
REALIS WORKING PAPER

How Published Criteria Acquire Force

A Candidate Architecture and Pre-Registered Research Design

Version 0.1 | Issued by Realis Institute



REALIS
INSTITUTE

realisinstitute.com

STATUS: Pre-results working paper, v0.1 skeleton. This document specifies a research program and publishes its commitments before any scoring occurs. The candidate architecture presented here is a hypothesis under investigation. Sections marked GATED await their research gates as defined in the governing Research Architecture, which is available to collaborators on request. Contributions to this document are recorded in the Contribution Record with provenance, and are graded under the same source discipline as any other evidence.

1. The Problem

The 1904 Baltimore fire is standard-setting's founding parable. Engine companies from eight cities arrived to help and could not connect to Baltimore's hydrants, and a national thread specification was published within the year. The parable usually ends there, with the specification treated as the solution.

The record continues. By 1914, 287 of the roughly 8,000 American cities and towns had conforming couplings and hydrants [NISTIR 7158]. Adoption moved where mutual aid compacts had made neighboring departments answerable to one another for interoperability, and it stalled nearly everywhere else. In 1991, 87 years after Baltimore, the Oakland firestorm killed 25 people while mutual aid crews struggled to connect to Oakland's hydrants. Oakland adopted the national standard afterward. Several major departments run non-standard threads today. [Currency check open per Research Architecture §11.6.]

The case generalizes. A criterion can be technically sound, publicly available, and endorsed by every body with standing to endorse it, and still change nothing for decades while the failures it addresses recur. In cases of this kind, publishing the knowledge is not sufficient. Something in the surrounding environment determines whether a published criterion becomes decision-relevant, and that environment has received less deliberate attention than the criteria themselves.

This paper specifies a research program on that environment. The working question: what structural conditions allow a publicly specified criterion to acquire distributed practical force without centralized enforcement, and under what conditions does an available, published, technically sound criterion remain inert?

A scope note is owed at the outset. The program's payloads are published criteria. Practitioners in evidence-based fields describe a resembling condition from the knowledge-production side, where well-supported findings fail to become decision-relevant to the institutions positioned to act on them. Whether that is the same structure with a different payload or a different structure with a similar surface is an open question of the program (Q6), held open deliberately: the equivalence would be worth discovering and is worth nothing asserted.

2. The Candidate Architecture

The program begins from a six-part candidate model: six functions that may have to be present for a published criterion to acquire practical force.

#	Function	What it requires
L1	Legibility	An observer can recognize that a condition described by the criterion may exist, without mastering the full apparatus
L2	Record	Recognition becomes a durable artifact that survives the observer
L3	Portability	The artifact stays intelligible across organizational and professional boundaries
L4	Interest	A receiving actor has an existing exposure the evidence touches: liability, underwriting risk, fiduciary duty, procurement eligibility, professional obligation, statutory oversight
L5	Consequence	The receiver possesses or connects to a lever: premiums, coverage terms, contract conditions, discovery, audit findings, licensing, publication, purchasing
L6	Feedback	Anticipated consequence alters upstream behavior before failure; the criterion becomes decision-relevant in advance

The hypothesis under test is that inertness traces to specific broken links rather than diffuse weakness.

The model's open questions are part of the design, and the coding pass is built to answer them: whether L1 and L3 are distinct functions or one property observed at two boundaries; whether L6 is a link or a name for the whole chain operating; whether the chain is a sequence or an unordered conjunction; and whether the interest function can be engineered, as the design of aviation's confidential reporting system suggests, or only found where it already exists. The last question governs what any eventual construction

guidance can honestly claim.

This model is a candidate in the strict sense. It has not been tested. The program this paper specifies is designed to test it, with the conditions of its failure committed in advance.

3. Method

Matched pairs. The primary evidence design is within-criterion variation: the same specification, propagated in one arm and inert in the other, so the criterion is constant while the surrounding structure varies. Heterogeneous success stories are deliberately demoted; a collection of propagation narratives cannot discriminate between models.

The failure arm is primary. A model of propagation earns little from propagated cases, because a plausible pathway is always reconstructable in hindsight. The evidential weight of this program sits on inert cases: for each, a prediction of which specific functions were absent is registered before coding (Section 5), and the model survives only if honest coding matches those predictions or misses in a pattern that names a better model.

Coding. Each case is coded on thirteen fields covering the criterion, its initial standing, recognition surfaces, evidence form, transmission path, receiving interest, consequence mechanism, feedback, propagation timeline, inert zones, substrate funding, coordination topology, and source grade. Each of the six functions is then scored Present, Absent, or Indeterminate against the coded fields, with citations. Indeterminate is an honest answer; a model that requires forcing Indeterminates to Present is failing.

Coder protocol and its limits. Two coders: the program's author, and a frontier language model coding blind with the framework vocabulary stripped. The model coder is correlated with the framework's authorship, and results from this protocol are graded suggestive rather than confirmatory. A genuinely external coder would upgrade the evidence grade, and recruiting one is an open aim of the collaboration this paper invites.

Stopping rule. A repair path and a stopping rule are maintained as distinct objects. Repairs (renaming a function, splitting one, merging L1 and L3, one addition, timescale conditioning) are allowed and logged. The stopping rule specifies the results under which the six-function model is wrong as such, and it is frozen before any scored pass. A model that can always be fixed cannot fail, and an unfailable model is not being tested.

Source discipline. No number or date appears in this program's outputs without verification against a primary or independent secondary source; unverified fields say so. Claims resting only on advocacy-grade sources cannot anchor a function score. The program studies why records fail to constrain decisions, and cannot run on the failure mode it studies.

4. Case Bank (design; selection finalizes after the prior-art pass)

Matched pairs. Fire hose threads, mutual-aid corridors against holdout departments, including the adapter workaround as a candidate coupling suppressor. International paper sizing, near-universal adoption against the North American installed base. The metric system in the United States, sovereign authorization since 1866 against continuing commercial inertness, with metricated sub-domains (science, medicine, military procurement, export manufacturing) as the internal contrast. SOC 2 against WebTrust, a within-issuer pair isolating procurement demand [UNVERIFIED beyond general shape; verification required before coding].

Mechanism singletons, one per transmission channel. Underwriters Laboratories (insurer-funded testing acquiring force through underwriting terms). The T.J. Hooper line (published precaution shifting the reasonableness baseline without any adoption event). The National Board of Fire Underwriters model building code [UNVERIFIED details]. GAAP through statutory incorporation. PCI-DSS, included deliberately as a hub-coordinated contrast case. Aviation's confidential reporting system, the existence proof that an interest coupling can be designed. Cyber-insurance security requirements, the modern speed benchmark. Credit ratings under regulatory incorporation, the cautionary case in which consequence persisted while accuracy degraded: force is a separate property from the criterion staying true, and any construction guidance must be designed against reproducing this.

Inert singletons. The Dvorak keyboard, the best-documented pure failure, with an adversarial literature on both sides. A second inert case, an unreferenced consensus standard from a formal body, to be selected during the prior-art pass with its prediction registered at selection, before deep reading.

5. Pre-Registered Predictions

For each inert case or arm, the absent functions are named before coding. Hose-thread holdouts: L4 and L5 absent, with adapters suppressing the coupling that mutual aid corridors possessed. The North American paper arm: L4 and L5 absent, switching cost distributed across parties none of whom holds a lever over the others. The metric system: L4 and L5 absent despite maximal initial standing, with metricated sub-domains predicted to show procurement or trade couplings supplying the missing functions. WebTrust: L4 absent, no purchasing population's exposure demanding the attestation. Dvorak: L4 absent and L1 contested. Falsification reading: if honest coding finds all six functions Present in two or more inert cases at independent source grade, the model does not discriminate and the stopping rule governs.

6. Prior Art and Positioning

[GATED: awaits completion of the prior-art sweep per Research Architecture §4. Initial wave, three works verified and characterized: the regulator-intermediary-target model (Abbott, Levi-Faur, Snidal 2017) as nearest known neighbor, addressing adopted rules with a rule-maker anchoring the system where this program addresses pre-adoption inertness without an operating center; Yates and Murphy (2019) as the institutional history of the standard-setting system, upstream of the force question; Cheit (1990) as method precedent for paired case design with a different pairing axis. Remaining sweep: private ordering, path dependence and its counter-literature, the negligence-and-custom line, transnational standards governance, responsive regulation, certification as social regulation, implementation science including the implementation-infrastructure and infrastructure-visibility literature, intermediary role-formation work, and recent standards-body material on implementation infrastructure and modular adoption. The contribution claim is located after the sweep, not before it.]

7. Results

[GATED: awaits the coding pass per Research Architecture §8, conducted against the frozen coding scheme and stopping rule. This section will report function scores with field citations, inter-coder agreement as a result in its own right, the error typology for every miss, the coordination-topology read, and the workaround-suppression read. If a stopping condition fires, this section reports the negative result under the same documentation discipline the program's parent corpus asks of the institutions it

studies.]

8. What This Program Is Not

The program is not adoption advocacy for any standard, including the one whose inertness problem prompted it. It requires no institution to adopt anything; the mechanisms under study operate severally, and several of them, the reasonableness baseline above all, operate without any adoption event occurring anywhere.

The program does not assume a coordinating center and is designed to test whether one is required. If the coded record shows that fast propagation always involved a hub, that tension is published rather than resolved by preference.

The program does not discharge the separate proof obligation concerning lead time, the demonstration that structural readings are available to an instrument before events make failures publicly undeniable. Prototype deployments in live channels are the only known route toward that class of evidence, and no single deployment supplies it.

Prototypes built during the program are experimental apparatus. Deployment and praise upgrade none of them; what a prototype produces is data, and standing is earned through the gates above.

9. Contribution

The program seeks specialists who occupy the channels under study: underwriting, procurement, counsel, journalism, regulation, accreditation, safety, standards work, and the research fields living the knowledge-side version of the condition. Five contribution types are sought, in rising order of commitment: challenges to the candidate model, which are welcome at any strength; cases the bank has missed, propagated or inert, with sources; prior art the sweep should characterize; channel constraints the coding scheme should capture; and prototype venues, where a question set, disclosure mechanism, or transmission pathway can run in a live channel and produce prospective observation.

Three rules govern every contribution. Provenance is recorded: who supplied what, in the Contribution Record. Source grading applies to contributed cases and claims exactly as to any evidence, with commercial interest noted. And the program's ledger discipline extends to contributors: nothing graduates from hypothesis by being convincing, whoever finds it convincing.

Contact Realis Institute to contribute. The full Research Architecture, including the case bank, coding scheme, registered predictions, and stopping rule, is available to serious prospective collaborators.

Appendix A: Contribution Record

Date	Contributor	Contribution	Disposition
(opens with first external contribution)			

Appendix B: Version History

v0.1 skeleton, August 2026. Writable sections drafted; Sections 6 and 7 gated.
Predecessor commitments: Research Architecture v0.1, same date, preserved immutable as the pre-registration record.