Longwall Method of Coal-Getting. *Human Relations*, 4(1):3–38, 1951. doi:[10.1177/001872675100400101](https://doi.org/10.1177/001872675100400101).
- [60] Michael L. Tushman and Elaine Romanelli. Organizational Evolution: A Metamorphosis Model of Convergence and Reorientation. In L. L. Cummings and Barry M. Staw, editors, *Research in Organizational Behavior*, volume 7, pages 171–222. JAI Press, Greenwich, CT, 1985. URL https://www.hbs.edu/ris/Profile%20Files/Tushman%20Vita%20with%20cases%20August%202025_c7de17b9-530d-46d1-9022-36fc6f2112da.pdf.
- [61] Ludwig von Bertalanffy. The Theory of Open Systems in Physics and Biology. *Science*, 111(2872):23–29, 1950. doi:[10.1126/science.111.2872.23](https://doi.org/10.1126/science.111.2872.23).
- [62] Emma Wiles, Megan Hsu, Julie Bedard, and Matthew Kropp. Putting AI on the Org Chart: Evidence on Delegation and Oversight. Technical report, Boston University and Boston Consulting Group, July 2026. URL https://www.emmawiles.com/storage/ai_employee.pdf. Working paper dated 17 July 2026.