DIAGNOSTIC AND REPLACEMENT SERIES · PAPER 6

THE ADMISSIBILITY PROBLEM

Why Better Substitutes Still Fail to Replace Worse Institutions

Jamie Forrester

Independent Systems Architect · Edinburgh, UK

hello@jamieforrester.com · jamieforrester.com

April 2026 · v3.0 · Working Paper

Systems are not replaced when they are outdesigned. They are replaced when better structure becomes admissible at the level of consequence.

Abstract

A superior substitute fails to replace a worse institution not because it is poorly designed, but because it has not crossed the admissibility threshold: the point at which consequence-bearing actors can safely rely on it without the incumbent's permission. This paper names the Admissibility Problem, defines the Admissibility Law, and identifies seven interdependent criteria — the Admissibility Stack — that must be satisfied in dependency order before a substitute can inherit institutional reality. The paper introduces the Recognition Gap (truth without a recognition channel), its counterpart the Recognition Channel, the Refusal Invariant (what the system cannot bypass without ceasing to be what it claims), the Successor Trap (decay back into the old topology at scale), and the Inter-Kernel Problem (the person as the seam between substitutes).

Three proofs are examined: Open Banking as the deliberate engineering of admissibility through regulatory force; the NHS Summary Care Record as a stall at the recognition criterion; and the Extended Project Qualification as a stall at the legal/operational interface. The paper closes with a ten-question Admissibility Test, a practitioner handoff for each criterion failure, a falsifiable prediction, and the paper's own admissibility acknowledgment. The governing tagline: truth without a recognition channel is still peripheral.

For System Builders and Institutional Designers — Read This First

You are facing the Admissibility Problem if your substitute is lighter, more coherent, and less dependent on hidden burden than the institution it could replace — yet consequence-bearing actors still route through the incumbent. The problem is not that your substitute is wrong. The problem is that correctness alone does not make it admissible.

A substitute becomes admissible when consequence-bearing actors can safely rely on it without the incumbent's permission. That requires satisfying seven criteria in dependency order: Trust Kernel, Refusal Integrity, Contestability, Governance Legibility, Fallback Integrity, Legal/Operational Interface, and Recapture Resistance. Each depends on the one before it. The chain runs from the inside out.

Three questions reveal where the stall is:

- Which of the seven criteria is currently absent or partial?
- Is the recognition channel open — can adjacent systems accept your outputs without rebuilding truth from scratch?
- What structural mechanism prevents recapture as you scale?

If the recognition channel is closed and you do not control it, admissibility must be forced open at each criterion, not merely mandated. The Admissibility Test in Section XIV provides the instrument. The practitioner handoff specifies what to do when each criterion fails.

Series Position

This is Paper 6 in the Diagnostic and Replacement Series. Papers 1–3 extracted the diagnostic laws: interface-legitimacy mismatch, function collapse under scale, and the structural reason why correction loops cannot generate redesign from within the same frame. Papers 4–5 named the replacement mechanics: burden transfer as the migration driver, and the five-stage pipeline from recurring failure to buildable substitute form. Paper 6 asks the next question: once a substitute is correctly built, what allows it to become institutionally real rather than remain correct at the edge? Paper 7 names the Governed Correction Sequence — the meta-layer that shows what Papers 1–6 are all instances of. Paper 8 applies the full toolkit to a new failure class: externally imposed primitive compression.

Reader Map

Reading Path	Sections
Full read — institutional designers, founders, policymakers	All sections and appendices
Executive read — CTOs, platform leads, reform programme directors	For System Builders box, then Sections I, IV, VI, VII, XIV, XVI
Series reader	Sections I and XVI only

Section I The Admissibility Name Problem

A mother arrives at a new school in September carrying a diary. It contains a more truthful record of her child than the institution itself holds: what triggered last year's anxiety, what intervention worked, which seating arrangement prevented escalation, what the previous teacher learned too late, and what must not be forgotten. The new teacher thanks her, sets the

diary aside, and opens the school's official system, which contains almost none of it. The diary is more continuous, more truthful, and more useful. It counts for nothing officially.

That is the admissibility problem.

Before asking why better substitutes stall, the thresholds must be separated. They are not a single scale from worse to better. They are distinct conditions, each of which a substitute can satisfy without satisfying the next.

Threshold	What It Means
Usability	A person can use the system for a task
Adoption	Some users choose the system over the incumbent
Migration	Reliance begins to shift over time toward the substitute
Legitimacy	The system appears rightful or trustworthy to observers
Institutional inheritance	The substitute begins carrying binding consequence previously held by the incumbent
Admissibility	The condition under which a substitute becomes sufficiently trustworthy, inspectable, governable, contestable, and recognisable to carry consequence without the incumbent's re-validation

A system can be adopted without being admissible. It can appear legitimate without inheriting institutional reality. These are different thresholds. The admissibility problem is the gap between the last threshold and the one before it.

> **Admissibility Failure**
>
> A substitute enters admissibility failure when it can reduce burden, preserve truth, or coordinate better than the incumbent, but cannot be safely relied upon by consequence-bearing actors without the incumbent's recognition, re-entry, or validation.
>
> Admissibility failure is not the same as product failure. A technically correct and user-preferred substitute is in admissibility failure when the world still routes consequence through the incumbent. The diary is true. The teacher opens the school system anyway.

This is a structural paper, not a reform programme. It does not prescribe transition policy, venture strategy, or lobbying tactics. It names a recurring constraint that becomes visible only after replacement has become conceptually possible and technically buildable: superior substitutes do not displace incumbents merely because they are better. They displace incumbents when they become admissible.

Section II The Promise That Better Systems Win

The common promise is simple: if a substitute is lighter to inhabit, preserves continuity more truthfully, and reduces burden more directly, displacement will follow.

This promise is incomplete.

Many systems are visibly better for users while remaining structurally peripheral. They reduce friction locally, preserve continuity more coherently, and yet fail to become the recognised place where consequence lives. The problem is not that the substitute is wrong. The problem is that correctness alone does not make it admissible.

Papers 1–5 explain why institutions are structurally incoherent and how better ones can be built. Paper 6 asks what allows a correctly built substitute to become real. The missing layer is not diagnosis, redesign, migration logic, or manufacture. It is admissibility: the threshold at which a substitute becomes trustworthy, inspectable, governable, contestable, recognisable, and resistant to recapture enough to carry consequence in the real world.

Section III The Observable Reality: Superior Substitutes Stall

The recurring reality is familiar. Better systems are built. Users at the edge may prefer them. Burden may fall locally. Coherence may improve. Yet the incumbent remains central. Why?

Because the incumbent still holds official recognition, legal standing, procedural finality, interoperability with adjacent institutions, symbolic authority, and default trust under risk. The substitute works. The world still refuses to treat it as the place where reality can safely count.

Non-admissible substitutes — the recurring pattern

Substitute	Better Than Incumbent At	Why Still Peripheral
Mother’s diary	Continuity truthfulness; developmental context	Not institutionally recognised; counts for nothing officially
Patient-held health record	Completeness, portability, patient control	Clinicians cannot rely without re-entering into official system; no binding obligation to consult
Extended Project Qualification (EPQ)	Genuine capability signal over A-level grades	Admissions grammar discounts it; does not fit UCAS tariff processing
Google Health (2008)	Patient control over records	No refusal integrity; stored continuity without ability to reject incoherence
Early Uber / Airbnb	Lighter coordination; lower friction	Recapture failure at scale; old topology returned inside the successor

In each case: the substitute is technically superior. The incumbent retains the recognition channel through which outputs become official. The substitute remains correct at the edge.

Section IV The Admissibility Law

The Admissibility Law

A substitute inherits institutional reality not when it surpasses the incumbent in design, but when consequence-bearing actors can safely rely on it without the incumbent's permission.

Migration reduces burden. Manufacture builds correctness. Admissibility makes both historically real.

A substitute becomes admissible when its truth can be relied on under consequence, inspected without blind faith, governed without opacity, contested without prohibitive cost, recognised by adjacent systems, and protected against recapture as it scales. If those conditions do not hold, substitution stalls regardless of design quality.

Hard line: systems are not replaced when they are outdesigned. They are replaced when better structure becomes admissible at the level of consequence. The admissibility threshold is the missing layer between correct substitutes and institutionally real ones.

Section V The Recognition Gap and the Recognition Channel

Legacy institutions do not merely coordinate. They hold the recognition channel through which outputs become official. A diagnosis counts because it is institutionally recognised. A credential counts because it is institutionally recognised. A judgment counts because it is institutionally recognised. A record counts because adjacent systems recognise it as valid.

A substitute may possess truth without possessing a recognition channel. That is why technically superior systems can remain socially unreal.

The Recognition Gap

The recognition gap exists when a substitute holds a more truthful record than the incumbent, but consequence-bearing actors cannot act on that record without first routing through the incumbent's own interface.

A patient-held health record may be factually complete, but if a specialist refuses to act on it without re-entering every field into the hospital's system, the recognition gap is present. The truth exists. The channel does not.

The Recognition Channel

The route through which a substitute's outputs become acceptable to adjacent consequence-bearing systems. Without it, truth remains private, local, or advisory — correct at the edge but structurally peripheral.

The admissibility problem begins precisely where truth and recognition diverge. A substitute cannot inherit institutional reality until it possesses a recognition channel that does not require the incumbent's permission.

The Legitimacy Inertia Problem

Institutions persist not only because they coordinate, but because they remain the recognised location of consequence. This inertia has several structural components: symbolic authority (institutions inherit trust from history, ritual, title, and embedded status); risk asymmetry (actors prefer recognised failure to unrecognised correctness); legal incumbency (the legacy route retains formal standing); credential monopoly (incumbents often control who counts as qualified to act); and mandatory recognition structures (substitutes may be unusable unless law or incumbent systems recognise their outputs as valid).

Medical licensing is one such structure. A continuity tool may preserve a patient's history better than the health system itself, yet remain inadmissible if only institutionally recognised clinicians, records, or interfaces can trigger binding action. The substitute is not merely competing with the incumbent. It is competing with the incumbent's control over recognition.

A substitute becomes admissible when it is more defensible to rely on than the failing legacy system it could replace.

Section VI The Admissibility Stack

Seven criteria determine whether a substitute is admissible. They are not parallel features. Each depends on the one before it. The dependency runs in one direction: from the inside out. A substitute that satisfies contestability without first satisfying refusal integrity offers a legible route to dispute outputs that cannot be trusted. A substitute with governance legibility but no trust kernel makes the rules of a system visible without guaranteeing what that system is actually doing. The chain must be satisfied in order.

Criterion	Required Question	Depends On	Failure Mode if Absent
1. Trust Kernel	What exactly is being guaranteed, by whom, and how can that guarantee be verified?	(Base — no prior dependency)	No basis for reliance; all downstream criteria are void. Cannot satisfy any other condition without a defined kernel.
2. Refusal Integrity	Can the substitute structurally refuse invalid states, privileged exceptions, or silent modification?	Trust Kernel	The trust kernel has no enforcement layer. Stores continuity without being able to reject incoherence. First point where most substitutes fail.
3. Contestability	Can participants dispute outputs through a legible and accessible route?	Refusal Integrity	Outputs can be rejected but cannot be contested. Unchallengeable authority that erodes trust without recourse.

Criterion	Required Question	Depends On	Failure Mode if Absent
4. Governance Legibility	Can affected actors understand who changes rules, how, and under what boundaries?	Contestability	Rules exist but cannot be understood or anticipated. Opaque power that makes reliance structurally unsafe.
5. Fallback Integrity	If the substitute fails, can participants revert without catastrophic loss, hidden lock-in, or continuity destruction?	Governance Legibility	Failure creates lock-in. Risk-averse actors cannot safely rely on a system with no safe exit.
6. Legal/Operational Interface	Can adjacent consequence-bearing systems recognise the substitute's outputs without fully absorbing it?	Fallback Integrity	Correct but unusable. Admissible internally but invisible to the institutions whose recognition matters.
7. Recapture Resistance	What prevents the substitute from re-fusing functions, re-exporting burden, or locking in participants as it scales?	All above	The Successor Trap: the old topology returns inside the new system. Most substitutes eventually fail here, even those that cleared all prior criteria.

The Admissibility Stack operates as a one-way dependency chain: each criterion unlocks the next. Clearing a criterion does not relax the one below it. A substitute cannot skip the stack. It can only build through it. The dependency runs inward. The failure runs outward.

Admissibility is not a launch threshold. It is a preservation requirement. Every criterion must hold not just at deployment, but as the substitute scales.

Section VII The Refusal Invariant

The Refusal Invariant

A refusal invariant is a rule the system cannot bypass without ceasing to be what it claims to be. A substitute does not become trustworthy because its operators seem reasonable. It becomes trustworthy when it can structurally refuse what should not be allowed — and when that refusal can be inspected without requiring blind trust.

Legacy trust often depends on discretion: the hope that a person will behave correctly. A bank manager can override a fee for a friend. A clinician can omit an adverse reaction from a record. A teacher can adjust a grade without notation. Substitute trust must be different. It must be grounded in inspectable refusal, not goodwill.

In Spectrum Registry, uniqueness functions this way: no two valid claims can occupy the same canonical coordinate. No operator can waive that rule for convenience or influence. The

trustworthiness of the system rests on refusal, not permission. In an admissible learning record, a refusal invariant could take a different form: no grade may be altered without an append-only justification and a timestamped notification to the learner. The system cannot silently rewrite educational history.

Google Health, launched in 2008, allowed patients to hold and share their own records but imposed no structural refusal on incorrect inputs: clinicians and patients alike could enter contradictory states, and the system had no mechanism to reject them. It stored continuity without being able to refuse incoherence. The trust kernel had no enforcement layer. It was admissible in appearance, inadmissible in structure.

Any admissible substitute must specify what it will not allow, who cannot override that refusal, and how outsiders can verify it. That verification layer is the audit surface: not the user interface, but the inspectable surface through which a third party can verify the substitute's invariants without requiring unrestricted access to private state.

Section VIII The Successor Trap

The Successor Trap

A substitute may clear the admissibility threshold initially and still decay into the topology it replaced. The successor trap is not a risk of failure. It is the default trajectory of success without structural protection.

As substitutes scale, function fusion can return, burden transfer can return, governance can re-centralise, audit memory can displace operational memory, users can become translators of bad categories again, and success itself can reward the erosion of the kernel.

Uber is useful here. It launched as a lighter coordination route between drivers and riders. Over time it reproduced recruitment grammar mismatch, algorithmic qualification gates that drivers could not inspect or contest, burden transfer to drivers who absorbed cost without authority, fused responsibility at the driver node (insurance, safety, service quality, customer relations), and compliance load pushed into the same human point. That is not a different problem. It is the original failure topology returning inside the successor.

Airbnb shows the same pattern from a different angle. Early Airbnb satisfied refusal integrity and contestability adequately: its rating system enforced mutual accountability and disputes had a legible route. As it scaled, dynamic pricing opacity, superhost qualification gates, and host-side burden transfer reproduced the asymmetries of the hotel industry it was claimed to replace. Recapture resistance is not a launch condition. It must be structurally enforced as the system grows, or the successor trap is already operating.

Admissibility without recapture resistance is not stable inheritance. It is early-stage substitution.

Section IX Power, Gatekeeping, and the Legal Interface

Admissibility is not judged in a neutral field. It is often controlled by the very institutions being displaced. This happens through licensing, accreditation, procurement rules, standards control, legal monopoly over finality, mandatory recognition laws, and professional boundary enforcement.

A better substitute may fail not because it is incoherent, but because admissibility itself is institutionally gated. Open Banking makes this visible. Better account-access architectures did not become real because incumbents voluntarily accepted them. They became real because regulation forced incumbent participation, standardised interfaces, and opened the recognition channel. This is not a policy detail. It is proof that where gatekeepers control the recognition channel, superior design alone is insufficient.

Where incumbents control the recognition channel, a substitute has three structural options: wait for regulation to force the channel open; build a translation layer that makes the substitute's outputs legible in the grammar adjacent systems already accept (as UCAS tariff mapping would do for the EPQ); or build within a domain where the incumbent's recognition monopoly is weaker. None of these is primarily a design problem. All of them are admissibility engineering problems.

If consequence-bearing recognition is monopolised, then admissibility must be designed against, through, or around that monopoly. Superior design alone is insufficient.

Section X The Inter-Kernel Problem

The Inter-Kernel Problem

One substitute is never the whole world. Health, education, identity, finance, and legal finality intersect in lived use. If substitutes are built separately and their coordination burdens fall back onto the person, the hidden carrier problem returns at the seam.

The person must never again become the seam.

The mother with the diary also keeps a separate medication log for the hospital. The school's continuity kernel does not know the medication log exists. The hospital's kernel does not know what the school has recorded about anxiety triggers. She is the only integration layer between them — translating, summarising, remembering what each institution has forgotten to tell the other. When she is not present, the seam fails. If a substitute ecosystem reproduces this structure at a higher level of abstraction — better kernels, same person carrying the gap — it has not solved the problem. It has rebuilt it in cleaner components.

This is not merely a technical interoperability issue. It is a governance problem at the boundary. The wrong answer is a new monolith. The right answer is a minimal coordination grammar.

Rule	Requirement
Must	One kernel may assert facts that another may verify or refuse
Must	Cross-kernel exchange must be signed, inspectable, and append-only

Rule	Requirement
Must	Seams must preserve custody rather than collapsing it upward
Must	Kernels must publish their assertion and refusal rules
Must not	No kernel may silently modify another kernel's core state
Must not	No kernel may require the user to carry coordination burden that should live at the seam

These are the minimum conditions for cross-kernel admissibility. A substitute ecosystem that violates them does not escape the hidden carrier problem. It relocates it.

A system that removes institutional burden but preserves human coordination burden has not solved the problem. It has relocated it.

Section XI Primary Proof I: Open Banking

Open Banking matters because it shows that admissibility is not produced by technical elegance alone. It was engineered through deliberate design of each admissibility condition — criterion by criterion, through regulatory force.

PSD2 Mechanism	Admissibility Criterion Satisfied	How
Mandatory incumbent participation	Legal/Operational Interface	Banks could not refuse to expose recognised APIs; the recognition channel was forced open
Standardised API interfaces	Recognition Channel	Third-party providers could accept outputs without rebuilding truth from scratch
Defined fallback obligations	Fallback Integrity	If an access route failed, structured reversion paths existed with legal force
Liability rules and counterparty obligations	Contestability	Dispute routes were legible and enforceable; counterparty obligations defined
Strong Customer Authentication requirements	Refusal Integrity	The system could structurally reject unauthorised state changes
FCA regulatory oversight	Governance Legibility	Rule changes required defined process; affected actors could inspect the boundary

This is not a coincidence of good policy. It is proof that admissibility can be deliberately engineered by mapping regulatory requirements onto structural criteria. Better coordination did not become real because the market spontaneously rewarded it. It became real because each admissibility condition was addressed by design.

The lesson is structural: where gatekeepers control the recognition channel, admissibility must be forced open at each criterion, not merely at the level of general mandate. Open Banking is the clearest available proof that the Admissibility Stack can be satisfied deliberately, criterion by criterion.

Proof verdict: Open Banking is the engineered success case. Every admissibility criterion was satisfied through deliberate design. Admissibility was not earned — it was constructed.

Section XII Primary Proof II: Healthcare Admissibility

The question here is not whether a better continuity architecture can exist. It is why it still fails to become central.

The NHS Summary Care Record is the clearest near-miss. It was a genuine attempt at patient-centred continuity: a record that could travel with the patient across care settings, containing medication history, allergies, and adverse reactions. It satisfied several admissibility criteria: a legal interface through NHS infrastructure, fallback integrity through records that persisted independently of any single provider, and a bounded trust kernel covering defined data types and access rules.

It stalled on the recognition criterion. Clinicians had no binding obligation to consult it and no liability for ignoring it. A GP who missed a drug interaction because they did not check the Summary Care Record faced no formal consequence for the omission. The truth existed. The channel did not require anyone to pass through it.

This is the recognition criterion failing in practice. A continuity system is not admissible because it stores continuity. It becomes admissible when continuity can be safely relied on, acted on, and inspected under consequence — and when ignoring it carries cost. Until that cost exists, routing through the incumbent remains the safer professional choice.

Proof verdict: Criterion failed — Legal/Operational Interface (recognition binding obligation absent). Trust Kernel, Fallback Integrity, and Legal Interface were partially satisfied. Contestability and Recapture Resistance were not tested at scale. The stall was at the recognition channel: truth held, channel not binding.

Section XIII Primary Proof III: Education Admissibility

A truthful learner-state architecture may preserve development more faithfully than transcripts, grades, and annual resets. That alone is insufficient.

The Extended Project Qualification in England illustrates where the failure sits. The EPQ is a richer, more genuine assessment of learner capability than A-level grades: it requires independent research, project management, and extended writing or production. It demonstrably captures capabilities that A-level grades do not. Universities systematically discount it in admissions — not because it is inaccurate, but because it fails the legal or operational interface criterion.

Admissions systems are built around UCAS tariff points. The EPQ does not fit the processing grammar that adjacent consequence-bearing systems use to make decisions. It is correct without being admissible. The recognition channel is closed not by malice but by architectural mismatch: the substitute's output format cannot be processed by the systems whose recognition matters.

Two structural routes exist. Change the processing grammar of adjacent systems — which requires incumbent co-operation or regulatory force. Or build a translation layer that maps EPQ outputs into UCAS tariff terms — which preserves the substitute's truth while making it legible to the incumbent's interface. Neither route is primarily a design problem. Both are admissibility engineering problems.

A learning substitute becomes central only when truthful learner state becomes more institutionally acceptable than the old proxy surfaces used to stand in for it. The EPQ is correct. The recognition channel is closed.

The EPQ is not the problem. It is a visible instance of a broader class: truthful learner-state architectures that fail to map into legacy processing grammars. The failure is not at the substitute level. It is at the interface level.

Where recognition channel engineering is required and incumbents cannot be compelled, a translation layer becomes the structural option: a mechanism that maps the substitute's high-fidelity outputs into the grammar adjacent systems already accept, while preserving the truth kernel behind the interface. For the EPQ: a UCAS-compatible equivalence mapping that expresses genuine capability in tariff points, without destroying the richer assessment behind it. The substitute remains truthful internally. It becomes legible externally.

Proof verdict: Criterion failed — Legal/Operational Interface (processing grammar mismatch). The substitute holds truth. Adjacent consequence-bearing systems cannot process it. Recognition channel engineering is the structural intervention required.

Three Proofs — Criterion Summary

Proof	Type	Criteria Satisfied	Criterion Failed	Required Intervention
Open Banking	Engineered success	All seven — Trust Kernel through Recapture Resistance	None — all satisfied by PSD2 design	Regulatory force opened every criterion deliberately
NHS Summary Care Record	Near-miss — stall	Trust Kernel, Fallback Integrity, partial Legal Interface	Legal/Operational Interface — recognition not binding; no liability for non-consultation	Binding consultation obligation with liability for omission
Extended Project Qualification	Near-miss — stall	Trust Kernel, Refusal Integrity (assessment is sound)	Legal/Operational Interface — UCAS processing grammar mismatch	Translation layer (UCAS tariff equivalence) or upstream grammar change

Section XIV The Admissibility Test

Answer each question for the substitute under examination. A ❌ on any question means the substitute is not yet admissible at that criterion. Two or more ⚠️ means high risk of stall or successor trap. Run this test before committing to a domain. Run it again after the first year of operation. The score should improve over time. If it does not, the successor trap is already operating.

Q	Question	What the Answer Reveals
1	What is the substitute's truth kernel — what exactly is being guaranteed, by whom?	Without a defined kernel, no downstream criterion can operate. If the answer is vague, the trust kernel is absent.
2	What consequence is the substitute trying to carry?	Names the domain and stakes. Consequence-bearing actors must be identified, not just users.
3	Who are the consequence-bearing actors who must rely on it — not just use it?	Identifies who must be convinced for admissibility. Users and consequence-bearing actors are often different populations.
4	What invalid states can the substitute structurally refuse — not just flag, but reject?	Tests Refusal Integrity. If the system can only flag, not refuse, the trust kernel has no enforcement layer.
5	Can that refusal be inspected by a third party without requiring unrestricted access to private state?	Tests the audit surface. Inspectable refusal is the difference between trust and hope.
6	How can affected actors contest outputs through a legible, accessible route?	Tests Contestability. If the answer is 'contact us,' contestability is absent.
7	Who can change the governing rules, how, and within what boundaries?	Tests Governance Legibility. Opaque rule change is a structural failure, even if current operators are trustworthy.
8	If the substitute fails, can participants revert without catastrophic continuity loss?	Tests Fallback Integrity. Risk-averse actors cannot safely rely on a system with no safe exit.
9	Which adjacent consequence-bearing systems must recognise its outputs — and is the recognition channel open?	Tests Legal/Operational Interface. A closed channel means correct but peripheral. This is where most proofs stall.
10	What structural mechanisms prevent the substitute from re-fusing functions or re-exporting burden as it scales?	Tests Recapture Resistance. Where most substitutes eventually fail, even those that cleared all prior criteria.

Scoring and Threshold

Score each question: 0 = absent, 1 = partial, 2 = fully satisfied. Maximum score: 20. Run this test before committing to a domain. Run it again after the first year of operation.

Score	Status	Action
18–20, no zeros	Admissible — proceed to scale	Run the test again at twelve months. Monitor specifically for recapture signals in Q10.
15–17, or any zero	Not yet admissible	Address every zero criterion before proceeding. Zeros in Q4 (Refusal), Q9 (Interface), or Q10 (Recapture) are the highest-risk failures.
Below 15	High successor trap risk	Redesign before launch. The structural failure is upstream of the test. Return to the Admissibility Stack and locate which criterion is absent.

Decision Flowchart

Step 1 — Has the trust kernel been defined? No → Stop. No downstream criterion can operate without it. Yes → Step 2.

Step 2 — Can the substitute structurally refuse invalid states, and can that refusal be inspected? No → Refusal Integrity absent. Address before proceeding. Yes → Step 3.

Step 3 — Are contestability, governance legibility, and fallback integrity all satisfied? No → Identify which criterion is absent and apply the practitioner handoff below. Yes → Step 4.

Step 4 — Can adjacent consequence-bearing systems recognise the substitute's outputs without re-entry or incumbent permission? No → Recognition channel closed. This is the most common stall point. Apply the legal/operational interface intervention. Yes → Step 5.

Step 5 — What prevents recapture as the substitute scales? If the answer is governance goodwill rather than structural enforcement → Recapture resistance is nominal, not real. Design it in before launch, not after. Proceed to the Admissibility Stack audit at the one-year mark.

Practitioner Handoff: If Admissibility Fails

Criterion Failure	Structural Intervention	Primary Failure Mode to Avoid
Trust Kernel absent	Define explicitly: what is being guaranteed, by whom, under what conditions, and how can it be verified? Do not build features until the kernel is named.	All downstream criteria are void. No basis for reliance exists.
Refusal Integrity absent	Encode refusal invariants in the system's core before launch: append-only logs, uniqueness constraints, no operator override, timestamped audit trail. Design refusal before designing features.	The trust kernel has no enforcement layer. The system stores truth without being able to reject incoherence.
Contestability absent	Specify the dispute route before launch: who can contest, through what process, at what cost, within what timeframe. Build the route into the system, not into the support queue.	Unchallengeable authority that erodes trust as scale increases.
Fallback Integrity absent	Define the reversion path before launch: data portability, offline mode, what participants hold after exit. Risk-averse actors will not rely on a system with no safe	Lock-in that prevents adoption by the consequence-bearing actors who matter most.

Criterion Failure	Structural Intervention	Primary Failure Mode to Avoid
	exit.	
Legal/Operational Interface absent	Two routes: force the channel open through regulation (Open Banking model) or build a translation layer that maps the substitute's outputs into the grammar adjacent systems already accept (UCAS tariff mapping model). Name which route before building.	Remaining correct but peripheral indefinitely. This is the most common and most costly failure.
Recapture Resistance absent	Treat admissibility as a preservation requirement, not a launch condition. Design periodic recapture audits, sunset clauses for function fusion, and structural separation of functions that must not re-fuse. Monitor the Correction Decay Spectrum (Paper 7).	The Successor Trap: the old topology returns inside the new system. Success itself rewards the erosion.

Section XV Limitations and Boundary Conditions

The Admissibility Law and Admissibility Stack apply specifically to systems attempting to inherit consequence-bearing institutional reality from an incumbent that holds recognition authority. The framework may not apply, or may apply differently, in the following conditions:

Condition	How Admissibility Changes
Greenfield (no incumbent controls recognition)	Admissibility is structurally easier; the recognition channel must be built, not forced open. Focus shifts to recapture resistance from day one.
Substitute not yet buildable (Paper 5 pipeline incomplete)	Admissibility is moot; complete the replacement pipeline first. Admissibility engineering on an unbuilt substitute is premature.
Incumbent already non-functional (no one uses it)	Institutional inheritance is nearly automatic; but this condition is rare. More commonly, the incumbent remains used because exit costs are high.
Consequence is trivial (low-stakes coordination)	Admissibility criteria can be relaxed. Recapture resistance and refusal integrity remain relevant; recognition channel may not be binding.
Incumbent holds legal monopoly on finality	Superior design alone is insufficient. Admissibility must be forced through regulation, built around the monopoly through translation, or deferred until the monopoly weakens.

The three proofs in this paper (Open Banking, NHS SCR, EPQ) all involve domains where incumbents hold significant recognition authority. Transfer to domains with weaker incumbent recognition, or where no incumbent exists, requires adjustment. The Inter-Kernel Problem applies wherever substitutes must coordinate across domains with separate recognition channels.

Section XVI Conclusion

What this paper adds to the series

Instrument	Definition
The Admissibility Law	A substitute inherits institutional reality not when it surpasses the incumbent in design, but when consequence-bearing actors can safely rely on it without the incumbent's permission
The Admissibility Stack	Seven interdependent criteria in dependency order: Trust Kernel → Refusal Integrity → Contestability → Governance Legibility → Fallback Integrity → Legal/Operational Interface → Recapture Resistance
The Recognition Gap	The condition in which a substitute holds a more truthful record than the incumbent, but consequence-bearing actors cannot act on it without routing through the incumbent's interface
The Recognition Channel	The route through which a substitute's outputs become acceptable to adjacent consequence-bearing systems. Without it, truth remains private, local, or advisory
The Refusal Invariant	A rule the system cannot bypass without ceasing to be what it claims to be. The enforcement layer of the trust kernel
The Successor Trap	A substitute that clears the admissibility threshold initially but decays into the old failure topology as it scales. The default trajectory of success without structural protection
The Inter-Kernel Problem	The condition in which substitutes built separately reproduce the hidden carrier problem at the seam between them. The person must never again become the seam
The Admissibility Test	A ten-question diagnostic instrument for locating which criterion is absent and where the stall is occurring
Practitioner handoff	Structural interventions for each criterion failure, with primary failure modes to avoid

Admissibility engineering is what becomes structurally possible once Paper 7's frame-level Gate is reached. Paper 7 names the sequence; Paper 6 specifies what crossing Move 4 must produce. The gate opens the question of replacement. Admissibility determines whether the answer becomes real.

Falsifiable Prediction

Falsifiable Prediction

A substitute that satisfies all seven criteria of the Admissibility Stack will inherit consequence-bearing centrality in its domain within a bounded period of operation, provided no incumbent holds an unbreakable legal monopoly on finality. If a substitute satisfying all

seven criteria fails to achieve institutional inheritance, the framework requires revision.

The claim would be weakened if a substitute were found that satisfied all seven criteria and still failed to achieve institutional inheritance despite the absence of a legal monopoly. Open Banking, at the criterion level, provides confirmatory pattern evidence for the framework. The NHS SCR and EPQ cases confirm which criterion stalls which substitute. The framework is structurally falsifiable at the criterion level.

The series' own admissibility gap

This framework possesses a recognition gap of its own. It holds a truth kernel — the diagnostic sequence from misclassification through to admissibility — but it does not yet possess a recognition channel. It cannot itself be officially relied upon under consequence. Its refusal invariants are not structurally enforced. Its contestability remains limited to whoever reads it. By its own criteria, it is a substitute that is correct at the edge.

Whether it earns admissibility depends on whether the structures built from it hold — and whether those structures are lighter to inhabit than the ones they are meant to replace.

Design makes the substitute possible.
Migration makes the substitute attractive.
Admissibility makes the substitute real.

The mother's diary was true. It remained peripheral until someone built a recognition channel that did not require the school's permission. That is the work admissibility names.

The substitute becomes historically real not when it exists, but when reality can safely pass through it without asking the old institution for permission.

Jamie Forrester · hello@jamieforrester.com · April 2026

If this maps to a substitute you are building or a reform programme you are designing, you can reach me at hello@jamieforrester.com

Appendix A — Evidence Reference

Tier	Source Type	Contents and Weight
A	First-party regulatory and legal (highest)	PSD2 (EU Payment Services Directive 2); FCA Open Banking guidelines; NHS Summary Care Record programme documentation and evaluation reports; UCAS tariff framework specifications
B	First-party operating	NHS SCR adoption and clinician consultation data; EPQ UCAS tariff value

Tier	Source Type	Contents and Weight
	signals (high)	and university admissions discount patterns; Uber/Airbnb terms of service and platform evolution documentation
C	Public symptom evidence (convergent pattern)	Clinician reporting on SCR consultation rates; university admissions office guidance on EPQ weighting; gig economy driver burden reporting; creator and host community reporting on platform recapture
D	Legal and regulatory consequence	EU Digital Markets Act implications for Open Banking; NHS Digital continuity programme reviews; OFCOM platform investigation findings; App Store litigation outcomes
E	Structural inference (framework-derived)	Applications of the Admissibility Law, Recognition Gap, Refusal Invariant, Successor Trap, and Inter-Kernel Problem. Clearly marked as inference throughout.

Appendix B — Glossary of Named Concepts

Term	Definition and Source
The Admissibility Law	A substitute inherits institutional reality not when it surpasses the incumbent in design, but when consequence-bearing actors can safely rely on it without the incumbent's permission. [Paper 6]
Admissibility	The condition under which a substitute becomes sufficiently trustworthy, inspectable, governable, contestable, and recognisable to carry consequence without the incumbent's re-validation. Distinct from usability, adoption, migration, and legitimacy. [Paper 6]
The Admissibility Stack	Seven interdependent admissibility criteria in dependency order: Trust Kernel → Refusal Integrity → Contestability → Governance Legibility → Fallback Integrity → Legal/Operational Interface → Recapture Resistance. Each depends on the one before it. [Paper 6]
Recognition Gap	The condition in which a substitute holds a more truthful record than the incumbent, but consequence-bearing actors cannot act on it without routing through the incumbent's interface. [Paper 6]
Recognition Channel	The route through which a substitute's outputs become acceptable to adjacent consequence-bearing systems. Without it, truth remains private, local, or advisory. [Paper 6]
Trust Kernel	The base criterion of the Admissibility Stack: what exactly is being guaranteed, by whom, and how can that guarantee be verified? Without a defined kernel, no downstream criterion can operate. [Paper 6, Paper 5]
Refusal Integrity	The capacity of a substitute to structurally refuse invalid states, privileged exceptions, or silent modification — and to make that refusal inspectable by third parties. [Paper 6]
Refusal Invariant	A rule the system cannot bypass without ceasing to be what it claims to be. The enforcement layer of the trust kernel. [Paper 6]
Audit Surface	The inspectable surface through which a third party can verify a

Term	Definition and Source
	substitute's invariants without requiring blind trust or unrestricted access to private state. [Paper 6]
Successor Trap	A substitute that clears the admissibility threshold initially but decays into the failure topology it replaced as it scales. The default trajectory of success without structural protection. [Paper 6]
Inter-Kernel Problem	The condition in which substitutes built separately in adjacent domains reproduce the hidden carrier problem at the seam between them — the person becomes the integration layer. [Paper 6]
Legitimacy Inertia	The structural tendency of incumbents to retain recognition authority through symbolic authority, risk asymmetry, legal incumbency, credential monopoly, and mandatory recognition structures — even after their operational coherence has failed. [Paper 6]
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]
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]

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. [This paper]

Forrester, J. (2026g). The Governed Correction Sequence. SSRN Working 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 6 v3.0 · April 2026