DIAGNOSTIC AND REPLACEMENT SERIES · PAPER 7

THE GOVERNED CORRECTION SEQUENCE

Why Correction Frameworks in High-Consequence Domains Converge on the Same Structure — and Where They Structurally Stop

The Meta-Layer

Jamie Forrester

Independent Systems Architect · Edinburgh, UK

hello@jamieforrester.com · jamieforrester.com

April 2026 · v3.0 · Working Paper

Correction without a frame-level gate is review-rich, exit-poor.

Abstract

A correction framework enters the failure class this paper names when its sequence contains gates for object-level decisions — fix or rollback, operate or don't operate, charge or decline — but lacks a formal mechanism for the frame-level equivalent: should the system performing the correction itself be replaced? This paper names that failure class, identifies the structural sequence all high-consequence correction frameworks converge upon, and proves that every known instantiation stalls at the same move — Move 4 at the frame level — for the same structural reason Paper 3 already established: correction loops built from within a frame cannot generate the question that would require that frame to change.

The paper introduces four instruments: the Governed Correction Sequence (a ten-move structural sequence), the Frame-Level Gate distinction (object versus frame at Move 4), the Correction Decay Spectrum (five observable stages of mismatch progression), and the Frame-Gate Test (a ten-question diagnostic instrument). Three primary domain proofs — medicine, aviation, and legal prosecution — are followed by software/SRE and a pre-modern transfer test (manuscript culture and print) that confirms structural depth beneath domain vocabulary. The paper closes with a falsifiable prediction: any correction framework without Move 4 at the frame level will produce corrections that hold temporarily and fail in the same structural way the original failure failed. Five NHS reform cycles and four financial regulation cycles provide confirmatory pattern evidence.

For Institutional Leaders — Read This First

You are in the failure class this paper names if your system meets three conditions: it has a formal correction framework for recurring failure; that framework contains a gate for object-level decisions (fix or rollback, operate or don't operate, charge or decline); and it contains no formal mechanism for the frame-level equivalent — should the architecture itself be replaced?

The dashboard is not the problem. The audit is not the problem. The review cycle is not the problem. These are downstream of the break. The break is at Move 4 — the gate between correcting within the current frame and asking whether the frame should change. Without it, every correction routes back into the same failure-producing architecture.

Three questions reveal whether the gate is absent:

- When the same failure returns after correction, what formally happens next?
- Who has authority to ask whether the frame — not the instance — should change?
- Where is the ghost correction structure forming — and who is absorbing the burden the missing gate should be routing?

If no one formally holds that authority, the gate is absent by design, not by accident. Completing the sequence does not guarantee replacement — admissibility (Paper 6) remains a separate threshold. But without the gate, replacement never initiates. The Frame-Gate Test in Section XIII provides the diagnostic instrument. The Correction Decay Spectrum in Section X names the five stages of progression.

Series Position

This is Paper 7 in the Diagnostic and Replacement Series. Papers 1–3 extracted the diagnostic laws: interface-legitimacy mismatch, function collapse under scale, and why correction loops cannot generate redesign from within the same frame. Papers 4–6 named the replacement mechanics: burden transfer, the five-stage replacement pipeline, and admissibility. Paper 7 names what all six papers are instances of — the Governed Correction Sequence, the structural sequence every high-consequence correction framework independently converges upon, and the point at which every examined instantiation structurally stalls. Paper 8 applies the complete toolkit to a new failure class: externally imposed primitive compression. Papers 1–6 name failures internal to the system or its correction loops. Paper 7 names the sequence those failures traverse. Paper 8 names what happens when that sequence is blocked from outside.

Reader Map

Reading Path	Sections
Full read — researchers, institutional designers, system architects	All sections and appendices
Executive read — CTOs, heads of reliability, policymakers	For Institutional Leaders box, then Sections I, IV, V, X, XIII, XIV
Series reader	Sections I and XIV only

Section I The Framework Name Problem

PDCA was invented in manufacturing. OODA emerged from military combat operations. DMAIC from semiconductor production. SBAR from acute clinical communication. The Toyota A3 from post-war Japanese quality engineering. The SRE incident response cycle from running internet infrastructure at scale. Six frameworks. Six domains. Six independent vocabularies.

They are treated as different methodologies and competing approaches. That treatment is a category error.

Strip the vocabulary — Plan-Do-Check-Act, Observe-Orient-Decide-Act, Define-Measure-Analyse-Improve-Control — and the remaining skeleton is identical: failure is captured, ambiguity is reduced, a correction is designed and executed, the result is verified, and the cycle restarts.

The point is not that the frameworks are identical in expression. They are not. The point is that they are performing the same kind of work: moving a system from visible failure toward controlled correction. Their surface language differs because their domains differ. Their structural sequence does not.

Before asking which framework is better, a prior question must be asked: is the sequence those frameworks share complete? If there is a structural move the sequence requires but none of the frameworks contain — not through oversight, but through a structural limitation Paper 3 already proved — then the debate about which framework is best is asking the wrong question.

The domain names the field. The sequence names the structure. The gap names the failure.

Section II The Promise Correction Frameworks Make

Correction frameworks are encountered first as promises. When an organisation adopts PDCA, DMAIC, OODA, A3, or an SRE incident model, it is not committing to a process diagram. It is accepting a claim: when something fails, the system will notice, understand it, correct it, and prevent recurrence.

That promise has four parts.

Detection. Relevant signals, anomalies, exceptions, recurring burdens, and weak warnings will be captured before they accumulate into something irrecoverable.

Identification. Causes will be correctly located — not the most visible symptom, not whatever surfaced most recently, but the structural source producing the recurring pattern.

Legitimacy. An intervention will be designed, authorised, and carried into the world in a form capable of changing the producing condition.

Hold. What was changed will remain changed under pressure. The system will not drift back toward the same failure, and if new failure emerges, the cycle will re-enter cleanly.

Every framework makes this promise in its own vocabulary. Check in PDCA. Control in DMAIC. Act in OODA. The terminal word in every framework name is the hold promise.

The question is whether the sequences those frameworks provide are structurally capable of keeping the fourth part. A promise of hold made by a sequence that cannot produce hold is not a failure of sincerity. It is a structural insufficiency in the sequence itself.

Section III The Observable Reality

Across domains, the operational loop of correction follows the same movement. A signal enters: a symptom, an alert, a complaint, an anomaly, a recurring incident, a pattern that will not resolve. Ambiguity is reduced. The signal is classified, scoped, and investigated until a structural cause is located rather than a surface presentation. An intervention is selected, executed under consequence, and verified. The system is monitored after release. If failure recurs, the signal re-enters.

The surface vocabulary differs. The structural movement does not.

But there is a second observable pattern alongside the operational one. Institutions under correction pressure intensify the tools they already have without removing the dependence that keeps people inside them. Critique accumulates. Reviews multiply. Dashboards expand. Policies are rewritten. Process grows denser.

What does not reliably happen is displacement of the failure-producing structure.

Systems can respond without resolving. They can become more legible to observers while remaining just as heavy for participants. They can increase documentary correction while failing to produce structural correction. Paper 4 named the result precisely: institutions become review-rich and exit-poor. The correction loops are functioning. They are functioning inside the frame that is generating the failure.

The systems are not inactive. They are active at the wrong layer.

Section IV The Governed Correction Sequence

Three layers — held separately

This paper works across three distinct layers. They must not be conflated.

Layer 1 — The descriptive archetype. What has been independently observed across high-consequence domains. Multiple correction frameworks in distinct fields converge on a common sequence. This is pattern recognition, not design.

Layer 2 — The structural limitation. Why that sequence cannot complete itself. Paper 3 established the mechanism: correction loops built for execution and diagnosis preserve the frame they operate inside. They can refine it, defend it, and densify it. They cannot generate the question that would require the frame to change. The absence of the frame-level gate is not a design omission. It is the structural consequence of running correction inside the thing being questioned.

Layer 3 — The completion layer. What a complete sequence requires: a formal gate between correction and replacement, and a governing layer above the loop that makes that gate structurally possible. This is this paper’s contribution — not a description of something that already exists, but a named structural absence and a specification of what fills it.

Conflate these three layers and the paper weakens immediately. The archetype is not the completion layer. The limitation explains why they differ.

The ten moves and their irreducibility

Every functioning correction sequence in a high-consequence domain performs the same ten moves — not optionally, but as a structural requirement of the work. If a system must receive failure signals, reduce ambiguity, act under consequence, validate the action, and maintain governability after release, these ten moves exist somewhere in its sequence, whether clearly separated or collapsed into one another. They cannot be removed without the function reappearing elsewhere in a less legible form.

Move	Name	Required Question	Failure Mode if Absent
1	Signal	What failure has entered the system?	Failure invisible; no correction cycle opens
2	Assess	What is the bounded object of correction?	Scope drifts; diagnosis is premature
3	Diagnose	What structural source is producing recurrence?	Correction addresses symptoms; recurrence guaranteed
4	Gate	Correct within the frame — or replace the frame?	All failures route to correction; structural failures recur indefinitely
5	Design	What is the specified correction or replacement?	Execution drifts from intent
6	Execute	Has it been built and deployed?	Design remains theoretical
7	Verify	Does output satisfy structural truth, not completion criteria?	Corrections that fail under real conditions are released
8	Deploy	Has it entered live operation?	Real-world conditions untested
9	Monitor	Is it watching for recapture, drift, and burden relocation?	Degradation undetected until failure recurs at scale
10	Restart	Does future failure re-enter from Signal?	Future failures treated as new events, not recurrence

The order is fixed because each move requires the output of the previous one. Signal must precede Assess because there is nothing to scope without a signal. Diagnose must precede Gate because the gate question cannot be answered without knowing the structural cause. Design must follow Gate because the design must respond to that decision. The dependency order is the logic of the work itself, not an arbitrary choice.

The Governed Correction Law

Where systems must correct recurring structural failure under ambiguity and consequence, multiple independent correction frameworks converge on the same ten-move sequence regardless of domain vocabulary. The sequence is sufficient for correction within a frame, but insufficient for determining whether the frame itself should persist. That insufficiency is structural. It was predicted by Paper 3 before this paper was written. Without a formal frame-level gate, structural failures return as improved versions of themselves.

When Move 4 at the frame level is formally reached and a structurally sound replacement is built, Substitution by Gravity becomes the mechanic that determines whether the replacement becomes institutionally real — addressed in Section XII. Completing the sequence does not by itself guarantee replacement. Whether a substitute becomes institutionally real remains a separate admissibility question addressed in Paper 6.

Section V The Missing Gate

Object-level versus frame-level — the precise distinction

Move 4 is present at the object level in every domain examined. Medicine has operate or don't operate. Legal prosecution has charge or decline to charge. SRE has fix or rollback. Aviation has ground, reroute, repair, or continue. Each is a real gate with real consequence.

What none contain is Move 4 at the frame level — a formal mechanism that asks whether the system performing the correction should itself be replaced. The operate-or-don't-operate decision asks whether this patient warrants intervention, not whether the health system's architecture should change. The charging decision asks whether this case should proceed, not whether the adversarial frame is the right structure for finding truth.

Systems can decide whether to act on an object. They cannot, from within the same correction loop, decide whether the system itself should be replaced.

Why the gate is structurally absent — Paper 3's mechanism

This is exactly the limitation Paper 3 predicted. Correction loops built for execution and diagnosis preserve the frame they operate inside — because they are built from that frame's categories, assumptions, interfaces, and success conditions. They can refine the frame. They can defend the frame. What they cannot do is generate the question that would require the frame to change.

The absence of the frame-level gate is not a design oversight by the creators of PDCA, DMAIC, or OODA. It is the predictable structural consequence of running correction inside the thing being questioned. Paper 3 identified four structural biases that produce this result.

Execution bias. Systems are organised to preserve flow, not suspend it. The frame is treated as a background condition because questioning it requires a kind of interruption the system is not built to host.

Defensibility bias. Records, reviews, reports, and dashboards accumulate around what was done and why it was justified — not around whether the structure requiring the action was

coherent. The system becomes better at defending itself. It does not become better at seeing itself.

Recognition bias. Systems can only route and legitimise the forms of work their interface can recognise. Redesign signals arrive at tools built for something else. They show up as complaint, delay, abstraction, or non-executable thinking. Move 4 has no native lane in any correction framework because the interface was built for execution and diagnosis, not for the question that would force the frame to change.

Incentive bias. People inside correction systems are rewarded for making the current machine function: clearing queues, restoring service, closing cases, hitting targets. The incentive structure selects for operational competence inside the frame, not for the attention that would make the frame visible as a problem.

Together, these four biases mean the gap at Move 4 is not a contingent omission. It is the expected output of how correction loops are constituted. The gap reproduces itself wherever correction is built from within the frame it would need to question.

Cross-framework gap analysis

Framework	M1–3	M4 (object)	M4 (frame)	M5–8	M9–10
PDCA	Partial	✗	✗	✓ correction	Partial
DMAIC	✓	✗	✗	✓ correction	Partial
OODA	Partial	Partial*	✗	✓ correction	Partial
Toyota A3	✓	✗	✗	✓ correction	Partial
Ford 8D	✓	✗	✗	✓ correction	Partial
SRE Incident	✓	Partial†	✗	✓ correction	✓
Governed Correction Sequence	✓	✓	✓	✓ correction or replacement	✓

*OODA’s Orient step involves frame-adjustment at the cognitive level but does not contain a formal gate capable of routing the system to replacement of its own frame.

†SRE’s fix-or-rollback is a real gate at the feature level. Architectural review triggered by postmortems exists in mature organisations but sits outside the incident loop as an informal downstream consequence, not a formal gate inside the sequence.

A note on carriers as transfer tests

The following proofs use carriers from distinct domains. A carrier is a familiar object or pathway that preserves the structure of a system without explaining it. Carriers are not teaching aids — they are transfer tests. When a pattern maps cleanly onto a carrier from a domain with no shared vocabulary, no shared institutional history, and no shared consequence profile, that transfer is evidence of structural depth below the level of local language. Each of the following sections is a transfer test, not a decorative example. If the ten-move sequence and the frame-

level gate absence cannot be cleanly mapped onto a domain's carrier, the claim needs revision. The carrier does not explain the structure. The carrier tests whether the structure is real.

Section VI Primary Proof I: Medicine

Medicine is the strongest single proof. The correction sequence is explicit, consequence-bearing, and visible in ordinary language. And medicine approaches Move 4 at the object level more closely than any other domain — which is precisely why its failure to complete the sequence at the frame level is so instructive.

The carrier: hospital triage and surgical pathway

Move	Medical Instantiation
Signal	Patient presents with recurring or acute symptom
Assess	Triage: urgency, pathway, severity
Diagnose	Tests, imaging, history; structural cause sought, not just symptom
Gate (object)	Operate or don't operate — the patient-level gate
Gate (frame)	Absent
Design	Surgical plan: structural intervention, not symptom management
Execute	Operation
Verify	Post-operative assessment: did the intervention address the source?
Deploy	Discharge
Monitor	Follow-up, rehabilitation, complication monitoring
Restart	Readmission: new diagnostic cycle opens

The operate-or-don't-operate decision is real. A surgeon who declines to operate on a patient who does not require surgery is making a structurally correct gate decision with irreversible consequence. This is Move 4 at the patient level. Medicine has it.

When the system under examination is not the individual patient but the institution treating patients, Move 4 at the frame level is absent.

The NHS has applied five distinct classes of correction across primary care continuity since the 1980s: the Summary Care Record, the Shared Care Record, the NHS App, ARRS workforce expansion, and the 2024/25 contractual continuity requirement. Paper 3's four biases explain exactly why each remained a correction within the inherited frame. Execution bias preserved the clinical route under pressure. Recognition bias gave no native lane to the diagnosis that the architecture itself should change. Defensibility bias expanded the audit layer. Incentive bias selected for keeping care moving. The governing diagnosis was available. The frame question was never formally entered into the correction sequence.

The result is what Paper 4 named: motion without exit. The system moved. The dependence did not. Formal correction did not reduce transferred burden in lived use.

Medicine proves the sequence. Health-system reform history proves the cost of the missing gate.

Section VII Primary Proof II: Aviation

Aviation has gone further than almost any other domain in separating functions, bounding roles, and formalising live control. That makes it an excellent proof: if the gap still appears here, it is unlikely to be a local weakness of sloppy systems.

NTSB and AAIB reports produce some of the most rigorous structural diagnoses available in any field. They locate the first break with precision, trace contributing factors carefully, and separate surface symptoms from structural causes. Moves 1–3 are executed at an extremely high standard.

The recommendations route to correction within the existing regulatory and operational frame.

Why? Because the NTSB investigates incidents. It has no formal authority to ask whether the certification architecture, the operational structure, or the human-automation boundary design should be replaced. Those questions are answered through separate regulatory and legislative processes entirely outside the incident correction loop. Investigation and governance are deliberately separated — which is exactly right for avoiding the function collapse Paper 2 describes. But it also means the incident correction loop has no formal connection to the frame-level gate.

This is aviation's near-miss: it separates investigation from operation more explicitly than almost any other domain, a genuine structural advance. But the same separation that prevents function collapse severs the path from the correction loop to the frame-level question. The incident loop diagnoses with precision. Legislators and regulators govern. The gate sits in neither.

The same structural failure modes — controlled flight into terrain, crew resource management failures, handoff failures at the human-automation boundary — reappear across decades of incident data. Each is diagnosed correctly. Each is corrected within the frame that produced it. Paper 3's recognition bias applies: the incident loop has no interface for 'should this operational architecture be replaced?' That question routes to legislators and regulators. The loop does not contain the gate.

A domain can separate functions correctly and still lack the frame-level gate. Separation defers collapse. It does not complete the sequence.

Section VIII Primary Proof III: Legal Prosecution

Legal prosecution offers a proof of a different kind — slower, more procedural, and universally legible. Its consequence profile is singular: a wrong verdict cannot always be undone. That makes it a strong boundary confirmation: if the gap appears where formal procedure is most elaborate, it appears everywhere.

The charging decision is Move 4 at the case level. It asks whether the evidence is sufficient and the frame of the prosecution is correct before proceeding. A prosecutor who declines to charge when the case does not support a conviction is making a structurally correct Move 4 decision with irreversible consequence.

The legal system has no equivalent gate at the frame level. No formal mechanism inside the prosecutorial or judicial correction sequence asks whether the adversarial frame should be replaced — whether the burden of proof allocation is structurally correct, whether the categories of offence map onto the harms they are meant to govern. Those questions are handled outside the operational correction loop, in scholarship, commission work, and political process.

When a verdict is contested, the correction mechanism available is appeal, which is adversarial by design. The same frame that produced the original judgment is the frame through which the judgment must be challenged. The correction loop does not step outside the adversarial structure. It re-runs it under additional scrutiny. An adversarial correction mechanism can question whether the procedure was followed correctly. It cannot question whether adversarial procedure is the right structure for finding the truth. The frame is preserved precisely because the correction mechanism is built from the same materials as the original judgment.

Miscarriages of justice are almost always Move 4 failures at the case level: the frame was wrong, and every subsequent move executed correctly inside a broken primitive. The trial was procedurally fair. The verdict was sound. The sentence was applied. The frame was wrong. The correction was technically perfect and structurally unjust.

Where the frame is adversarial by inheritance, procedural intelligence alone does not amount to redesign. More intelligence inside a bad frame defends it.

Section IX Primary Proof IV: Software and SRE

Software incident response is the most modern proof and the one that comes closest to completion. That proximity is what makes it useful: it shows what near-completion looks like, and why near-completion is still not completion.

The fix-or-rollback gate is real. In mature SRE organisations, postmortems generate architectural review items, acknowledging that the current system architecture may need redesign. This is the nearest software approximation to a frame-level gate.

But the architectural review sits outside the incident loop. It is a downstream consequence of a sufficiently strong postmortem recommendation, activated in a separate governance process. The question ‘should this architecture be replaced?’ is not a formal gate inside the incident correction sequence. It routes elsewhere, if it routes at all.

Paper 3’s four biases are especially clear here: execution bias keeps the incident loop oriented toward service restoration; defensibility bias produces postmortems that explain what happened but do not alter architecture; recognition bias routes the frame-replacement question to a roadmap process that is structurally separate; incentive bias rewards uptime restoration before structural redesign. All four operate exactly as Paper 3 predicted. The result is the same.

Technical debt is software’s own name for the accumulated result: repeated corrections inside a frame that should have changed earlier, compounding until replacement can no longer be avoided. The sequence ran without completing Move 4. The price was paid at replacement, under duress.

Software proves not only that the sequence transfers. It shows what near-completion looks like, and why near-completion is still not completion.

Section X Pattern Extraction

Four domains. One mechanism.

Four domains. Different tempos, consequence profiles, and institutional histories. The same ten-move sequence. The same stall point at Move 4 at the frame level. The same downstream consequence: corrections that hold temporarily and fail in the same structural way the original failure failed.

> **The Tempo Theorem**
>
> The tempo of the correction sequence varies by domain. Medicine is acute and compressed. Legal prosecution runs in months and years. Software incident response operates in minutes. The consequence profile differs: immediate and irreversible in medicine, delayed and redistributed in software, catastrophic and singular in aviation.
>
> The structural gap remains constant. Mechanism is what transfers. Tempo is what varies. The Governed Correction Sequence is not a claim about pace. It is a claim about structure.

The extractable consequence, stated directly: formal correction counts only if it reduces transferred burden in lived use. Where the frame-level gate is absent, correction cycles through the front of the sequence, densifies the frame, and produces review-rich, exit-poor systems. Paper 4's Migration Law belongs here naturally: institutions persist while the burden of their incoherence can be transferred outward; they begin to be replaced when a superior coordination architecture makes that burden avoidable. That is the institutional-scale consequence of Move 4 never being reached. Burden transfer is what happens when the correction loop cycles indefinitely. Avoidability is what a formally completed sequence enables.

Pre-Modern Transfer Test: Manuscript Culture and Print

The cross-domain claim strengthens when tested against a case with no shared vocabulary with medicine, aviation, or software — and no modern quality framework of any kind.

Manuscript culture did not lose centrality because scribal systems were criticised. It lost centrality when print made the burden of text reproduction and distribution avoidable through a lighter route. The scribal bottleneck — few producers, slow copying, high gatekeeping, high replication cost — did not lose relevance because someone named its inefficiency. It lost relevance when a lower-burden route existed and people used it. The frame changed from outside, through replacement. That is precisely the pattern this paper predicts.

No modern quality framework was present. No shared technical vocabulary existed. The pattern holds. This is what structural depth looks like when tested against a reality surface with no shared language.

The Correction Decay Spectrum

The following spectrum maps five observable stages of correction failure. Each stage is reached when the prior correction layer proves insufficient and burden migrates outward.

Stage	What It Looks Like	Burden Status	Domain Example
Functional Correction	Failures detected and resolved within a coherent frame. The sequence runs cleanly and the hold promise is kept.	Paid by the system	High-reliability organisations with separated functions (post-1970s aviation safety)
Correction Densification	Reviews, audits, dashboards, and policies multiply. The frame is defended and refined, not questioned.	Paid by documentation	NHS primary care 1990–2010; financial regulation between crisis cycles
Correction Collapse	Correction moves fuse into overloaded human nodes. The same person must observe, scope, diagnose, decide, and keep the system running simultaneously.	Paid by individuals	NHS post-2010 GP overload; SRE hero developer
Ghost Correction	Informal carriers perform what the formal sequence cannot distribute. The ghost structure becomes the only coherence layer.	Paid by hidden labour	Family-held patient records; informal clinical triage; developer as living architecture map
Frame Exhaustion	The system can only repeat correction or be replaced. The legal surface has opened. Migration pressure is externally visible.	Paid and accumulating visibly	Financial regulation post-2008; NHS post-2024; every domain at the moment of structural replacement

Compounding institutions

A small number of systems have approached completion of the sequence more fully than others. Toyota’s Andon mechanism is the closest physical instantiation of a frame-level gate in any production environment: any worker on the line can pull the cord and stop the entire frame from running, triggering an immediate decision about whether the current process should continue or change. It is Move 4 made real, with a physical interface and consequence authority held at the point of production. The contrast with most correction frameworks is stark: PDCA, DMAIC, and OODA all have the equivalent of a suggestion box. The Andon cord stops the line. The suggestion box does not.

Aviation safety separated investigation from operation more explicitly than its predecessors. Mature SRE formalised postmortems and architecture review consequences more rigorously than earlier software operations. None complete the sequence at the full frame level. All show structural improvement that compounds more than it cycles — visible in quality curves, accident rate data, and platform reliability histories. That correlation is itself evidence for the law.

Secondary Confirmation: Financial Regulation

Financial regulation provides four correction cycles across four decades, each confirming the pattern. The 1987 Black Monday response intensified circuit-breaker mechanisms and market oversight within the existing trading architecture. The 1998 Long-Term Capital Management collapse triggered new risk reporting requirements within the same leverage-permitting frame. The 2008 crisis produced the largest regulatory response in a generation — Dodd-Frank, Basel III, stress-testing regimes — all of which densified the frame rather than replacing it. The 2023 Silicon Valley Bank collapse revealed that the stress-testing frame had not been applied to the class of risk it produced. Each cycle: correction within the frame. No frame-level gate reached. The failure returned as an improved version of itself.

Section XI Correction Collapse and Meta-Substitution

Paper 2 showed that when sensing, interpretation, authority, and memory collapse into the same human node under consequence, the institution fails structurally while appearing to function — because a ghost structure of informal carriers absorbs what the architecture cannot hold. The visible story is always about individuals under pressure. What disappears from view is the topology that placed multiple forms of institutional work in the same location.

The same mechanism appears one level higher.

Correction collapse is the condition in which the ten correction moves are not distributed across separated roles, tools, and time. The same people or teams must observe, scope, diagnose, decide, design, execute, verify, deploy, monitor, and re-enter while still keeping the existing system running. Paper 3 named the result: the system becomes data-rich and structure-poor. More is captured. More is recorded. More is traceable. What the system cannot do is locate the structural source upstream of the loops doing the accounting. It responds harder and harder while remaining structurally stuck.

Meta-substitution is what happens when correction collapse persists. The hero developer who is the only route through a class of recurring failures. The experienced nurse who triages informally because the formal process has been bypassed. The administrator who is the living map of a system that official process cannot navigate. These are not exceptional individuals. They are the ghost structure of the correction layer — carrying the correction work the formal sequence has not been built to distribute.

When they leave, the organisation discovers it had no structural correction capacity. It had a person carrying it outside the formal sequence.

The cascade follows Paper 2’s pattern: responsibility is personalised; documentation inflates; process sprawls; informal substitution absorbs the gap; the system becomes more legible to observers while navigability worsens for participants.

Paper 2 describes collapse at the operational layer. This paper describes collapse at the correction layer. Together they form a complete diagnostic: the system fails to operate correctly, then fails to correct correctly, then transfers the burden outward while persisting on the strength of what it has not acknowledged.

Section XII Substitution by Gravity

When Move 4 at the frame level is formally reached and a structurally sound replacement is built, the replacement does not succeed by argument. It succeeds by gravity.

Users do not move through reasoned preference. They move because the burden of remaining in the broken system — the burden currently absorbed by hidden carriers — exceeds the burden of entering the new one.

> **The Law of Substitution by Gravity**
>
> Users do not leave broken systems because they are convinced by argument. They leave when the structural burden of remaining in the broken frame exceeds the friction of entering the new one. The burden differential is the gravitational field. Structural honesty at the primitive level is the gravitational constant.

Three clarifications prevent misreading.

First, gravity does not activate merely because a substitute is better designed. A replacement may be structurally more coherent and still fail to matter if it does not reduce the burden users are actually carrying. The burden differential must be real, not theoretical.

Second, gravity may remain latent while the ghost structure is still absorbing the burden successfully. This is why weak systems can survive so long. The burden exists, but it is being carried elsewhere. When those carriers withdraw, the differential becomes visible all at once. Migration can activate quickly when it was building quietly for years.

Third, gravity alone is not enough. A substitute may possess structural truth without possessing a recognition channel. A patient-held record may be more accurate than the official system and still fail to change care if clinicians cannot act on it without re-entering every field into the hospital's own system. The truth exists. The gravity does not activate.

The Gate-Admissibility Boundary

Move 4 at the frame level produces the decision to pursue replacement. It does not guarantee that replacement becomes institutionally real. That is Paper 6's domain.

A structurally complete intervention requires two separate thresholds. The gate threshold — Move 4 — is the formal mechanism deciding whether to replace the frame. The admissibility threshold — Paper 6 — specifies the conditions under which a substitute can safely carry consequence.

Without the gate, replacement never initiates. Without admissibility, initiated replacement fails to become real. Both are necessary. Neither is sufficient alone.

Move 4 opens the possibility. Gravity governs migration. Admissibility determines whether migration can carry consequence. All three are necessary. None collapses into the others.

Section The Frame-Gate Test

XIII

Answer each question for the correction framework or system under examination. An answer that reveals absent frame authority, fused correction moves, and ghost correction structures confirms the failure class.

Q	Question	What the Answer Reveals
1	What is the recurring failure?	Not the most recent symptom. The structural source producing the same pattern across correction cycles.
2	Which correction sequence currently receives it?	Name the framework in operation: PDCA, DMAIC, incident loop, audit cycle, review board.
3	Which of the ten moves are present, absent, fused, or informal?	Fused moves are where correction collapse accumulates. Absent moves are where burden migrates.
4	Where is the object-level gate?	The fix/rollback, operate/don't operate, charge/decline decision. Who holds it and with what consequence authority?
5	Does a formal frame-level gate exist?	A mechanism that can route the failure to frame replacement rather than frame refinement — separate from the operational correction loop and with authority to act.
6	Who has authority to reject the current frame?	If no one formally holds this authority, the gate is absent by design, not by accident.
7	What happens when redesign signals enter the system?	If they become complaints, delays, or non-executable thinking, recognition bias is operating at the gate.
8	Has formal correction reduced transferred burden in lived use?	Or has it made the institution more review-rich while leaving the structural failure intact?
9	What ghost correction structure has formed?	Who is informally carrying the correction work the formal sequence cannot distribute? What happens when they leave?
10	If the frame were replaced, which admissibility conditions must the substitute satisfy?	Paper 6 specifies these. No exit becomes institutionally real without admissibility. Name the conditions before designing the exit.

Decision Flowchart

Step 1 — Does the system face recurring structural failure that formal correction has not resolved? No → framework does not apply. Yes → Step 2.

Step 2 — Can you identify the ten moves in the current correction sequence? Which moves are present, absent, fused, or informally held? → Proceed to Step 3.

Step 3 — Does the sequence contain a formal frame-level gate — a mechanism capable of routing the failure to frame replacement, with authority to act? No → failure class confirmed. Apply the Frame-Gate Test. Yes → Step 4.

Step 4 — Has that gate been formally triggered? If yes: has a replacement been specified and subjected to Paper 6's admissibility conditions? If no: the gate exists nominally but without consequence authority. → Step 5.

Step 5 — Where is the burden accumulating? Locate the case on the Correction Decay Spectrum (Section X). Stage determines urgency. Proceed to Section XII for the replacement mechanic and the Gate-Admissibility Boundary.

Practitioner Handoff

If the Frame-Gate Test confirms the gate is absent, three structural moves are available — in order of prerequisite, not preference.

Pattern Identified	Structural Intervention	Trap to Avoid
Ten correction moves fused into the same person or team; correction collapse operating	Separate the ten moves: ensure no single person or team holds more than three consecutive moves. Distribute Signal, Diagnose, Gate, and Monitor across distinct roles with distinct authority.	Correction collapse: fused moves recollapse under pressure unless the separation is structurally enforced, not just assigned on an org chart.
Frame-level gate absent; same structural failure class recurring across correction cycles	Install a formal frame-level gate with authority to recommend frame replacement — not just flag concerns. Trigger: same structural failure class recurs across three correction cycles. Authority: external to the operational correction loop.	Nominal gate: the mechanism exists on paper but routes to no one with authority to act. Confirm the gate has consequence authority, not only process.
Replacement required but not yet built; burden differential approaching Migration Threshold	Build a replacement in parallel before the gate triggers. Substitution by Gravity means the replacement must be ready when burden differential exceeds entry friction. Paper 6's admissibility conditions apply before deployment.	The Successor Trap (Paper 6): a replacement that clears admissibility initially but drifts back toward the failure topology as it scales. Design refusal integrity into the translation layer from the start.

What this paper adds to the series

Instrument	Definition
The Governed Correction Sequence	A named ten-move structural sequence that multiple high-consequence correction frameworks independently converge upon
The Governed Correction Law	The sequence is sufficient for correction within a frame, insufficient for determining whether the frame should persist. Without a formal frame-level gate, structural failures return as improved versions of themselves
The Frame-Level Gate distinction	Object-level gates act on a case or instance. Frame-level gates ask whether the system performing correction should itself be replaced
The Tempo Theorem	Mechanism is what transfers across high-consequence domains; tempo is what varies

Instrument	Definition
The Correction Decay Spectrum	Five observable stages of mismatch progression: Functional Correction → Correction Densification → Correction Collapse → Ghost Correction → Frame Exhaustion
The Frame-Gate Test	A ten-question diagnostic instrument for locating the missing gate and identifying where its absence transfers burden
Correction Collapse	Collapse at the correction layer: the ten correction moves fuse into the same node. Distinct from operational collapse (Paper 2)
Meta-Substitution	The ghost structure at the correction layer: informal carriers performing correction work the formal sequence cannot distribute
The Law of Substitution by Gravity	Users migrate when the burden differential between the broken frame and the replacement exceeds entry friction. Argument is not the mechanism
A falsifiable prediction	Any framework lacking Move 4 at frame level will produce corrections that hold temporarily and fail in the same structural way. Five NHS cycles and four financial regulation cycles confirm it

Section XIV — Conclusion — The Gate Is Where the Sequence Either Completes or Returns

Boundary conditions — when the claim does not apply

The Governed Correction Law applies specifically to high-consequence domains where failure causes significant harm or institutional failure, and where correction is formally separated from governance. It may not apply to very low-consequence domains where failure is cheap and replacement trivial; systems with explicit constitutional amendment processes formally integrated into the correction loop; or small, informal groups where frame replacement is implicit and uncontested. No counterexample has been found among the domains examined. The absence of a counterexample does not establish universality. It justifies treating the pattern as an archetype and the cross-domain transfer as evidence of structural depth.

The falsifiable prediction

Falsifiable Prediction

Any correction framework that does not contain Move 4 at the frame level will produce corrections that hold temporarily and then fail in the same structural way the original failure failed. The NHS provides five cycles of confirmatory pattern evidence. Financial regulation provides four cycles across 1987, 1998, 2008, and 2023. The stall point is the same in each cycle. The frame was not questioned. The failure returned.

The claim would be weakened if a correction framework were found with a formal frame-level gate inside the correction sequence that reliably routes recurring structural failure to replacement rather than refinement. Across the domains examined, no such mechanism

has been found. Every examined instantiation lacks the gate at the frame level. The pattern repeats.

What changes when the sequence is seen

What changes is not immediate salvation. It is the availability of better questions.

A policymaker who can locate the absence of Move 4 before commissioning reform will ask whether the proposed intervention addresses the frame or merely improves within it. An institutional designer who can name correction collapse before writing governance will separate the correction phases rather than loading them onto the most senior available human node. A founder who understands Substitution by Gravity — and the recognition gap preconditions — before building a replacement will know that structural lightness is necessary but not sufficient; recognition and admissibility still determine whether consequence can move.

This paper's own admissibility acknowledgment

Paper 6 closed by acknowledging that the series itself possesses a recognition gap: it holds a diagnostic truth kernel but does not yet possess a recognition channel through which it can be officially relied upon under consequence. The same acknowledgment applies here. This paper names the sequence, locates the gate, and explains why the gap is structural rather than incidental. That diagnostic contribution stands. Whether it becomes institutionally real depends on the same admissibility conditions it describes. The sequence has been independently evolved for centuries. It has not been formally completed at the frame level. The rest is construction.

Correction names the failure.
Governance names the gate.
Replacement names the exit.

The domain names the field. The sequence names the structure. The gate is where the sequence either completes — or returns.

Jamie Forrester · hello@jamieforrester.com · April 2026

If this maps to a correction sequence you are running, a reform programme you are designing, or an institution you are building against, you can reach me at hello@jamieforrester.com

Appendix A — Framework Reference

Framework	Domain of Origin	M4 Object	M4 Frame	Primary Strength	Structural Limit
PDCA	Manufacturing	✗	✗	Iterative	Restarts within the

Framework	Domain of Origin	M4 Object	M4 Frame	Primary Strength	Structural Limit
	/ quality			improvement cycle	same frame; no gate
DMAIC	Semiconductor / six sigma	✗	✗	Rigorous statistical diagnosis	Define phase takes frame as given; no replacement route
OODA	Military combat operations	Partial	✗	Speed of adaptation to changing conditions	Orient adjusts cognitively; no formal replacement gate
Toyota A3	Post-war Japanese manufacturing	✗	✗	Root cause + countermeasure discipline	Andon partially instantiates a gate; A3 does not
Ford 8D	Automotive defect resolution	✗	✗	Team-based, discipline-structured problem closure	Closes the object-level problem; does not question frame persistence
SRE Incident Response	Internet infrastructure at scale	Partial	✗	Blameless postmortem; architecture review triggered	Review sits outside the loop; no formal frame-level gate inside the sequence
SBAR	Acute clinical communication	✗	✗	Rapid situational clarity at point of consequence	Communication framework; not a correction sequence

Appendix B — Glossary of Named Concepts

Term	Definition and Source
Governed Correction Sequence	The ten-move structural sequence multiple high-consequence correction frameworks independently converge upon. Sufficient for correction within a frame, insufficient for determining whether the frame should persist. [Paper 7]
The Governed Correction Law	Where systems must correct recurring structural failure under ambiguity and consequence, independent frameworks converge on the same ten-move sequence. Without a formal frame-level gate, structural failures return as improved versions of themselves. [Paper 7]
Object-Level Gate	Move 4 at the case or instance level: fix or rollback, operate or don’t operate, charge or decline. Present in every framework examined. [Paper 7]
Frame-Level Gate	Move 4 at the architectural level: a formal mechanism asking whether the system performing correction should itself be replaced. Absent from every framework examined. [Paper 7]
The Tempo Theorem	Mechanism is what transfers across high-consequence domains; tempo is what varies. The structural gap is constant beneath variation in pace and consequence profile. [Paper 7]

Term	Definition and Source
Correction Decay Spectrum	Five observable stages: Functional Correction → Correction Densification → Correction Collapse → Ghost Correction → Frame Exhaustion. [Paper 7]
Correction Collapse	The ten correction moves fuse into the same node rather than distributing across separated roles and time. Distinct from operational collapse (Paper 2). [Paper 7]
Meta-Substitution	The ghost structure at the correction layer: informal carriers performing correction work the formal sequence cannot distribute. [Paper 7]
The Law of Substitution by Gravity	Users leave broken systems not through argument but when burden differential exceeds entry friction. Gravity governs migration; admissibility determines whether migration can carry consequence. [Paper 7]
Gate-Admissibility Boundary	The distinction between Move 4 (which produces the decision to pursue replacement) and admissibility (which determines whether that replacement becomes institutionally real). Without the gate, replacement never initiates. Without admissibility, initiated replacement fails to become real. Both are necessary; neither is sufficient alone. [Paper 7, Paper 6]
Frame-Gate Test	A ten-question diagnostic instrument for locating the missing frame-level gate and identifying where its absence transfers burden. [Paper 7]
Carrier	A familiar object or pathway that preserves the structure of a system without explaining it. Used as a transfer test: clean mapping across unrelated carriers is evidence of structural depth below domain vocabulary. [Paper 7]
The Redesign Law	Execution improves performance within a frame. Diagnosis improves control within a frame. Redesign changes the frame itself. The first two do not reliably produce the third. [Paper 3]
Migration Law	Institutions persist while the burden of their incoherence can be transferred outward. They begin to be replaced when a superior coordination architecture makes that burden avoidable. [Paper 4]
Admissibility	The conditions under which a substitute can safely carry consequence. Completing the sequence opens Move 4. Admissibility determines whether the gate's output becomes institutionally real. [Paper 6]

Appendix C — Series Reference

Forrester, J. (2026a). The Expertise Illusion in AI Task Marketplaces. SSRN Working Paper.

Forrester, J. (2026b). The Four-Function Law of Scalable Institutions. SSRN Working Paper.

Forrester, J. (2026c). Why Systems Can't Fix Themselves: The Missing Redesign Layer. SSRN Working Paper.

Forrester, J. (2026d). Institution Migration: How Better Coordination Makes Legacy Institutions Unnecessary. SSRN Working Paper.

Forrester, J. (2026e). The Institutional Replacement Pipeline. SSRN Working Paper.

Forrester, J. (2026f). The Admissibility Problem: Why Better Substitutes Still Fail to Replace Worse Institutions. SSRN Working Paper.

Forrester, J. (2026g). The Governed Correction Sequence. SSRN Working Paper. [This paper]

Forrester, J. (2026h). Host-Constrained Primitive Mismatch: The Structural Decay of Relationship-Coded Systems Inside Transaction-Coded Hosts. SSRN Working Paper.

---

hello@jamieforrester.com · jamieforrester.com · Edinburgh, UK · Paper 7 v3.0 · April 2026