

# Signal-Response Distance

Andre Olivier · Soaring Wings

Working Paper · Version 0.3 · 2026-05-10

- [Signal-Response Distance](#)
  - [v0.3 reformulation summary](#)
  - [§0 What SRD does — orientation for the practitioner reader](#)
  - [Abstract](#)
  - [§1 The problem](#)
  - [§2 Defining the construct](#)
  - [§3 Lineage](#)
  - [§4 The SRD diagnostic matrix](#)
  - [§5 AI's role in SRD reduction](#)
  - [§6 How this gets tested](#)
  - [§7 Open questions](#)
  - [§8 Close](#)
  - [Appendix A — Application evidence summary](#)
  - [Appendix B — Reading List](#)
  - [Appendix C — Document control](#)
  - [Appendix D — Capability Stack derivation rule \(promoted from Playbook §D.5\)](#)
  - [Appendix E — Construct Frame Library \(promoted from Playbook §D.5\)](#)
  - [Appendix F — Status thresholds and calibration anchor](#)

# Signal-Response Distance

**A diagnostic construct for AI-era organisational responsiveness Working Paper · Version 0.3**  
**Andre Olivier · Soaring Wings Consulting 2026-05-10**

---

## **v0.3 reformulation summary**

v0.3 is a **review-driven** reformulation of v0.2 — distinct from v0.2's empirical-driven reformulation of v0.1 / v0.1.1. The v0.2 working paper was reviewed at structured peer-review depth in May 2026 (one critical issue, eleven major issues, eighteen minor issues across two independent reviews; review artefacts in [Papers/SRD/Reviews/](#)). v0.3 absorbs the substantive review findings rather than carrying them as appended caveats.

The bar AJ originally set for v0.3 was empirical: a third application against a materially different organisational profile (size, regulation, or stack) than the industrial conglomerate or the bank. That bar is **not yet met**. v0.3 is therefore a *theoretical-revision* reformulation, not an *empirical-third-case* reformulation. The empirical bar is preserved for v0.4 (or v1.0 if the third case yields no new framework-level findings).

**What v0.3 establishes (relative to v0.2):**

- A **strengthened anti-confirmation gate** with post-engagement consequentiality test, explicit fail-conditions, and a named enforcement mechanism (review C1).
- An **operational signal-identification protocol** as a §2.0 prior step, addressing the gap that signal selection precedes  $t_0$  measurement and was previously analyst-relative (review M-5).
- A  **$t_0$  anchoring protocol with reference standard** (“the maturity median of an industry peer group”) and a worked-example chain of reasoning (review M-6 + manual M2a).
- **Threshold operationalisation:** the binary SHORT/LONG matrix axis is replaced with continuous SRD measurement + status thresholds harmonised with the Playbook’s `defaultLongSRDThresholdDays = 30` calibration anchor (review M-3).
- **Worst-signature claim demoted** from “the worst SRD signature” to “a class of pathology of structural significance,” with a comparative ranking framework distinguishing the categorical claim from the operational worst-case aggregation rule (review M-4 + manual M1 first bullet).
- **t-state-space exhaustiveness/exclusivity** addressed with explicit scope statements and an operational rule for ambiguous assignments (review M-8).
- **Inter-analyst reproducibility** treated as a methodological gap with a named v0.4+ test protocol (review M-9).
- **Lineage integration** with time-based competition (Stalk; Stalk & Hout), speed of strategic decision-making (Eisenhardt 1989; Judge & Miller; Baum & Wally), absorptive capacity (Cohen & Levinthal; Zahra & George), information processing (Galbraith), process research methodology (Langley; Pentland), Boyd contemporary scholarship (Osinga), and dynamic-capabilities measurement literature (Helfat et al. 2007; Pavlou & El Sawy 2011; Wilden et al. 2013) — addressing review M-1 + manual M3 + manual m4.
- **Diagnostic-vs-theory disambiguation:** the bifurcation discipline at §5.5 is extended to distinguish diagnostic claims (the matrix, the t-state space, the protocol) from implicit theoretical claims (worst-signature severity, coupling-shape exhaustiveness, profile-driven affordances), with falsifiability commitments stated for each (review M4).
- **Generalisability section** treating SA-only / regulated-only / incumbent-only convenience-sample limitations substantively rather than as a one-line open question (review M5).
- **Open question 5 reinstated** (deliberate non-detection): the v0.2 retirement was a category conflation with the `uncommitted` t-state; the political-non-detection question is genuinely open (review M6).
- **Quadrant naming retired:** “AI Frontier” (Q2) and “Strategic Blind Spot” (Q4) updated to less era-bound and less judgemental labels (review N2 + N3).

**What v0.3 does NOT establish:**

- Predictive utility against ex-ante / ex-post performance (still v1.0 bar; needs three closed cases).
- Generalisability to non-South-African / non-regulated / greenfield / public-sector / small-firm contexts (still requires at least one materially-different case).
- Inter-analyst reliability across two or more independent analysts (still v0.4+; the protocol is specified, the empirical test is pending).
- Confirmation of the worst-signature pathology claim across multiple cases (still v0.4 bar).

The contribution claim is unchanged from v0.1 / v0.2: SRD's contribution is **operational**, not conceptual. Existing literature names the gap between environmental change and organisational response without instrumenting it; SRD provides the measurement convention, the diagnostic matrix, the t-state space, the within-engagement triangulation method, and (now in v0.3) the strengthened falsifiability standard, the operational signal-identification protocol, and the integration with the broader strategic-management speed/responsiveness/absorptive-capacity literatures.

---

## §0 What SRD does — orientation for the practitioner reader

This paper develops SRD technically across §§1–8. For the practitioner reader who wants the value story before the technical anatomy, this section is the orientation.

**SRD is a diagnostic instrument that tells executives where their organisation is structurally too slow to respond to its environment, why the slowness lives there, and what to do about it.**

The construct measures a phenomenon strategy literature has named without instrumenting for forty years — the temporal gap between environmental signal availability and committed organisational response. Strategic failure shows up in post-mortems as “we knew, but not soon enough” or “we knew, but couldn't move.” That is the third failure mode of strategy: not *wrong strategy* (well theorised); not *right strategy, can't execute* (well theorised by the strategy-execution-gap literature); but *right strategy, can't execute fast enough* (largely undertheorised, despite being the binding constraint on most modern strategic responsiveness). SRD operationalises that third mode.

**An SRD diagnostic engagement produces, for the executive audience:** (a) a heatmap placing each measured signal domain on the four-quadrant matrix (Augmented Frontier · Slow-Complex Zone · Instrumentation Gap · Well-Instrumented), surfacing where signal-response distances actually live across the firm's strategic surface; (b) a Capability Stack (Sensing / Seizing / Reconfiguring / Absorptive-acquisition / Absorptive-assimilation / Reflexive measurement) translating per-interval lag measurements into board-recognisable construct language anchored in Teece dynamic capabilities, Cohen-Levinthal absorptive capacity, and Haeckel sense-and-respond; (c) intervention prescriptions matched to where the lag lives (detection vs interpretation vs implementation), with operational caveats about measurement completeness; and (d) a v0.3 binding falsifiability commitment via the Output 3b anti-confirmation counter-check section, including a post-engagement consequentiality follow-up at 6 and 12 months. An engagement without Output 3b cannot claim SRD-diagnostic status (§6.1).

**For an executive, the value proposition is five-fold:** SRD replaces a vague sense of organisational slowness with a measured number; it tells them whether AI investment is well-targeted (AI compresses detection lag on complex signals; AI does not compress governance-bound implementation lag); it surfaces structural couplings between problems that look independent at the per-domain level; it produces an honest read of measurement-infrastructure maturity (most organisations score “Absent” on the reflexive-measurement capability today); and it commits the engagement to a falsifiable post-engagement consequentiality test that distinguishes diagnostic insight from confabulation.

**For the academic reader,** the contribution is operational rather than conceptual: existing literature (§3 lineage) has named responsiveness without instrumenting it; SRD provides the measurement convention, the chain-point decomposition, the diagnostic matrix, the t-state space, the within-engagement triangulation method (§6.5), the strengthened falsifiability standard (§6.1), the operational protocols (signal-identification at §2.0;  $t_0$  anchoring at §2.7; threshold operationalisation at §2.3; inter-analyst reproducibility at §6.7), and the integration with the broader speed/responsiveness/absorptive-capacity literatures (§3 — Stalk; Eisenhardt; Cohen-Levinthal; Helfat / Pavlou & El Sawy / Wilden; Galbraith; Langley/Pentland; Osinga). The paper claims none of these as theoretical novelty; what they amount to collectively is the operational instrument the existing literatures have not consolidated.

The rest of the paper develops the construct (§§1–2), positions it in the literature (§3), specifies the diagnostic matrix and prescriptive output (§4), characterises AI’s role (§5), specifies the testing path (§6), names the open questions (§7), and closes with a statement of what v0.3 establishes vs what is on the v0.4 / v1.0 path (§8).

## Abstract

**Signal-response distance (SRD)** is a measurable construct for the elapsed time between the moment a strategically-relevant signal becomes available in an organisation’s environment and the moment the organisation’s response is committed and observable. The construct is operational, not conceptual: existing strategy literature (Stalk on time-based competition; Boyd on the OODA loop; Haeckel on sense-and-respond; Teece on dynamic capabilities; Eisenhardt on speed of decision-making; Cohen & Levinthal on absorptive capacity; Snowden on signal complexity) has named the importance of organisational responsiveness across forty years, but the temporal gap that responsiveness depends on has been consistently named without being instrumented. SRD provides the measurement, decomposed across three primary intervals (detection, interpretation, implementation), placed against signal complexity in a four-quadrant diagnostic matrix, and anchored in AI-era measurement infrastructure as the technology stack that makes systematic SRD reduction tractable for the first time. v0.3 (this paper) refines the construct against two prospective applications (the industrial conglomerate, federated holding company; the bank, tier-2 South African bank) and against a structured peer-review pass that surfaced one critical and eleven major issues. v0.3 strengthens the construct’s falsifiability standard, operationalises signal identification and  $t_0$  anchoring, harmonises threshold operationalisation with the Playbook’s calibration anchor, distinguishes diagnostic claims from implicit theoretical claims, treats generalisability substantively, and integrates the speed/absorptive-capacity/process-research literatures the v0.2 lineage analysis missed. The contribution remains the operationalisation of a construct existing

literature names without quantifying; v0.3 sharpens what the framework can defensibly claim and what remains pending empirical contact.

---

## §1 The problem

Banking executives describe their strategic failures, almost without exception, as failures of speed. They saw the signal and didn't act fast enough. They acted fast but on the wrong signal. They acted on the right signal but the action arrived after the window closed. Post-mortems read like variants of a single sentence: *we knew, but not soon enough*, or *we knew, but couldn't move*.

This is not a strategy problem in the conventional sense. The strategies were often correct. The problem is mechanical — a property of the organisation's nervous system rather than its brain. Yet the strategy-execution-gap literature focuses on cognitive, political, and capability failures. The temporal infrastructure of signal-detection-and-response remains largely undertheorised, despite being the substrate on which strategic responsiveness actually depends.

The gap between when a signal is available in the relevant environment and when a response lands is what this paper calls **signal-response distance**. The argument is that SRD is the binding constraint on dynamic capability for most firms, that it has been historically difficult to measure or reduce, and that AI-era measurement infrastructure changes both.

v0.3 inherits the v0.2 evidence base — the industrial conglomerate (federated holding company; deployment-gating, divisional CIO veto, divergent platform choices) and the bank (tier-2 South African bank; EA practice maturity, workforce & culture allocation, master data management) — and adds the review-driven theoretical revisions detailed in the v0.3 reformulation summary above.

**Scope statement (new in v0.3, addresses review M5).** SRD as currently developed is grounded in two prospective applications, both South African, both regulated (industrial / financial), both incumbent firms with substantial legacy technology stacks, both tier-1 to tier-2 in scale. The construct is hypothesised to generalise across organisational shapes, but the evidence to date does not test that hypothesis against materially different size, regulation, jurisdiction, or technology-stack profiles. v0.3 addresses generalisability in §7.4. Until further empirical contact, applications outside this profile should treat the framework as transferable-with-caveat rather than validated-across-shapes.

---

## §2 Defining the construct

### §2.0 Signal — operational definition (NEW in v0.3)

The construct depends on  $t_0$  — “the moment a signal becomes available.”  $t_0$  presupposes that we have identified what counts as a signal. v0.2 defined SRD precisely but left signal identification ad hoc. v0.3 makes signal identification an explicit prior step.

**Operational definition of a signal (for SRD purposes):** an environmental event (external to the firm, or intra-organisational per §2.1, or political-deliberation-class per §7.5) that meets all three of the following tests:

1. **Defensible  $t_0$  test.** The event has a defensible environmental moment of availability — a moment at which it could in principle have been registered by a reasonably-instrumented organisation, given accessible data in the relevant environment. Events that exist only in the organisation’s internal cognition (e.g., a strategy team’s spontaneous worry) do not have an environmental  $t_0$  and are not signals for SRD purposes (they are responses to upstream signals, or they are not signals at all).
2. **Named performance outcome test.** The signal corresponds to a named organisational performance outcome that the response is intended to affect. “We should consider expanding into Region X” is not a signal in the SRD sense; “competitor Y has launched a product in Region X that erodes our market share” is. The signal is the environmental event; the named outcome is the performance dimension SRD measures the gap against.
3. **Same-domain time window test.** There exists, in principle, a same-domain time window against which lag is interpretable — either a peer-organisation benchmark, a longitudinal series within the firm, or a discipline-defined threshold (per §2.3 and the Playbook’s `defaultLongSRDThresholdDays = 30` calibration anchor for non-categorised signals). Events for which no time window is in-principle definable are not SRD-tractable signals.

Events failing any of the three tests are **not signals for the SRD diagnostic**. Whether they are interesting strategic objects in some other sense is outside SRD’s scope. The diagnostic operates only on events that pass all three tests.

**Worked test of the protocol against v0.2’s the industrial conglomerate and the bank applications:**

Domain	Test 1 — defensible $t_0$	Test 2 — named outcome	Test 3 — time window	Signal?
the industrial conglomerate D1 (deployment-gating)	yes — divisional adoption gate established at Programme Plan v.1	yes — group-level platform adoption rate	yes — peer holding-company benchmark	✓
the industrial conglomerate D2 (divisional CIO veto)	yes — CIO position emerged in evidence at named meeting	yes — group platform consolidation	yes — discipline-defined (capability-decisions horizon: weeks–months)	✓
the industrial conglomerate D3 (divergent platforms)	yes — divisional platform choices documented at acquisition close	yes — synergies-realised rate	yes — peer holding-company benchmark	✓
the bank D1 (EA practice maturity)	yes — R1.02 release date documented	yes — architectural-coherence outcome	yes — banking-industry peer benchmark	✓
the bank D2 (W&C allocation)	yes — board portfolio-allocation cycle	yes — architectural-stewardship outcome	yes — discipline-defined	✓
the bank D3 (MDM)	yes — MDM-need-detected event in 2022	yes — operational-data-quality outcome	yes — banking-industry peer benchmark	✓
the bank Container ADR (control case)	yes — Container ADR initiation	yes — same as D1	yes — same as D1	✓

All seven measured signal domains across both engagements pass the v0.3 signal-identification protocol — providing post-hoc reassurance that the v0.2 applications were operating on legitimate SRD-tractable signals even though the protocol was not yet specified.

**Operational consequence.** Two analysts running SRD on the same engagement records, applying the v0.3 protocol, will produce more comparable signal sets than two analysts working from v0.2. Inter-analyst reproducibility (§6.7) is gated on this protocol being applied consistently.

## §2.1 The relevant environment

SRD measurement begins not when the organisation detects the signal but when the signal becomes observable in the relevant environment. The relevant environment may be:

- **External** — the firm’s market, regulatory landscape, or competitive surroundings (the canonical case; v0.1’s framing).

- **Intra-organisational** — for **federated organisational structures** (holding companies with autonomous divisions; matrix organisations; joint ventures with shared capabilities), the relevant environment relative to a given decision-making node is the rest of the organisation. Signals propagate across federation boundaries the same way external signals propagate across the firm boundary.

The intra-organisational class was forced on the framework by the the industrial conglomerate application: two of three the industrial conglomerate signal domains had triggers internal to the industrial conglomerate (the holding reacting to its own divisional structure). The framework’s external-environment-only anchor produced forced unknown markings on these domains rather than meaningful measurements — a definitional gap, not a measurement gap.

**Operational rule.** When applying SRD to a federated organisation, the analyst must explicitly identify per signal domain which environment supplies  $t_0$ . Mixing external and intra-organisational  $t_0$  across signal domains in a single analysis is permitted; conflating them within a single signal domain is not.

A signal that exists in the relevant environment for ninety days before the organisation registers it has already imposed ninety days of distance, regardless of how fast the subsequent response is.

## §2.2 Response is committed, observable, and (where applicable) landed

v0.1 said “response is executed and observable.” v0.2 sharpened to “committed and observable” because the the bank application demonstrated that responses can be uncommitted indefinitely without being formally rejected. v0.3 retains the v0.2 sharpening but extends it to address the manual review’s m1 finding: the v0.2 sharpening accommodates **uncommitted** cleanly but loses the  $t_3$ -as-implementation-completion specificity that some pathologies require.

### v0.3 distinguishes three sub-events:

Sub-event	Definition	Relationship to t-states
<b><math>t_2</math> committed</b>	Decision made and resources allocated	$t_2$ in <b>known</b> , <b>pending</b> , or <b>open / in-flight</b> state
<b><math>t_3</math> observable</b>	Response materially visible — implementation is in flight or has landed	$t_3$ in <b>open / in-flight</b> or, more strongly, <b>landed</b> (new sub-state — see §2.4)
<b><math>t_3</math> landed</b> (new sub-state)	Implementation has fully landed and outcomes are observable	$t_3$ specifically in the <b>landed</b> sub-state

The headline SRD measure ( $t_3 - t_0$ ) uses  **$t_3$  observable** by default, consistent with v0.2’s “committed and observable” definition. For pathologies where the gap between commitment and landing is the binding constraint (allocate-but-not-execute), v0.3 admits **landed-SRD** =  $t_3$  landed –  $t_0$  as a derived second-class measure parallel to elapsed-unaddressed time. landed-SRD is uncomputable when  $t_3$  is in any state other than **landed**; for those cases, the headline SRD remains the  $t_3$ -observable form and landed-SRD is reported only when applicable.

This restores the “executed and observable” specificity the v0.1 → v0.2 sharpening lost, without re-introducing the framework-incompatibility with `uncommitted` decisions.

### §2.3 Elapsed time and threshold operationalisation

SRD is a temporal measure, denominated in time units appropriate to the signal type. v0.2 noted units were context-dependent (“milliseconds for transaction fraud, hours for incident response, days for operational shifts, weeks for capability decisions, months for strategic repositioning”) but did not specify the protocol for determining short-vs-long SRD within a single application. This was the M-3 finding from the v0.2 review: without a defensible threshold protocol, the matrix collapses to “structured framework + analyst judgment call.”

**v0.3 threshold operationalisation (NEW; harmonises with Playbook calibration anchor).** SRD is a continuous measurement; status thresholds map continuous measurements to categorical status labels for visualisation and prescription. The protocol:

1. **Per-signal-class threshold.** Each signal class carries a discipline-defined threshold:

Signal class	Default LONG threshold	Provenance
Transaction-time signals (fraud, real-time risk)	seconds–minutes	industry standard (card auth latency benchmarks)
Incident-response signals	hours	banking ops convention; ITIL incident SLAs
Operational signals	days	management-reporting cadence; weekly ops review benchmark
Capability-decision signals	weeks	management-team decision-cycle benchmark
Strategic-repositioning signals	months	board-cycle benchmark
<b>Non-categorised default</b>	<b>30 days</b>	<b>Playbook calibration anchor ( <code>SRDDiagnostic.defaultLongSRDThresholdDays = 30</code> )</b>

1. **Per-engagement override.** The analyst may override the default per-signal-class threshold using either:

- **Peer-organisation benchmark.** Quartile-relative thresholds within a named peer-organisation set (e.g., “longer than 75th percentile of peer banks’ MDM-implementation times”).
- **Longitudinal series.** Standard-deviations relative to a longitudinal series within the client (e.g., “more than  $2\sigma$  longer than the firm’s own historical incident-response distribution”). The override must be named in the engagement record alongside its provenance.

2. **Continuous-measurement output.** SRD is reported as a continuous time measurement. The status mapping (Functional / Partial / Constrained / Absent) is applied to the **derived capability score** (per the Playbook §D.5 Capability Stack derivation rule) rather than directly to SRD time:

Capability score	Status label	Threshold
≥ 65%	Functional	derived from SRD relative to threshold per Playbook §D.5
45–65%	Partial	derived from SRD relative to threshold per Playbook §D.5
15–45%	Constrained	derived from SRD relative to threshold per Playbook §D.5
< 15%	Absent	derived from SRD relative to threshold per Playbook §D.5
insufficient data	Unknown	flag for measurement extension

1. **The four-quadrant matrix becomes a heatmap visualisation.** v0.2’s binary SHORT/LONG axis is replaced with the continuous score. The four-quadrant *naming* is preserved (with v0.3 renaming per §4.1) for diagnostic reference, but quadrant assignment is now soft (a domain near the SHORT/LONG boundary is reported as a heatmap point rather than as a binary quadrant assignment).
2. **Same-domain benchmark requirement is now a method, not a precondition.** v0.2 said SRD was “meaningful only when measured against a same-domain benchmark or longitudinal series.” v0.3 supplies the operational mechanism: the per-signal-class default threshold table provides the benchmark for non-categorised signals, with the per-engagement override available where richer benchmark data exists.

This operationalisation closes the M-3 finding. The SRD diagnostic now has a defensible threshold protocol that does not collapse to analyst judgment; the threshold is explicit, traceable, and overridable with named provenance.

## §2.4 The t-value state space

v0.1 treated t-value recoverability as a binary (known / unknown). v0.2 admitted a richer state space derived from the the bank application. v0.3 retains the v0.2 schema and addresses the M-8 finding (exhaustiveness/exclusivity) with explicit scope statements and an operational rule for ambiguous assignments.

**t<sub>2</sub> state space (5 states):**

State ( $t_2$ )	Definition	Engagement exemplar
<b>known</b>	t-value documented with confidence	Container ADR $t_2 = 2025-10-20$ 12:22:35
<b>unknown</b>	t-value exists but is undocumented or unrecoverable from records	the industrial conglomerate D2 $t_2$ (veto date undated)
<b>pending</b>	decision in front of decision-maker; scheduled review window known	(rare in observed engagements)
<b>uncommitted</b>	response is deferred without schedule; not “won’t happen” but “no commitment to a happening”	the bank D1 $t_2$ , D2 $t_2$ , D3 $t_2$
<b>open / in-flight</b>	decision made; implementation in progress	the industrial conglomerate D1 $t_3$ (early state); Container ADR $t_3$

**$t_3$  state space (6 states; v0.3 adds **landed** per §2.2):**

State ( $t_3$ )	Definition
<b>known</b>	$t_3$ recorded with confidence, prior to or at landing
<b>unknown</b>	$t_3$ exists but is undocumented
<b>pending</b>	implementation scheduled but not yet begun
<b>uncommitted</b>	implementation deferred without schedule
<b>open / in-flight</b>	implementation in progress, not yet landed
<b>landed</b> ( <i>new in v0.3</i> )	implementation fully landed; outcomes observable

**$t_1$  state space (4 states):**

State ( $t_1$ )	Definition	Engagement exemplar
<b>known</b>	discrete recognition event; signal entered awareness on a date	the bank D1 $t_1 = 2025-09-30$ (R1.02 release)
<b>unknown</b>	signal entered awareness but the date is undocumented	the industrial conglomerate examples
<b>pending</b>	signal moving toward awareness; not yet recognised	(rare)
<b>continuous</b>	signal has been continuously known since at least $t_0$ with no discrete recognition event	the bank D3 $t_1$ (MDM continuously known 2022–2025)

**Exhaustiveness statement (v0.3, addresses M-8).** The state spaces are exhaustive **for the sub-population of organisational responses with discrete (or recoverable-as-discrete) t-values**. The schema does not yet accommodate three classes of state encountered in principle but not yet in practice:

Future-state candidate	Description	v0.3 disposition
re-opened-after-close	a $t_2$ that had reached known close, was then formally re-opened	flagged for v0.4; requires empirical contact to operationalise
partially-committed	a $t_2$ that is concurrently committed for some sub-population and uncommitted for others (e.g., committed for digital channel but uncommitted for branch)	flagged for v0.4; the operational rule is to record the sub-population scope and assign per-sub-population t-states
concurrent-pending-and-uncommitted	a $t_2$ that is pending on the formal record but uncommitted in stakeholder behaviour (a scheduled review window the organisation systematically defers)	v0.3 operational rule below

**Mutual-exclusivity rule (v0.3, addresses M-8).** Where an analyst encounters ambiguity between two states, assign **the weaker, defensible label**. v0.2 applied this rule specifically to uncommitted vs absent (preferring uncommitted because available evidence does not establish the strong claim about future state); v0.3 generalises the rule to all state assignments. Specifically:

- pending vs uncommitted (where the formal record says pending but stakeholder behaviour suggests uncommitted): default to uncommitted and note the formal-record discrepancy, *unless* the analyst can establish active intent to honour the pending schedule.
- unknown vs anything else: prefer unknown if the available evidence does not unambiguously establish the assigned state.
- landed vs open / in-flight : prefer open / in-flight unless the analyst can establish that outcome observation has occurred.

The rule reflects the framework’s epistemic discipline: refuse to invent state where evidence is incomplete; prefer the label that does not over-claim.

**Worst SRD signature — v0.3 demotion.** v0.2 named continuous  $t_1$  + uncommitted  $t_2$  as *the* worst SRD signature. v0.3 demotes this to “**a class of pathology of structural significance, observed in one case to date (the bank D3 MDM, 43-month duration); subject to confirmation in further engagements**” per the M-4 finding (claim asserted, not argued;  $n=1$  evidence base).

**Comparative ranking framework for state-space pathologies (NEW in v0.3, partial response to M-4).** The comparative argument the v0.2 paper did not provide:

Signature	Conservatism of bound	Structural-difficulty-of-recovery	Typical elapsed-unaddressed magnitude	Diagnostic priority
continuous $t_1$ + uncommitted $t_2$	conservative bound = elapsed-unaddressed since $t_0$ ; $t_3$ never arrives in observed pathology	high — discretising into recoverable-t-state form requires governance change	multi-year (the bank D3: 43 months)	<b>HIGH</b> — empirically anchored, single case
unknown $t_0$ + uncommitted $t_2$	bound degrades to elapsed-unaddressed since $t_1$	medium — $t_0$ may be recoverable with archaeology	indeterminate	medium (no anchor)
pending $t_2$ that has stalled	bound = scheduled review minus current date	low — re-engaging the pending schedule is mechanical	weeks–months	low–medium
open / in-flight $t_3$ that has stalled	bound = current date minus $t_2$	low–medium — execution review available	months	low–medium
unknown $t_0$ + unknown $t_1$ + uncommitted $t_2$	bound = elapsed-unaddressed since the most recent recoverable anchor (engagement start, prior measurement)	high — full archaeology required	indeterminate	medium

The ranking is provisional. v0.3 admits the comparative ranking as a candidate framework rather than a load-bearing claim. **v0.4 should test the ranking against further engagements; if a third case produces a higher-priority pathology not on the list, the ranking is extensible.**

## §2.5 SRD and elapsed-unaddressed time

When all four t-values ( $t_0$ ,  $t_1$ ,  $t_2$ ,  $t_3$ ) are recoverable in a known/in-flight state, **SRD =  $t_3 - t_0$**  is computable as v0.1 specified.

When  $t_2$  or  $t_3$  is uncommitted, total SRD is uncomputable — but **elapsed-unaddressed time** is computable as a second-class derived metric:

$$\textit{Elapsed-unaddressed time} = (\textit{measurement-point}) - t_0$$

where measurement-point is engagement close, current date, or any future re-measurement. When  $t_0$  is also unknown, fall back to **(measurement-point) –  $t_1$**  with the substitution noted.

**Conditioning-regime note (v0.3, addresses skill review m-4 / manual review m5).** The two forms — (measurement-point –  $t_0$ ) and (measurement-point –  $t_1$ ) — operate under different conditioning regimes:

- The (measurement-point –  $t_0$ ) form is anchored in the environmental availability moment. It is the conservatively-correct measure: always  $\leq$  true SRD (since  $t_0 \leq$  measurement-point  $\leq t_3$ ).
- The (measurement-point –  $t_1$ ) form is anchored in the organisation’s awareness moment. It is also conservatively correct (since  $t_1 \geq t_0$ , so (measurement-point –  $t_1$ )  $\leq$  (measurement-point –  $t_0$ )  $\leq$  true SRD), but it tracks awareness-side uncommitment rather than environment-side uncommitment.

When the substitution from  $t_0$  to  $t_1$  is made, **flag the conditioning-regime change in the engagement record**. Two signal domains compared with mixed conditioning regimes (one  $t_0$ -anchored, one  $t_1$ -anchored) are not directly comparable; convert to a common regime or flag the comparison as conditional.

**For genuinely non-terminating processes (manual review m5 fix, NEW in v0.3).** The bound argument “always  $\leq$  true SRD because  $t_0 \leq$  measurement-point  $\leq t_3$ ” presumes  $t_3$  exists and will arrive. For genuinely non-terminating processes — where an **uncommitted** decision may never close (organisation has chosen not to act, formally or de facto) —  $t_3$  is undefined and the bound argument does not hold. For such cases:

- Elapsed-unaddressed time becomes the **canonical metric** (no SRD reference is meaningful).
- The metric grows linearly with measurement deferral and has no upper bound.
- The diagnostic should flag the case as non-terminating and report elapsed-unaddressed time as the strategic-cost-of-non-termination figure rather than as a conservative SRD bound.

The the bank D3 (MDM) case at 43 months elapsed-unaddressed is conservatively a bound on true SRD (assuming MDM eventually lands); if MDM is treated as a non-terminating-in-practice process (the organisation has not committed to a programme), the 43-month figure is the canonical metric, and it grows linearly with each month of further deferral.

**Fallback rule for triple-uncomputable cases (skill m-5 fix, NEW in v0.3).** When elapsed-unaddressed time is *also* uncomputable —  $t_0$  unknown,  $t_1$  also unknown — the diagnostic falls back to:

***(measurement-point) – (most recent recoverable anchor)***

where the anchor may be engagement start, a prior measurement, or any documented organisational moment after  $t_0$  but before measurement-point. The fallback is reported with the anchor explicitly named; it is the strictly weakest measurement available and should be marked as such.

## §2.6 Decomposition

SRD decomposes across three primary intervals, with an optional fourth:

- $t_0 \rightarrow t_1$  — **detection lag**. Signal becomes available in the environment  $\rightarrow$  organisation notices it.
- $t_1 \rightarrow t_2$  — **interpretation lag**. Signal noticed  $\rightarrow$  meaning correctly understood, decision made.
- $t_2 \rightarrow t_3$  — **implementation lag**. Decision committed  $\rightarrow$  response materially lands and outcomes are observable.
- (Optional) **feedback lag** — time to surface the result of the response back to the organisation. v0.3 retains v0.2's framing: not required by the construct; admitted as a v0.4 candidate.

**Reflexive measurement is a meta-discipline, not a fourth chain point** (preserved from v0.2 / Playbook v0.4.0 reformulation; resolves the v0.1.1 stress-test M5 critique). It wraps all three chain points and produces an engagement-level capability score (per Playbook §D.5 Capability Stack derivation rule). It is captured as a binary engagement-level input (instrumented vs observer-captured), not as a t-interval.

The decomposition is the diagnostic; the total is the headline. A 12-month total SRD with 11 months in  $t_2-t_1$  has very different prescriptive implications from one with 11 months in  $t_1-t_0$ .

## §2.7 Working measurement protocol

The seven-step protocol (v0.2 baseline, v0.3 augmented):

1. **Identify the signal class** (e.g. customer preference shift, fraud pattern, regulatory change, federated-internal directive). Apply the §2.0 signal-identification protocol — events failing the three tests are not signals for SRD.
2. **Identify the relevant environment** for  $t_0$  (external or intra-organisational; v0.3 §2.1).
3. **Establish  $t_0$**  — the earliest moment at which the signal could in principle have been registered, given accessible data in the relevant environment. Apply the  **$t_0$  anchoring protocol** (NEW in v0.3, addresses M-6 + manual M2a — see below).
4. **Measure  $t_1$**  — the moment at which the signal entered organisational awareness; record the  $t_1$  state per §2.4 ( `known` / `unknown` / `continuous` ).
5. **Measure  $t_2$**  — the moment at which an explicit decision was made; record state per §2.4 ( `known` / `unknown` / `pending` / `uncommitted` ). Apply the §2.4 mutual-exclusivity rule for ambiguous cases.
6. **Measure  $t_3$**  — the moment at which the response was committed and observable, with sub-state for `landed` per §2.2 if outcome observation has occurred; record state per §2.4.
7. **Compute SRD** where computable; **compute elapsed-unaddressed time** per §2.5 where  $t_2$  or  $t_3$  is `uncommitted`; **compute landed-SRD** per §2.2 where the gap between commitment and landing is the binding constraint and  $t_3$  has reached `landed`. Apply the §2.5 fallback rule for triple-uncomputable cases.

Where any t-value is `unknown`, **name the source you would need to recover it. Do not estimate.** This discipline is binding under §6.1's anti-confirmation gate.

**$t_0$  anchoring protocol (NEW in v0.3, addresses M-6 + manual M2a)**

$t_0$  is a counterfactual measurement: it asks the analyst to determine, retrospectively, when the signal *could in principle* have been detected. Counterfactuals are a known hard problem (Pearl 2009; Morgan & Winship 2014). The v0.2 protocol did not engage the difficulty; v0.3 introduces a three-component anchoring protocol to make  $t_0$  determinations reproducible across analysts.

**Reference standard for “accessible data” (manual review M2a fix).** v0.2 said  $t_0$  was the earliest moment “given accessible data in the relevant environment.” v0.3 specifies the reference standard:

*“Accessible data” means data that would have been available to an organisation operating at the maturity median of its industry peer group, using publicly available or industry-standard internal data infrastructure. The reference standard is the median, not the leading edge or the laggard, to avoid the framework rewarding either heroic instrumentation or excused under-instrumentation.*

**Three-component  $t_0$  anchoring protocol:**

Component	Specification
<b>(a) Environmental data source</b>	Name the data source(s) consulted to establish $t_0$ . Source attribution must be specific (named publication, named system, named regulatory filing, named market data feed); generic references (“industry data”) are insufficient.
<b>(b) Chain of reasoning</b>	Walk the chain that identifies the earliest-in-principle registration moment within the named source(s). The reasoning must be explicit enough that a second analyst working from the same source could reach the same conclusion.
<b>(c) Alternatives considered and rejected</b>	Name at least two alternative $t_0$ candidates (typically: an earlier-but-too-weak candidate; a later-but-overcautious candidate). Explain why the chosen $t_0$ wins both comparisons.

**Worked example (the bank D3 MDM, 2022-06-13):**

Component	Application
(a) Source	(illustrative — actual source attribution lives in the engagement record) Banking industry MDM-maturity benchmark report 2022; the bank’s own 2021 architecture-review minutes.
(b) Reasoning	The June 2022 industry benchmark report named MDM as the third-most-cited operational-data-quality concern across tier-2 banks; the bank’s 2021 architecture review had already flagged MDM gaps but at lower severity. The June 2022 publication is the earliest moment at which a peer-median-maturity bank would have had unambiguous environmental signal that MDM was a tier-2 priority. Earlier dates would presume above-median instrumentation; later dates would presume below-median instrumentation.
(c) Alternatives considered	(i) 2021-Q3 internal architecture review: rejected because internal moments are not external-environment signals; the 2021 review is a $t_1$ candidate, not a $t_0$ candidate. (ii) 2023-Q2 banking regulator’s MDM guidance: rejected because by 2023 the signal had already escalated past the peer-median-maturity availability moment.

**Reproducibility implication.** The protocol does not eliminate analyst judgment, but it makes the judgment **explicit, auditable, and contestable**. A second analyst arriving at a different  $t_0$  must surface the disagreement on a named axis (different data source consulted; different reasoning chain; different alternative-rejection rationale), which converts a soft variable into a structured-disagreement variable. v0.4 should report inter-analyst agreement rates on  $t_0$  determinations as part of the §6.7 reproducibility test.

### §3 Lineage

SRD does not claim novelty as a theoretical insight. The intuition that the gap between environmental change and organisational response is strategically decisive has been articulated, in different vocabularies, across at least seven traditions. v0.3 expands the v0.2 lineage analysis from four traditions to seven, addressing the M-1 (skill review) and M3 (manual review) findings about missing closest-precursor literatures. The construct’s claim to contribution remains operational, not conceptual: what SRD adds is measurement, an AI-era extension, a strengthened falsifiability standard, and the operational protocols (signal identification,  $t_0$  anchoring, threshold operationalisation, inter-analyst reproducibility) the existing literature has not consolidated.

The seven traditions most directly relevant are summarised below.

Tradition	Core contribution	What SRD inherits	What SRD adds or diverges
<p><b>Boyd — OODA Loop</b> (1976+; Osinga 2007 for contemporary scholarship)</p>	<p>Observe-Orient-Decide-Act as a cycle; whoever cycles faster than their adversary wins. Osinga 2007 extends OODA into organisational and non-adversarial contexts.</p>	<p>Temporal framing; loop duration as the strategic variable; “getting inside” an opponent’s loop; via Osinga, organisation-level applicability.</p>	<p>Original Boyd is military, adversarial, individual-cognitive. SRD is organisational, environmental rather than adversarial, measurement-mediated. SRD also separates signal-to-detection from detection-to-response, which OODA collapses.</p>
<p><b>Stalk — Time-Based Competition</b> (1988; Stalk &amp; Hout 1990 <i>Competing Against Time</i>)</p>	<p>Speed-as-strategic-variable in product-market competition; time-compression as a primary source of competitive advantage.</p>	<p>The recognition that temporal compression is itself the strategic move, not a means to other ends; the explicit framing of speed as a measurable competitive variable.</p>	<p>Stalk’s framework operates at the firm-vs-firm competitive level; SRD operates at the firm-vs-environment-signal level. Stalk measures total cycle time; SRD measures the chain points within the cycle and identifies which specific lag interval is the binding constraint.</p>
<p><b>Haeckel — Sense-and-Respond</b> (1999)</p>	<p>Organisations should shift from make-and-sell to sense-and-respond postures.</p>	<p>The organisational frame; responsiveness as a structural property; the shift from forecasting to detection.</p>	<p>Haeckel was pre-AI and pre-pervasive instrumentation. He could describe the posture but not measure the gap. SRD operationalises and quantifies what Haeckel argued for conceptually, and locates the AI contribution specifically in collapsing the gap.</p>
<p><b>Teece — Dynamic Capabilities</b> (1997, 2007); subsequent measurement literature: Helfat et al. 2007; Pavlou &amp; El Sawy 2011; Drnevich &amp; Kriauciunas 2011; Wilden et al. 2013</p>	<p>Sensing, seizing, reconfiguring as the meta-capabilities that distinguish firms that adapt. The measurement literature post-2007 has measured dynamic capabilities at the firm</p>	<p>The most rigorous adjacent framework; sensing is named as the upstream primitive; the measurement literature establishes that firm-level dynamic-capability</p>	<p>The post-Teece measurement literature measures <i>capabilities at the firm level</i>; SRD measures the <b>temporal gap between sensing and seizing for a specific signal in a</b></p>

Tradition	Core contribution	What SRD inherits	What SRD adds or diverges
	level via survey-based and archival methods.	measurement is feasible and informative.	<b>specific organisational context.</b> The contribution is more specific than “measurement of dynamic capabilities” (already done by Helfat / Pavlou / Wilden); it is the per-signal, per-context temporal-gap measurement. Wilden et al. 2013 is particularly relevant as it appeared in <i>Long Range Planning</i> ; its measurement instrument and SRD’s chain-point decomposition are complementary, not competitive.
<b>Eisenhardt &amp; related — Speed of Strategic Decision-Making</b> (Eisenhardt 1989 <i>Making Fast Strategic Decisions</i> ; Judge & Miller 1991; Baum & Wally 2003)	Empirical literature on decision-speed in TMTs; identification of organisational and environmental antecedents of fast vs slow decision-making.	The empirical foundation that decision-speed <i>varies systematically</i> across firms and contexts; the methodological precedent for studying decision-speed via time-stamped event analysis.	The Eisenhardt-tradition work measures decision-speed (loosely SRD’s $t_1 \rightarrow t_2$ interval) with limited treatment of detection ( $t_0 \rightarrow t_1$ ) and implementation ( $t_2 \rightarrow t_3$ ). SRD’s contribution is the full chain-point decomposition; the literature already established that $t_1 \rightarrow t_2$ measurement is feasible.
<b>Cohen &amp; Levinthal — Absorptive Capacity</b> (1990); Zahra & George (2002)	The firm’s ability to recognise, assimilate, and apply external knowledge; foundational treatment of the upstream side of $t_1 \rightarrow t_0$ (detection lag) in organisational-learning vocabulary.	The structural framing that detection capacity is itself a learnable, investable property; the distinction between potential and realised absorptive capacity (PACAP/RACAP per Zahra & George).	Cohen & Levinthal characterise the <i>capacity</i> ; SRD measures the <i>elapsed-time consequence</i> of the capacity. SRD’s $t_0 \rightarrow t_1$ measurement is the temporal reflection of absorptive-capacity-in-action. The Construct Frame Library (Appendix E) makes the

Tradition	Core contribution	What SRD inherits	What SRD adds or diverges
			linkage explicit at the prescription level: detection-lag-dominant findings tag to the absorptive-capacity construct.
<b>Snowden — Cynefin / Sense-Making</b> (1999+)	Different problem domains require different sense-making postures and response patterns.	The recognition that signal complexity is a variable, not a constant; this directly informs the horizontal axis of the SRD matrix.	Cynefin is qualitative and epistemological. SRD is quantitative and operational. SRD treats signal complexity as one axis (Snowden's contribution) and adds the temporal axis to produce a measurable construct.

**Methodological lineage (NEW in v0.3, addresses skill review M-1).** The chain-point decomposition method draws on process research in management studies (Langley 1999; Pentland 1999; Pentland & Feldman 2007). Process research treats organisational phenomena as sequences of events with identifiable temporal anchors; SRD's  $t_0 / t_1 / t_2 / t_3$  decomposition is a specialised application of the process-research methodological frame to the signal-response phenomenon specifically. Acknowledging the methodological lineage strengthens the framework's claim to methodological rigour: the chain-point decomposition is not invented from nothing; it is process-research-with-a-specific-application.

**Adjacent literatures informing SRD without being primary lineage.** Information processing theory (Galbraith 1977 *Organization Design*) provides the organisational-design substrate within which SRD's chain-point decomposition is embedded — Galbraith's framework treats organisations as information-processing systems and frames responsiveness as an information-processing capacity. SRD operationalises the temporal dimension Galbraith did not formalise. Real options reasoning (Bowman & Hurry 1993; McGrath 1999) contributes the framing that organisational responsiveness is itself a valuable option, the price of which is set by SRD. Decision latency literature (largely a practitioner literature in supply chain and BI) overlaps SRD's  $t_2-t_1$  component but typically excludes the detection lag. The recently emerging Enterprise Nervous System literature (Tech Mahindra and others, 2026) describes the architectural pattern within which SRD reduction becomes implementable.

**Relationship to crisis-response / high-reliability / mindfulness literatures (NEW reading list, v0.4 candidate).** Weick & Sutcliffe (2007) on managing the unexpected; LaPorte & Consolini (1991) on high-reliability organisations; Weick & Roberts (1993) on collective mind in operations — these literatures inform SRD's treatment of signals under uncertainty and the detection-side antecedents of  $t_0 \rightarrow t_1$ . v0.3 flags the linkage; v0.4 should engage substantively if a future application surfaces a high-reliability or crisis-response context.

## §4 The SRD diagnostic matrix

The construct alone is descriptive. Operationalising it for diagnosis requires a second axis: not all signals carry equal SRD, and the appropriate intervention depends on signal complexity as well as response speed. The matrix plots SRD against signal complexity, producing four quadrants with distinct intervention patterns.

Signal complexity is defined here in a Cynefin-adjacent sense. Low-complexity signals have known causal structure and can be detected by threshold-based rules. High-complexity signals require pattern recognition, context integration, or interpretation — they exhibit Cynefin’s ‘complex’ or ‘complicated’ properties and resist simple instrumentation.

**The legibility frame** (carried from v0.2; the bank-derived): “detection happens; it is just not legible” is structurally distinct from “detection does not happen.” A signal can be complex *and detected* yet illegible to the engagement record (the industrial conglomerate’s CIO veto), or simple *and detected* yet structurally absent from the response infrastructure (the bank’s W&C zero-investment). Apply the legibility frame when scoring the X axis.

**v0.3 visualisation: heatmap, not binary quadrant grid.** v0.2 plotted the matrix as a binary SHORT/LONG × Low/High quadrant grid. v0.3 retains the four-quadrant *naming* (with renamings per §4.1 below) for diagnostic reference but plots signal domains on a continuous heatmap derived from the §2.3 threshold operationalisation. A domain near a quadrant boundary is reported as a heatmap point with named threshold and confidence; quadrant assignment is a visual aid, not a binary classification.

	COMPLEXITY →	
	Low	High
FAST SRD (under per-signal-class threshold; Functional / Partial status)	Q1 Well-Instrumented  Banking ex: – card auth – fraud rules – ATM uptime	Q2 Augmented Frontier  Banking ex: – real-time AML – dynamic credit – behaviour anom.
SLOW SRD (above per-signal-class threshold; Constrained / Absent / uncommitted status)	Q3 Instrumentation Gap  Banking ex: – branch NPS – ops RCAs – channel mix – attrition	Q4 Slow-Complex Zone  Banking ex: – customer pref. – regulatory shift – competitor BM – geopolitical fx

## **§4.1 Quadrant naming (revised in v0.3)**

v0.3 retires two of v0.2's quadrant labels per the manual review's N2 and N3 findings:

Quadrant	v0.2 name	v0.3 name	Rationale for change
Q1 (Fast SRD, Low complexity)	Instrumented	<b>Well-Instrumented</b>	minor sharpening; v0.2 label retained meaning but was non-specific
Q2 (Fast SRD, High complexity)	AI Frontier	<b>Augmented Frontier</b>	“AI” as a quadrant label is a 2025-2027 zeitgeist; in 5 years AI will be ambient (like “computing” today). The quadrant’s structural meaning (fast/complex) outlasts the AI-specific framing. AI remains the current-era exemplar named in the body of §4.3.
Q3 (Slow SRD, Low complexity)	Embarrassment	<b>Instrumentation Gap</b>	“Embarrassment” embedded judgement; the structural finding is that the signal is detectable-in-principle but not instrumented-in-practice. The renamed label is descriptive, not evaluative.
Q4 (Slow SRD, High complexity)	Strategic Blind Spot	<b>Slow-Complex Zone</b>	“Strategic Blind Spot” presupposes that slow response on complex signals is a failure. Some organisations <i>strategically</i> operate slowly on complex signals (long-horizon investment, deliberate cultural shift, regulatory caution). The renamed label is neutral; whether the slowness is a pathology depends on the organisation’s strategic posture, not on quadrant placement alone.

**Implication for downstream artefacts.** The Playbook §D.5 will be updated in lockstep (v0.5 release) to reflect the v0.3 renaming. Strategic report templates, the SMP SRDDiagnostic engine, and any client-facing artefacts will need synchronisation; the methodology repo’s v0.5 release coordinates the rename across the artefact graph.

#### §4.2 Q1 — Well-Instrumented (fast SRD, low complexity)

The well-served domain. Signals are simple, response loops are fast, and the technology to maintain them is mature. AI’s contribution here is incremental — it tunes thresholds, reduces false positives, and improves specificity. The marginal return on AI investment in Q1 is real but small.

#### §4.3 Q2 — Augmented Frontier (fast SRD, high complexity)

The strategic prize. Signals require interpretation but the response window is short. This is the quadrant where AI’s pattern-recognition capability does its hardest and most consequential work — converting complex signals into actionable detection in time for response to matter. Real-time adaptive AML scoring, dynamic credit risk on complex commercial portfolios, and behavioural anomaly detection at scale all live here. Q2 capability is rare, hard-won, and disproportionately valuable. Q2 is also the quadrant where the productisation opportunity for SRD-aware tooling is most concentrated.

The “AI” qualifier is era-specific — Q2 capability today is AI-mediated; Q2 capability in earlier eras involved different pattern-recognition technologies (skilled-judgement-at-scale; rule-based systems with human-in-the-loop). The structural property of the quadrant (fast/complex) is era-invariant.

#### §4.4 Q3 — Instrumentation Gap (slow SRD, low complexity)

Signals are simple and the slowness is unforced. The response loop is long not because the analysis is hard but because the measurement infrastructure is missing or the organisational pathway from signal to decision is poorly designed. AI is overkill in Q3 — the intervention is to instrument the signal and shorten the reporting pathway. Most Q3 problems can be solved with disciplined data engineering.

#### §4.5 Q4 — Slow-Complex Zone (slow SRD, high complexity)

Signals are complex and the response loop is long. This is where most strategic failure happens — *and also* where some organisations operate by deliberate design (long-horizon investment, regulatory caution, cultural change). v0.3 retains the v0.2 description of Q4 as where “AI’s highest-leverage use case is” *when* the slow response is unforced — but v0.3 admits that Q4 placement is not, by itself, evidence of pathology. Determining whether a specific Q4 placement is a pathology or a strategic posture is part of the §6.1 anti-confirmation gate’s contribution: the diagnostic must surface whether the slowness is unforced (pathology, intervention indicated) or chosen (posture, intervention contraindicated).

The hardest quadrant to build for: the signals are noisy, the ground truth is delayed or absent, and the response infrastructure depends on organisational change as much as technological capability.

## §4.6 Inter-signal-domain coupling and coupling-shape distinction

v0.1 plotted signal domains independently. v0.2 introduced the coupling-shape distinction (central-mechanism coupling pairwise; operating-model-property coupling cross-domain) and the the bank Pass 2.5 backfill demonstrated that the shapes are not mutually exclusive within an engagement.

Coupling shape	Level of analysis	Mechanism	Prescriptive implication	Exemplar
<b>Central-mechanism coupling</b>	Pairwise (within-engagement)	Multiple signal domains share a decision-making node (one bottleneck stakeholder or one decision body).	Surgical: address the central mechanism; coupling resolves automatically.	the industrial conglomerate's D1 → D3 (Servicely platform build queued behind divisional gates); the bank's D1 → D2 (chapter representation in EXCO); the bank's D1 → D3 (architectural stewardship)
<b>Operating-model-property coupling</b>	Cross-domain (higher-level)	Multiple signal domains share dependence on an operating-model property the organisation lacks.	Architectural: per-domain prescriptions are insufficient; the leverage is at the operating-model property.	the bank's three domains' shared dependence on run-mode-and-maturation governance

**Coupling-list extensibility (NEW in v0.3, addresses manual M1 second bullet).** v0.2 implicitly framed the two coupling shapes as exhaustive. The list is **two coupling shapes empirically observed to date; the list is extensible and other shapes are anticipated.** Candidate shapes for v0.4+ empirical contact:

- **Temporal coupling** — multiple domains share a time window (budget cycle; regulatory reporting cadence) that gates response across all of them simultaneously.
- **Cultural coupling** — multiple domains share a cultural antecedent (risk tolerance; consensus-decision norm) that produces correlated lags across domains without a named central mechanism.
- **Resource coupling** — multiple domains compete for the same scarce resource (architect bandwidth; capital allocation pool); the resource bottleneck couples lags without being a “central mechanism” in the v0.2 sense.

These are not load-bearing in v0.3 (no empirical contact yet). v0.3 names them as anticipated extensions to make the coupling-list-extensibility commitment explicit.

**Coupling shapes coexist at different levels of analysis (v0.2 refinement, retained).** An engagement can simultaneously exhibit: - *Pairwise central-mechanism couplings* between specific domains, and - *Operating-model-property coupling* across all measured domains.

These operate at different levels of analysis and have different prescriptive implications. They are not alternative readings; they coexist. **The diagnostic must run the pairwise pass first** — it surfaces the surgical interventions the operating-model-property reading cannot — **and then check for cross-domain operating-model-property coupling** as a separate, higher-level question. Skipping the pairwise pass and jumping to the operating-model reading produces an architectural prescription that misses available surgical leverage.

**Coupling-line convention (carried from v0.1.1).** When applying the SRD diagnostic to an organisation where two or more signal domains are linked, the heatmap shall annotate the coupling explicitly: a line connecting the coupled domains; a label naming the coupling mechanism; a directional indicator showing upstream/downstream relationship (or bidirectional with explanatory prose).

**Coupling-aware quadrant prescription variants (carried from v0.1.1, refined in v0.2, retained in v0.3).** Where coupling is present, the standard per-quadrant prescription is *replaced* by a coupling-aware variant for the downstream domain. The exact replacement depends on coupling shape:

- For **central-mechanism coupling**, the replacement is targeted at the central node. the industrial conglomerate's Q3-coupled-to-Q2 case prescribed *decoupling* (separate the platform build track from the divisional adoption gates) rather than the standard Q3 *instrument the signal first* prescription, because the signal was already instrumented.
- For **operating-model-property coupling**, the replacement is at the operating-model layer. the bank's three-of-three coupling case prescribed *discretise into chunks that flow through existing build-mode governance* (Path b) as a near-term leverage move, while flagging *develop the missing run-mode governance* (Path a) as the longer-horizon transition. The deliberate-phasing precondition (§6.1 anti-confirmation gate applied to the prescription) is binding before either path issues.

**Application heuristic (v0.2 — sequenced).**

1. **Pairwise pass first.** For each pair of signal domains ( $D_i, D_j$ ), check whether the lag in one is exogenous to the other's mechanism. Where coupling exists, name the central mechanism, indicate direction, and identify the surgical intervention available at that level. *Do not skip this step.* The the bank Pass 2 originally jumped past the pairwise pass to the higher-level operating-model reading, missing the within-engagement central-mechanism couplings until Pass 2.5 backfilled them.
2. **Cross-domain pass second.** After the pairwise pass, check whether all (or most) measured domains share dependence on an operating-model property the organisation lacks. If yes, name the operating-model-property coupling and identify the architectural intervention.
3. **Both shapes can coexist.** The pairwise central-mechanism couplings and the cross-domain operating-model-property coupling are independent findings; both are recorded; both inform prescriptions. The pairwise findings often produce the near-term-actionable surgical interventions; the cross-domain finding produces the longer-horizon architectural-transition question.

## §4.7 Measurement-completeness precondition

The SRD diagnostic produces two distinct outputs of differing epistemic robustness:

1. **Structural differentiation** — the matrix’s ability to surface that two signal domains belong in different quadrants, that a signal domain is unplottable, that two domains are coupled, or that the t-state space exposes a worst-signature signal class. **This output is robust to partial measurement** because the differentiation depends on relative position, not absolute SRD values.
2. **Quadrant-specific prescription** — the framework’s recommended intervention pattern. **This output is conditional on measurement completeness**: if SRD is partially measured, the placement is partial; if the placement is partial, the prescription’s authority is bounded.

Measurement state	Differentiation reliable?	Prescription reliable?
All four t-values measured per domain	yes	yes
Three of four measured (one unknown)	yes	yes if the unknown is $t_0$ and the domain’s other intervals are substantial; conditional otherwise
Two of four measured	yes (for relative position)	conditional — prescription should be marked “subject to measurement completion”
One of four measured	structural placement only	not reliable — prescription withheld
Zero of four measured	unplottable	not applicable
Continuous $t_1$ + uncommitted $t_2/t_3$ (worst signature)	yes (signature itself is the differentiator)	conditional — prescription path is to discretise into recoverable t-state forms

**Operational rule.** When producing the diagnostic’s prescriptive output, the analyst must report the measurement-completeness state per signal domain alongside each prescription. Prescriptions issued without their measurement-completeness caveat overstate the framework’s authority and fall under §6.1’s confirmation-bias prohibition.

## §4.8 Diagnostic use

An SRD diagnostic engagement plots a client’s critical signal domains onto the matrix and produces three outputs:

1. **A heatmap** showing SRD distribution across the firm’s strategic surface, with coupling lines per §4.6 and t-state-space tags per §2.4. Continuous heatmap visualisation per v0.3 §4 (replacing v0.2’s binary quadrant grid).
2. **Localised SRD decomposition** indicating where the lag actually lives — detection, interpretation, implementation, or (in the worst-signature case) sustained-uncommitment over multi-year span.

3. **A prescriptive intervention pattern** matched to each quadrant placement, with measurement-completeness caveats per §4.7 and coupling-shape-appropriate variants per §4.6.

The diagnostic is the artefact that turns SRD from a concept into a productisable consulting offering. The the industrial conglomerate and the bank applications are the canonical worked examples ( `Applications/industrial_conglomerate_Application_2026-05-05.md` ; `Applications/the_bank_Application_2026-05-05.md` ).

---

## §5 AI's role in SRD reduction

The argument that AI matters for organisational responsiveness is widely made and rarely precise. SRD permits a sharper analysis: AI's contribution operates at four distinguishable points along the signal-to-response chain. v0.3 retains the v0.2 chain-point analysis and adds named tooling exemplars per chain-point (skill review m-7) to strengthen the AI-era-infrastructure claim.

### §5.1 Signal availability ( $t_0 \rightarrow t_1$ ): collapsing the detection gap

AI shifts the detection model from query-driven to anomaly-driven: the system flags deviation from learned patterns rather than waiting to be asked. This reduces  $t_1-t_0$  from periodic to continuous and from known-unknowns to unknown-unknowns. Continuous behavioural baselining, unsupervised pattern detection on transaction streams, and embedding-based similarity search on customer behaviour all collapse  $t_1-t_0$  in ways conventional BI cannot.

#### Named tooling exemplars (NEW in v0.3):

- *Continuous behavioural baselining*: time-series anomaly detection libraries (Prophet; ADTK; Datadog Watchdog); transaction-graph embedding systems (PaymentsML platforms).
- *Unsupervised pattern detection*: clustering-based fraud detection (DBSCAN-on-transaction-streams); autoencoder-based AML pattern surfacing.
- *Embedding-based similarity*: customer-behaviour embedding pipelines (sentence-transformer-derived; behaviour2vec); vector-database-anchored similarity-search at customer scale.

These are exemplars, not endorsements. The point is that AI-mediated  $t_0 \rightarrow t_1$  compression is implementable today with named off-the-shelf or modestly-engineered tooling; it is not speculative.

### §5.2 Decision support ( $t_1 \rightarrow t_2$ ): compressing the interpretation cycle

LLMs producing plausible interpretations of detected signals — with linked evidence and proposed response options — compress  $t_2-t_1$  from days to minutes. The risk is hallucination and confident wrongness; the mitigation is human-in-the-loop validation with confidence scoring. Decision support is augmentation, not automation.

**Named tooling exemplars:**

- *Plausible-interpretation generation*: RAG-anchored LLM patterns (Anthropic Claude / OpenAI GPT-class with retrieval over engagement-specific knowledge bases); structured-output patterns enforcing decision-frame compliance.
- *Linked-evidence patterns*: citation-grounded generation (LLM patterns that refuse to assert without source-link); confidence-scored output (calibrated-confidence retrievers).
- *Human-in-the-loop validation*: approval-routing patterns wrapping LLM-generated interpretations with human-confirmation gates before downstream commitment.

**§5.3 Implementation orchestration ( $t_2 \rightarrow t_3$ ): closing the execution loop**

AI's contribution at  $t_2 \rightarrow t_3$  is workflow orchestration — ensuring decisions trigger downstream actions, handoffs happen, and implementation is observable in near-real-time rather than discovered in retrospective audit. Banking implementation latency is often a function of compliance handoffs, control validation, and approval routing — each can be partially automated with carefully scoped AI assistance.

**Named tooling exemplars:**

- *Workflow orchestration*: agent-orchestration patterns (LangGraph; AutoGen; Anthropic agent-SDK) wrapping decision-implementation chains.
- *Compliance handoff automation*: RAG-anchored compliance review (LLM scans implementation plan against control library; flags gaps before human review).
- *Implementation observability*: event-driven instrumentation surfacing implementation state changes (Kafka/Pulsar event streams with LLM-generated commentary on state-change semantics).

**§5.4 Feedback compression: making the loop reflexive**

AI-instrumented systems can measure their own SRD, decompose it, and surface where the loss is occurring. This makes the SRD construct itself reflexive — the platform that reduces SRD also measures it.

**Named tooling exemplars:**

- *Self-measuring SRD platforms*: event-stream-anchored SRD telemetry (the firm's own production systems emit  $t_0 / t_1 / t_2 / t_3$  events; SRD computes itself from telemetry); SMP's `SRDDiagnostic` engine is the analyst-side equivalent.
- *Loss-decomposition surfacing*: automated chain-point breakdown (LLM commentary on which interval is the binding constraint, with prescriptive-intervention generation per the v0.3 framework).

**§5.5 The bifurcation: SRD-as-diagnostic vs SRD-as-instrumentation, and the implicit-theoretical-claim disambiguation**

v0.1 conflated two readings; v0.2 committed to **SRD-as-diagnostic** as the framework's primary identity:

*SRD is a **diagnostic construct**: an analytical instrument for surfacing the latency profile of an organisation’s response to a defined set of signals, applied retrospectively to closed engagements or prospectively to live ones. The diagnostic produces a heatmap, latency decomposition, and quadrant-conditional intervention prescriptions, all subject to the measurement-completeness preconditions of §4.7.*

**SRD-as-instrumentation** remains a separable productisation track:

*An organisation that wishes to operate against SRD as a continuous metric — measuring its own SRD across all signals on an ongoing basis rather than running periodic diagnostics — requires platform infrastructure that v0.3 names as **out of scope**. The instrumentation track is a different and larger product surface (continuous external monitoring, anomaly-driven detection, workflow-integrated implementation observation, reflexive measurement). Per Appendix B, this is the productisation opportunity SWP / SMP can be specified against, but specifying it requires its own paper.*

**Diagnostic claims vs implicit theoretical claims (NEW in v0.3, addresses manual review M4).**

The diagnostic-vs-instrumentation bifurcation is necessary but not sufficient. The framework’s epistemological position requires a second disambiguation: **what the diagnostic descriptively claims** (and the falsifiability standard for those claims) vs **what the diagnostic implicitly theorises** (and the falsifiability standard for those claims).

Claim class	Examples	Falsifiability standard
<b>Diagnostic claims</b> (descriptive: the matrix, the t-state space, the protocol)	“This signal domain has continuous $t_1$ + uncommitted $t_2$ ”; “These two domains are coupled by central-mechanism coupling”; “This domain’s measurement-completeness state is two-of-four.”	The §6.1 anti-confirmation gate (strengthened in v0.3 per below): does the diagnostic surface structure not in engagement notes? Does the surfaced structure prove consequential to engagement outcome (post-engagement consequentiality test)?
<b>Implicit theoretical claims</b> (prescriptive: what the diagnostic recommends and why)	“Worst-signature pathology continuous- $t_1$ -plus-uncommitted- $t_2$ has high diagnostic priority”; “Coupling-shape exhaustiveness: all couplings are central-mechanism or operating-model-property”; “Engagement-profile typology: instrumentation-poor engagements force upstream framework refinement.”	These claims operate at the <b>theoretical</b> level. The v0.3 strengthened anti-confirmation gate addresses descriptive claims well; for theoretical claims, the standard is <b>comparative falsification across engagements</b> : does the claim survive contact with cases that would falsify it?

v0.3 commits to falsifiability for both classes, but with different operational tests:

- **Diagnostic claims** are tested per-engagement via the §6.1 strengthened anti-confirmation gate.
- **Implicit theoretical claims** are tested via cross-case synthesis (§6.3) and the engagement-profile typology comparison (§6.5). v0.3 demotes the strongest theoretical claims (worst-signature ranking; coupling-shape exhaustiveness; profile-driven affordance) to provisional status pending v0.4+ empirical contact.

The §5.1–§5.4 chain-point analysis applies to both readings, but the diagnostic reading uses AI at engagement time (the analyst applying the diagnostic uses AI tools to compress interpretation cycles); the instrumentation reading would use AI in production (the firm’s continuous SRD measurement infrastructure runs AI continuously). v0.3 specifies AI’s role in the diagnostic reading; the instrumentation reading is for v1.0+ work.

---

## §6 How this gets tested

v0.3 is the post-review-pass version. Validation continues through the multi-stage path v0.1 specified, with v0.3 strengthening the anti-confirmation gate, adding inter-analyst reproducibility, and tightening the criteria for cross-case synthesis.

### §6.1 The anti-confirmation gate (strengthened in v0.3)

v0.2’s anti-confirmation gate stated:

*The matrix’s marginal utility is whether plotting the signal domains onto it surfaces structure not derivable from the engagement notes. If the matrix produces only what was already known, the axes need revision.*

The v0.2 review identified this as a **CRITICAL** finding (manual review C1): the gate clears too easily because (a) engagement notes are an unusually weak baseline, (b) the gate does not require surfaced structure to be *correct* or *consequential*, only *present in analysis but absent in notes*. A diagnostic that consistently produced novel-looking-but-wrong structure would still pass v0.2’s gate.

v0.3 strengthens the gate with three additions: a post-engagement consequentiality test, explicit fail-conditions, and a named enforcement mechanism.

#### §6.1.1 The strengthened gate (3 sub-tests)

The v0.3 anti-confirmation gate now consists of three sub-tests that an engagement must pass to claim SRD-diagnostic-status:

##### **Sub-test A — Contribution presence (v0.2 baseline, retained).**

The diagnostic must surface structure not derivable from the engagement notes. If the matrix produces only what was already known, the axes need revision. This is the originally-named v0.2 test.

*Worked exemplar of A passing:* the industrial conglomerate D3 decoupling intervention (not in engagement notes); the bank Container ADR differential (not in engagement notes).

**Sub-test B — Independent verifiability (NEW in v0.3, addresses manual C1 acceptable-fallback).**

The diagnostic’s surfaced structure must be **independently verifiable by a third party** with access to the engagement record but not to the analyst’s reasoning chain. The verification asks: given the engagement record, does the surfaced structure correspond to evidence the third party would accept? Or could the surfaced structure equally be reached from the same record via an unrelated reasoning chain that produces a different intervention?

*Operational rule:* every gate-passing finding must include a named verification anchor — a section of the engagement record (specific document section, specific stakeholder interview, specific data point) that grounds the finding. Findings without a named verification anchor fail Sub-test B.

*Fail-condition for B:* “the gate fails if the surfaced structure can be shown via independent verification to be incorrect or non-actionable, OR if the surfaced structure has no named verification anchor.”

**Sub-test C — Post-engagement consequentiality (NEW in v0.3, addresses manual C1 preferred fix).**

A 6-12 month follow-up requiring the engagement to demonstrate that gate-passing findings were **acted upon AND produced measurable change**. This converts the gate from an output test (did the diagnostic surface novel structure?) to a process-and-output test (did the structure prove consequential?).

*Operational specification:* - Engagement reports must include a **commitment to follow-up review** at 6 and 12 months post-engagement-close. - The follow-up review records, per gate-passing finding: (i) was the finding acted on? (ii) was the action consequential to a named outcome? (iii) did the engagement-period predicted-effect match the observed-effect direction (even if magnitude differs)? - A finding is gate-passing-and-confirmed if all three follow-up checks return positive. - A finding is gate-passing-but-stale if (i) returns negative (not acted on); the finding is not falsified, only un-actioned. - A finding is **gate-failing** if (i) returns positive but (ii) returns negative (acted on, no consequential outcome) — the diagnostic surfaced structure that was not consequential. - A finding is gate-failing if (i) and (ii) are positive but (iii) is negative (acted on, consequential, but in opposite direction) — the diagnostic surfaced consequential structure with mis-predicted direction.

*Retroactive application:* v0.3 commits to retroactive Sub-test C application against the industrial conglomerate and the bank within 12 months of their engagement close. This is the v0.3-binding empirical commitment for the strengthened gate.

**§6.1.2 Enforcement mechanism (NEW in v0.3, addresses manual review m6)**

v0.2 stated the gate was “binding for any future application.” v0.3 specifies the enforcement mechanism:

*The strengthened anti-confirmation gate is binding in the sense that **engagement reports must include an explicit anti-confirmation counter-check section** (Output 3b in the engagement deliverable). The counter-check section names the gate-passing findings, their verification anchors*

*(Sub-test B), and the post-engagement consequentiality commitment (Sub-test C). **Absence of this section disqualifies the engagement from claiming SRD-diagnostic status.** For partner-channel work or external review, the counter-check must be reviewable by an independent peer.*

This makes the gate operationally enforceable. An engagement that does not produce the counter-check section is not a v0.3-compliant SRD diagnostic, regardless of how the engagement is otherwise marketed.

### §6.1.3 Evidence to date

Both the industrial conglomerate and the bank passed v0.2's Sub-test A: - The industrial conglomerate: D3 decoupling intervention was not in engagement notes. - the bank: Container ADR differential was not in engagement notes.

Sub-test B (independent verifiability) is satisfiable in principle for both: - the industrial conglomerate D3 verification anchor: the divisional-adoption-gate sequence in the Programme Plan v.1 (engagement record) corroborates that D3's lag is queue-position rather than detection-failure. - the bank Container ADR verification anchor: the ADR document and its differential characteristics (decision-shape: tactical-not-strategic; reversibility: high; cost-type: operational-not-capital) are recorded in engagement.

Sub-test C (post-engagement consequentiality) is the pending v0.3 commitment. the industrial conglomerate engagement closed 2026-Q1; the bank engagement closed 2026-01-31; the 6-month and 12-month follow-up reviews are scheduled for 2026-Q3 and 2027-Q1 respectively. v0.3.1 (a future patch release) should report the 6-month follow-up findings; v1.0 should report the 12-month follow-up findings.

## §6.2 Retrospective + prospective application

The industrial conglomerate and the bank are both retrospective-plus-prospective: the industrial conglomerate straddles closed-and-in-flight (the engagement is mid-flight; some  $t_3$  values are open / in-flight), the bank is closed-and-uncommitted (most  $t_3$  values are uncommitted).

Both demonstrate the diagnostic produces structure not in the engagement notes when the analyst maintains the strengthened anti-confirmation gate (per §6.1). v0.3 makes the gate binding for any future application: the diagnostic's contribution must be stated explicitly in the engagement's anti-confirmation counter-check section (Output 3b), with the engagement notes serving as the falsification baseline and the post-engagement consequentiality test (Sub-test C) as the binding follow-up commitment.

## §6.3 Prospective measurement

Build SRD measurement into a current or near-future engagement from the start. Define the four  $t$ -values per §2.7 for at least three signal domains; track the decomposition longitudinally; measure the impact of interventions on SRD reduction.

v0.3 adds the v0.2 commitments and extends them: - Explicit identification of the relevant environment (external or intra-organisational) per §2.1 at engagement start, not retrofitted later. - Application of the

§2.0 signal-identification protocol at engagement start; signals that fail the three-test protocol are explicitly excluded from the diagnostic. - Application of the §2.7  $t_0$  anchoring protocol at engagement start;  $t_0$  determinations record the (a)/(b)/(c) components from the protocol. - Pre-commitment to the §6.1 Sub-test C post-engagement consequentiality follow-up at 6 and 12 months.

The prospective measurement is the most demanding stage and produces the strongest evidence; time-to-evidence is months to a year depending on engagement length.

## §6.4 Cross-case synthesis

v0.1 set the bar at three cases. v0.2 reduced to two cases for the v0.2-to-v0.3 transition because two cases (the industrial conglomerate + the bank) had already produced eight grounded findings that justified reformulation. The two-case minimum was sufficient if and only if the second case exercised a different organisational shape than the first; if a candidate second case was structurally too similar, a third was required.

The industrial conglomerate and the bank exercised meaningfully different shapes (federated holding company vs hollowed-out tier-2 bank). Their findings clustered differently (the industrial conglomerate's three on the upstream-environment side; the bank's five on the response-state-space side) and did not redundantly produce the same finding. v0.2 treated this as the empirical justification for the reduced bar, not an in-principle change.

### v0.3 falsifiability statement for the bar reduction (NEW; addresses skill review m-6).

*The two-case bar (the industrial conglomerate + the bank) holds for v0.3 — but the reduced bar carries explicit fail-conditions: - **If a third case fails to differ structurally from the industrial conglomerate or the bank** (i.e., is also South African, regulated, with substantial incumbent stack, tier-1-to-2 in scale), the 3-case bar from v0.1 returns; v0.4 should not absorb findings from a structurally-redundant case unless they are method-level rather than framework-level. - **If a third case differs structurally but produces no new framework-level findings**, the framework is more stable than v0.2 / v0.3 currently claim; the implicit-theoretical claims (§5.5) demoted to provisional status in v0.3 should be re-evaluated as load-bearing. - **If a third case differs structurally and produces new framework-level findings that contradict v0.3 claims**, the contradicted claims must be retracted and the framework reformulated (v0.4 → v1.0 path may need extension).*

## §6.5 Internal-control natural-experiment as a measurement method-when-available (revised in v0.3)

v0.2 introduced the internal-control natural-experiment as a **first-class measurement method**:

*When multiple signal domains exist within a single organisation and at least one exhibits a fast/landing response while others exhibit slow/uncommitted responses, perform a differential analysis on candidate variables (decision-type, stakeholder, cost-type, urgency, scope, reversibility). The*

*variables that differ between fast-landing and slow/uncommitted domains identify the binding constraints on response capacity.*

v0.3 retains the method but **softens the “first-class” framing per manual review m3:**

*The internal-control natural-experiment is a powerful within-engagement triangulation method **when available**. Where the engagement contains a fast-response counterexample, prefer this method over cross-case extrapolation because it controls for organisational shape, governance machinery, and time period in ways cross-case cannot. Where the engagement does not contain a fast-response counterexample, fall back to cross-case synthesis (§6.4). The method is method-of-choice-when-applicable; it is not a peer to cross-case synthesis in the universal-applicability sense.*

The Container ADR vs D1/D2/D3 differential (the bank) is the canonical worked example. the industrial conglomerate could not produce this because the industrial conglomerate did not have a within-engagement fast-response counterexample.

## **§6.6 Engagement-profile typology (revised in v0.3 — claim language softened)**

The industrial conglomerate and the bank produce inverse instrumentation profiles. v0.2 admitted this as a sixth empirical finding (F6) and named it a “typology.” v0.3 retains the empirical observation but **softens the framing per manual review M1 third bullet:** a “typology” implies theoretically-grounded mutual-exclusivity and exhaustiveness; the evidence is two cases. v0.3 reframes as “an empirical observation about the two cases studied to date, suggestive of a typology that further work may confirm.”

Profile dimension	Instrumentation-poor (industrial-conglomerate exemplar)	Instrumentation-rich (the bank exemplar)
Upstream ( $t_0, t_1$ ) recoverability	Weak — frequently unknown	Strong-mixed
Downstream ( $t_2, t_3$ ) recoverability	Strong — decisions and implementations recorded	Weak — $t_2$ uncommitted, $t_3$ undefined across measured domains
Specifically opaque surface	Divisional-CIO-interaction surface	EXCO portfolio-allocation rationale; COO sign-off bandwidth allocation
Dominant sub-interval clustering	Upstream (detection / interpretation)	Downstream (decision-side / interpretation)
Engagement instrumentation richness	Instrumentation-poor	Instrumentation-rich
Resulting framework contribution	Discovery-class engagement finding (D3 decoupling) + 3 v0.1.1 findings	Refinement-class within-engagement triangulation (Container ADR differential) + 6 v0.2 findings

**The profiles are observed-as-inverse in the case-pair.** the industrial conglomerate could not see signals well but could see what was done with them; the bank can see signals well but the responses are uncommitted. This inversion *suggests* a typology rather than instantiating one:

**Provisional observation (subject to v0.4+ confirmation):** *instrumentation-poor engagements may force upstream framework refinement because  $t_0/t_1$  unknowns dominate the analytical work; instrumentation-rich engagements with response-side gaps may force downstream framework refinement because the  $t_2/t_3$  states reveal the response-state space.*

**Confirmation-bias guard (NEW in v0.3, addresses manual review m2).**

*Profile classification at engagement-design stage **must not preclude novel findings outside the predicted profile.** The engagement should be designed to exploit the predicted affordance while remaining alert to evidence that contradicts the classification. A pre-classification of “instrumentation-poor” must not steer the diagnostic exclusively toward upstream framework refinement; if the engagement surfaces strong response-state-space findings, the profile classification should be revised mid-engagement, not preserved at the cost of the novel finding.*

*Better still: run a brief diagnostic before profile classification, with classification then a hypothesis about engagement shape rather than a prescription. v0.4 should test whether mid-engagement profile-revision is operationally viable.*

**Operational rule (v0.2 baseline, retained).** At engagement-design stage, identify the candidate engagement's profile type using available pre-engagement signals (records density; documented decision artefacts; stakeholder-interaction logging discipline). The profile *predicts* which kind of framework contribution the engagement is likely to produce (with confirmation-bias guard above); design the diagnostic to exploit the predicted affordance rather than fighting against it. v0.4 should propose a profile-type-aware diagnostic-design protocol once a third engagement has confirmed or extended the typology.

### **§6.7 Inter-analyst reproducibility (NEW in v0.3, addresses skill review M-9)**

The diagnostic depends on the analyst's identification of (i) signals (per §2.0), (ii) environments (per §2.1), (iii) t-value states (per §2.4), (iv) central mechanisms (per §4.6 for coupling shape), (v) operating-model properties (per §4.6 for cross-domain coupling), (vi) thresholds (per §2.3), and (vii) gate-passing findings (per §6.1).

Each is interpretive. v0.2 did not address whether two independent analysts running SRD on the same engagement records would produce comparable diagnostics. v0.3 acknowledges the gap as the primary methodological limitation of a measurement-claim construct and specifies the test protocol.

**Inter-analyst reproducibility test protocol (specification; v0.4+ empirical execution):**

Component	Specification
<b>Setup</b>	Two independent analysts, both trained on the v0.3 framework, both blinded to the other’s analysis. Both work from the same engagement record (same documents, same access to interviews).
<b>Comparison points</b>	Per signal domain: (a) signal identification (does the §2.0 protocol produce the same signal set?); (b) t-value attributions (do the four t-values match within tolerance?); (c) state assignments (do the t-state classifications match exactly?); (d) coupling-line identifications (do both analysts identify the same coupling mechanism / direction / shape?); (e) quadrant placement on the heatmap (do both analysts place the domain in the same status zone — Functional / Partial / Constrained / Absent?); (f) prescriptive intervention pattern (do both analysts produce the same intervention class for the same domain?).
<b>Agreement metrics</b>	Cohen’s $\kappa$ for categorical assignments (state classifications, coupling shapes, status zones); intraclass correlation coefficient (ICC) for continuous t-value differences (treating the time differences between analysts’ t-values as the metric of interest). Per-component agreement rates reported separately.
<b>Disagreement-resolution protocol</b>	Where analysts disagree, surface the disagreement on a named axis (different data source consulted; different reasoning chain; different alternative-rejection rationale per the §2.7 $t_0$ anchoring protocol). The disagreement-as-structured-output is itself a methodological contribution: it converts a soft variable into a structured-disagreement variable.
<b>Threshold for “reproducible”</b>	$\kappa \geq 0.6$ (substantial agreement) for categorical comparisons; $ICC \geq 0.7$ (good agreement) for continuous comparisons. v0.4 binding test target.

**Informal cross-checks reported (v0.3 baseline).** No formal inter-analyst reproducibility test has been performed against the industrial conglomerate or the bank. Informal cross-checks:

- the industrial conglomerate Pass 2 → Pass 2.5 backfill: the analyst (single) ran a self-replication pass that re-checked pairwise couplings and discovered missed central-mechanism couplings. This is a within-analyst reproducibility check, not inter-analyst, but the fact that the second pass surfaced couplings the first pass missed is a calibration finding: even a single analyst running the framework twice on the same record produces additional structure. Inter-analyst variance is presumptively at least as large as within-analyst variance.

- the bank: no second-analyst review performed.

**Empirical commitment for v0.4+.** The first SRD engagement after v0.3 ships should run the inter-analyst reproducibility test as a binding empirical commitment. The two-analyst review is an explicit cost in engagement scope; the v0.4 engagement contract should price this explicitly.

### §6.8 Methodological note: gate count, not pass count

v0.1's working draft assumed a three-pass measurement structure (Pass 1 t-tables; Pass 2 decomposition; Pass 3 matrix placement and prescription). v0.2 clarified that the three-pass shape was inherited from a different methodology (pattern authoring) and is not derived from SRD's epistemic distinctions. The load-bearing piece is the **stop-gate review cadence**, not the pass count.

The actual epistemic distinctions in SRD measurement are:

- **Measurement** — establish t-values; refuse to invent unknowns.
- **Synthesis** — interpret the pattern of measurements + unknowns across domains; produce coupling-detection and instrumentation-profile findings.
- **Counter-check** — apply the §6.1 strengthened anti-confirmation gate; produce the Output 3b counter-check section per §6.1.2.

These three distinctions can be served by 1, 2, or 3 passes depending on engagement context. The industrial conglomerate and the bank applications used three passes by inherited convention; the bank Pass 3 was deliberately collapsed into a wrap-up because Pass 2 had already produced most of the analytical content.

### **Stop-gate operationalisation (NEW in v0.3, addresses skill review m-3).**

The “stop-gate review cadence” needs operational specification. v0.3 specifies what triggers a gate and how its pass/fail is determined:

Gate	Triggered by	Pass criterion	Fail action
<b>Measurement gate</b>	All in-scope signal domains have t-values (or named-source-for-recovery) recorded per §2.7	All four t-values per domain in <code>known</code> / <code>unknown</code> / <code>pending</code> / <code>uncommitted</code> / <code>open</code> / <code>in-flight</code> / <code>landed</code> / <code>continuous</code> (state-appropriate) per §2.4. No estimated values.	Re-engage with stakeholder/document review; surface the missing-source explicitly to the engagement sponsor before proceeding.
<b>Synthesis gate</b>	Cross-domain pattern interpretation begun per §4.6	Pairwise pass complete (every domain pair examined for central-mechanism coupling); cross-domain pass complete (operating-model-property coupling assessed). Findings recorded per pass.	Re-run the pass with a different analyst (if available) or with explicit second-pass review by the original analyst (per the the bank Pass 2.5 model).
<b>Counter-check gate</b>	Synthesis complete	§6.1 strengthened anti-confirmation gate Sub-tests A and B passed; Sub-test C commitment recorded. Output 3b counter-check section produced.	Re-engage Sub-test A (does the diagnostic surface novel structure?) — if not, the axes need revision and the engagement is incomplete.

The gates are sequential (synthesis cannot pass before measurement; counter-check cannot pass before synthesis) but each gate’s pass criterion is independent. Engagement reports should record gate-pass timestamps as evidence of methodological discipline.

## §7 Open questions

v0.1 listed five open questions. v0.2 retained four (revised), retired one as resolved by application evidence, and added two new ones from v0.1.1. Total in v0.2: six. **v0.3 reinstates the v0.1 question 5 retired in v0.2 (manual review M6), adds one v0.3-specific open question, and treats the generalisability question (v0.1 Q4) substantively per manual review M5.** Total in v0.3: eight.

### §7.1 Signal complexity operationalisation (v0.1 Q1, retained, sharpened)

Whether ‘signal complexity’ is best operationalised on a Cynefin-adjacent scale, on a data-structure scale (structured / semi-structured / unstructured), or on a hybrid measure. v0.2 admitted the legibility frame (“detection happens; it is just not legible”) as a complement to the complexity scale; v0.3 retains

the legibility frame and notes that the §2.0 signal-identification protocol's three-test definition is itself a partial operationalisation. Future versions need a definitional commitment that integrates Cynefin-style complexity, legibility, and the three-test protocol into a single operational rule.

### **§7.2 SRD as scalar vs vector (v0.1 Q2, retained)**

Whether SRD should be measured as a single elapsed time or as a vector of four sub-latencies ( $t_1-t_0$ ,  $t_2-t_1$ ,  $t_3-t_2$ , plus a feedback latency). v0.3 retains the single-elapsed-time headline with decomposition for diagnostic; a feedback-latency component is explicitly out of scope until v0.4.

### **§7.3 Matrix axis as continuous vs categorical (v0.1 Q3, partially resolved in v0.3)**

v0.2 noted the matrix's vertical axis was hybrid (categorical at the boundaries where `uncommitted` lives at long-end as a state-tag; continuous within the measurable range). **v0.3 partially resolves this** via the §2.3 threshold operationalisation: the underlying SRD measure is continuous; status thresholds (Functional / Partial / Constrained / Absent) are applied to the derived capability score; the four-quadrant matrix becomes a heatmap visualisation rather than a binary classification. The hybrid framing is now explicit and operationalised. The remaining open question is whether the per-signal-class default thresholds and the per-engagement override mechanism produce defensible cross-engagement comparability (v0.4+ test target).

### **§7.4 Generalisability across organisational shapes, sectors, sizes, jurisdictions (v0.1 Q4, expanded substantively in v0.3 — manual review M5 fix)**

How SRD relates to organisational size, regulatory environment, jurisdiction, and incumbent technology stack. the industrial conglomerate and the bank are both South African; both regulated; both with substantial incumbent stacks; tier-1 to tier-2 in scale. The two-case bar does not yet test SRD against materially different size/regulation/jurisdiction/stack profiles. v0.3 treats this as a substantive open question rather than a one-line note:

**Specific generalisability concerns the framework does not yet address:**

Boundary case	Concern	v0.4+ test
<b>Public-sector organisations</b>	SRD measurement may behave differently where $t_3$ commitment is delayed by political processes that are not well-modelled by federation-coupling or operating-model-property-coupling	Engage at least one public-sector / government-agency / regulated-utility engagement before claiming public-sector applicability
<b>Smaller organisations (&lt; 500 employees, single-site)</b>	Below a certain organisational size, the t-interval decomposition may be theoretically interesting but operationally meaningless (signals are detected, interpreted, and acted on by overlapping individuals; the chain points collapse)	Engage at least one mid-market engagement (50-500 employees) before claiming size-class generalisability
<b>Cross-jurisdictional organisations</b>	SA-only validation does not test whether the construct generalises to organisations operating across jurisdictions where regulatory environments produce different $t_0$ definitions (a pan-African bank; a multinational with EU operations)	Engage at least one cross-jurisdictional engagement before claiming jurisdiction-class generalisability
<b>Non-incumbent / greenfield organisations</b>	The cases are both incumbent firms with substantial legacy. Whether SRD applies meaningfully to greenfield firms with modern stacks is untested	Engage at least one digital-native / fintech / cloud-native engagement before claiming non-incumbent applicability
<b>Non-regulated industries</b>	the industrial conglomerate and the bank are both regulated (industrial / financial). SRD may behave differently in industries where regulatory cadence does not anchor $t_0$ for major signal classes	Engage at least one non-regulated industry engagement (SaaS; consumer goods) before claiming regulation-class generalisability

**The framework's intended scope (v0.3 explicit):** - **Validated:** South African / regulated / incumbent / tier-1-to-2 organisations. - **Hypothesised to apply (with caveats):** organisations sharing 2 or more of these properties. - **Untested:** organisations sharing 1 or 0 of these properties. - **Boundary cases (likely require framework extensions):** public sector; small firms; greenfield; non-regulated; cross-jurisdictional.

**Hypothesis (still testable, not yet tested at scale):** SRD scales non-linearly with organisational complexity, but specific firms can outperform their size class through deliberate measurement infrastructure investment.

## §7.5 Politically-deliberate non-detection (v0.1 Q5, REINSTATED in v0.3 — manual review M6 fix)

*v0.2 retired this question; v0.3 reinstates it. The v0.2 retirement was a category conflation.*

v0.1 raised the question: *whether the framework requires a fifth quadrant or layer for signals the organisation has chosen not to detect*. v0.2 retired the question on the grounds that t-state expansion in §2.4 absorbed uncommitted decisions. v0.3 corrects this: the retired question concerned **deliberate non-detection** (political/strategic decisions not to look at a signal class); the v0.2 cited resolution concerned **deferred-action-on-detected-signals** (the **uncommitted** t-state where a signal is detected but the response is deferred without schedule). These are different phenomena.

**Deliberate non-detection** — the organisation does not look in a particular direction. The signal exists in some environment but not in the organisation’s “relevant environment” per §2.1 (because the organisation has chosen not to make it relevant). A bank deliberately not surveilling a customer segment for political reasons. A firm strategically ignoring a competitor’s moves. A regulated entity choosing not to monitor a regulatory consultation pre-final-rule because the organisation has decided to wait for final guidance.

**Deferred response (the uncommitted state)** — the organisation has detected the signal and not acted. This is what the v0.2 / v0.3 **uncommitted** state captures.

### v0.3 disposition:

- **The framework currently has no machinery to surface that deliberate non-detection has been chosen** (versus the signal genuinely being outside the relevant environment via §2.1). A v0.4+ extension may be needed.
- Possible extensions: (a) an “anti-quadrant” or “outside the matrix” treatment — signals deliberately not measured; (b) a **chosen-not-to-detect** t-state at the  $t_0$  layer (signal exists in some environment, but the organisation has decided that environment is not its relevant environment); (c) a separate “scope of the relevant environment” decision that is itself a strategic decision and has its own SRD.
- v0.3 does not commit to one of these extensions; the question is genuinely open.
- Reading list (NEW for v0.4 work): strategic ambiguity literature (Eisenberg 1984); strategic ignorance literature (Williams & Stratman 2024 if available; otherwise the broader literature on willful organisational blindness); political non-detection in regulated industries.

## §7.6 Minimum viable measurement infrastructure (v0.1.1 Q6, retained)

What is the **minimum viable measurement infrastructure** an organisation must have for the SRD diagnostic’s prescriptive output to be reliable? §4.7 specifies preconditions but not the infrastructure investment those preconditions imply. v0.4 should propose a minimum-viable instrumentation specification.

## §7.7 Federated organisation propagation (v0.1.1 Q7, retained)

For federated organisations specifically, can SRD measurement happen at the central holding level alone, or must it propagate to division-level instrumentation? the industrial conglomerate’s evidence suggests the latter — the divisional CIO interaction surface, which is structurally invisible to the centre, is the surface on which SRD’s  $t_0 \rightarrow t_1$  chain runs through. v0.4 should engage this question explicitly.

## §7.8 Worst-signature comparative ranking (NEW in v0.3 — partial response to skill M-4)

The §2.4 comparative ranking framework names five candidate worst-signature pathologies and provisionally orders them. The ranking is provisional. v0.4+ should test the ranking against further engagements; if a third case produces a higher-priority pathology not on the list, the ranking is extensible. The “highest-priority pathology” claim of v0.2 has been demoted to “a class of pathology of structural significance”; v0.4 should determine whether the ranking is empirically stable and whether the ordering is a continuum or a typology.

These eight questions remain open; v0.3 names them rather than answering them.

## §8 Close

v0.3’s contribution claim is the same as v0.1 / v0.2’s, sharpened: SRD is a measurable construct for a phenomenon the existing literature has named without quantifying, situated in a diagnostic instrument that maps the construct against signal complexity, and connected to AI-era measurement infrastructure as the technology stack that makes SRD reduction tractable for the first time.

### v0.3 adds, relative to v0.2:

- A signal-identification protocol (§2.0) with three operational tests, addressing the prior step that v0.2 left ad hoc.
- A  $t_0$  anchoring protocol (§2.7) with reference standard (“the maturity median of an industry peer group”), three-component specification, and worked example. Counterfactual measurement is now reproducible in principle if not yet in empirical practice.
- A threshold operationalisation (§2.3) with per-signal-class defaults, per-engagement override mechanism, and continuous-measurement output mapped to the Playbook’s status thresholds. The four-quadrant matrix is now a heatmap visualisation; the binary SHORT/LONG axis is retired.
- A landed-SRD sub-state (§2.2) restoring the v0.1 “executed” specificity that v0.2’s “committed” sharpening collapsed.
- t-state-space exhaustiveness/exclusivity statements (§2.4) — explicit scope (“Type 1 sequential responses with discrete or recoverable-as-discrete t-values”); the mutual-exclusivity rule generalised from “uncommitted-vs-absent” to all state assignments; future-state candidates flagged for v0.4 (re-opened-after-close; partially-committed; concurrent-pending-and-uncommitted).

- A worst-signature demotion and comparative ranking framework (§2.4) — claim demoted from “the worst” to “a class of pathology”; five-row comparative ranking provided as a provisional framework subject to v0.4 confirmation.
- An elapsed-unaddressed-time conditioning-regime note and a triple-uncomputable fallback rule (§2.5) — addresses the manual review m5 and skill m-5 findings.
- A strengthened anti-confirmation gate (§6.1) with three sub-tests: contribution presence (v0.2 baseline); independent verifiability (NEW); post-engagement consequentiality (NEW). Plus a named enforcement mechanism: engagement reports must include an Output 3b counter-check section, absence of which disqualifies the engagement from claiming SRD-diagnostic status.
- An inter-analyst reproducibility section (§6.7) treating the gap as the primary methodological limitation of a measurement-claim construct, with named test protocol (Cohen’s  $\kappa$ ; ICC) and v0.4+ empirical commitment.
- An expanded lineage analysis (§3) integrating Stalk time-based competition; Eisenhardt / Judge & Miller / Baum & Wally on speed of decision-making; Cohen & Levinthal absorptive capacity; Helfat / Pavlou & El Sawy / Wilden et al. on dynamic-capabilities measurement; Galbraith information processing theory; Langley / Pentland process research; Osinga 2007 contemporary OODA scholarship.
- A diagnostic-vs-implicit-theoretical-claim disambiguation (§5.5) — separates falsifiability standards for descriptive claims (per §6.1) from falsifiability standards for theoretical claims (per cross-case synthesis and §6.6 typology comparison); demotes the strongest theoretical claims to provisional status pending v0.4+ empirical contact.
- Quadrant renamings (§4.1) — Q2 “AI Frontier” → “Augmented Frontier”; Q3 “Embarrassment” → “Instrumentation Gap”; Q4 “Strategic Blind Spot” → “Slow-Complex Zone”. The structural meaning of each quadrant is preserved; the labels are now era-neutral and less judgemental.
- Coupling-list extensibility (§4.6) — the two coupling shapes (central-mechanism; operating-model-property) are framed as empirically-observed-to-date with anticipated extensions (temporal coupling; cultural coupling; resource coupling) flagged for v0.4+.
- Engagement-profile typology language softening (§6.6) — the “typology” framing demoted to “an empirical observation about the two cases studied to date, suggestive of a typology that further work may confirm.” Plus a confirmation-bias guard.
- Internal-control-method softening (§6.5) — the v0.2 “first-class measurement method” framing softened to “a powerful within-engagement triangulation method when available.” Cross-case synthesis remains the universal-applicability fallback.
- A substantive generalisability section (§7.4) with explicit boundary-case enumeration and v0.4+ engagement targets per boundary class.
- A reinstated open question 5 (§7.5) on politically-deliberate non-detection — the v0.2 retirement was a category conflation with the `uncommitted` t-state.
- A worst-signature comparative ranking open question (§7.8) flagging the v0.4 confirmation target.
- A stop-gate operationalisation in §6.8 — measurement / synthesis / counter-check gates with explicit pass criteria and fail actions.

- Named tooling exemplars in §5.1–§5.4 — the AI-era-infrastructure claim is now grounded in implementable tooling, not speculation.

The next move is empirical. v0.4 will be written when the construct has survived contact with at least one materially different organisational profile (size, regulation, jurisdiction, or stack) than the industrial conglomerate or the bank, when the v0.3 strengthened anti-confirmation gate has been retroactively applied at 6-month follow-up, and when the §6.7 inter-analyst reproducibility test has been run against a fresh engagement.

v1.0 follows when the diagnostic has demonstrated predictive utility against ex-ante / ex-post performance across at least three fully-closed cases, when the inter-analyst reproducibility test has reported  $\kappa \geq 0.6$  /  $ICC \geq 0.7$ , and when the §7.4 generalisability section can be re-written from boundary-case-pending to boundary-case-tested.

Until then, this paper is a hypothesis with structure, two empirical applications, a strengthened falsifiability standard, and a substantively-revised lineage and operational-protocol set. It is not yet a framework with proven predictive utility; v0.3 sharpens what can be defensibly claimed and what remains pending empirical contact.

---

## **Appendix A — Application evidence summary**

Two applications, both 2026:

Application	Structural shape	Primary findings	Marginal utility
<p><b>the industrial conglomerate</b> ( Applications/ industrial_conglomerate_Application_2026-05-05.md )</p>	<p>Federated holding company; tier-1; JSE-listed industrial conglomerate</p>	<p>D3 decoupling intervention; three v0.1.1 framework findings (intra-organisational <math>t_0</math>; inter-domain coupling; measurement-completeness precondition)</p>	<p>Discovery-class — D3 decoupling not in engagement notes</p>
<p><b>the bank</b> ( Applications/the bank_Application_2026-05-05.md )</p>	<p>Tier-2 South African bank; hollowed-out architectural posture; greenfield digital + legacy core</p>	<p>Container ADR differential identifying decision-shape compatibility with build-mode operating model as binding constraint; Run-the-Bank reframing; D3 MDM 43-month elapsed-unaddressed time; six v0.2 framework findings (state-space; continuous-<math>t_1</math>; elapsed-unaddressed-time; internal-control method; coupling-</p>	<p>Framework-refinement-class — Container ADR differential not in engagement notes; one structural reframing; one quantification</p>

Application	Structural shape	Primary findings	Marginal utility
		shape coexistence; engagement- profile typology)	

Both applications passed the §6.1 anti-confirmation gate Sub-tests A and B. Sub-test C (post-engagement consequentiality) follow-up scheduled per §6.1.3.

### **A.1 — Finding-confidence taxonomy (NEW in v0.3, addresses skill review M-2)**

The eight v0.2 findings + v0.3-specific additions are not all at the same confidence level. v0.3 explicitly classifies them:

Finding	Origin	Confidence level	Rationale
Intra-organisational $t_0$ class	the industrial conglomerate	Load-bearing across cases	industrial-conglomerate-derived; the <i>complement</i> to the bank's external $t_0$ is structurally different from the industrial conglomerate's intra-organisational $t_0$ . Both cases together establish the external/intra-organisational distinction.
Inter-domain coupling (general principle)	the industrial conglomerate	Load-bearing across cases	industrial-conglomerate-observed; the bank also exhibits coupling (in different shape). The <i>existence</i> of inter-domain coupling is established across cases.
Measurement-completeness precondition	the industrial conglomerate	Load-bearing across cases	industrial-conglomerate-driven (5 of 12 t-values unrecoverable); the bank also exhibits partial measurement (uncommitted $t_2/t_3$ ). The <i>principle</i> applies across cases.
t-state space expansion (uncommitted; continuous; pending; open / in-flight)	the bank	Load-bearing across cases	the bank-derived; the states accommodate phenomena the industrial conglomerate also exhibits but did not surface. The expanded state space is structurally necessary across cases.
Elapsed-unaddressed time as derived metric	the bank	Case-specific, provisionally generalised	the bank-derived (43-month MDM example); the metric is mathematically defined and applicable wherever uncommitted states arise, but its diagnostic utility has been demonstrated in one case.
Coupling-shape distinction (central-mechanism vs operating-model-property)	the industrial conglomerate + the bank	Load-bearing across cases	The <i>distinction</i> is observable across both cases (the industrial conglomerate central-mechanism; the bank both shapes coexisting). Coexistence is the bank-derived.
Internal-control natural-experiment method	the bank	Case-specific, method when available	the bank-derived (Container ADR differential); the industrial conglomerate could not produce this because no fast-response counterexample. The method's <i>applicability</i> is bounded; v0.3 §6.5 softens "first-class" framing accordingly.

<b>Finding</b>	<b>Origin</b>	<b>Confidence level</b>	<b>Rationale</b>
Engagement-profile typology	the industrial conglomerate + the bank (inverse points)	Reflexive / meta — provisional	Both cases instantiate inverse points on the proposed dimension; the typology structure is suggestive but not confirmed. v0.3 §6.6 reframes from “typology” to “empirical observation suggestive of a typology.”
Worst-signature pathology (continuous $t_1$ + uncommitted $t_2$ )	the bank D3	Case-specific, provisional	One case (the bank D3, 43 months); v0.3 §2.4 demotes claim and provides comparative ranking framework.
<b>(v0.3 NEW) Comparative ranking framework for worst-signature</b>	v0.3 theoretical	Provisional; empirical confirmation pending	Five-row ranking introduced in §2.4; v0.4+ should test against further engagements.
<b>(v0.3 NEW) Strengthened anti-confirmation gate</b>	v0.2 review C1 + v0.3 theoretical	Specified; empirical confirmation pending	Three sub-tests specified in §6.1; Sub-test C (post-engagement consequentiality) commitment binding for the industrial conglomerate and the bank at 6/12-month follow-up.
<b>(v0.3 NEW) Inter-analyst reproducibility protocol</b>	v0.2 review M-9 + v0.3 theoretical	Specified; empirical confirmation pending	Test protocol specified in §6.7; v0.4 binding test target.
<b>(v0.3 NEW) Threshold operationalisation</b>	v0.2 review M-3 + Playbook calibration anchor	Specified; harmonised with Playbook; cross-engagement comparability pending	Per-signal-class defaults + per-engagement override mechanism; Playbook <code>defaultLongSRDThresholdDays = 30</code> as non-categorised default.
<b>(v0.3 NEW) Signal-identification protocol</b>	v0.2 review M-5 + v0.3 theoretical	Specified; tested post-hoc against the industrial conglomerate and the bank (all signal domains pass)	Three-test protocol in §2.0; reproducibility across analysts pending v0.4 §6.7 test.

### Confidence levels:

- **Load-bearing across cases:** observed/derivable from both the industrial conglomerate and the bank, OR derived from one case but corroborated by the structural difference of the second case. Treat as established for v0.3 purposes.

- **Case-specific, provisionally generalised:** observed in one case, plausible elsewhere, awaiting third-case test. Treat as candidate for v0.4 confirmation.
- **Reflexive / meta:** derived from contrast between cases rather than from either alone. Treat as provisional pending v0.4+ contact with materially different organisational profile.
- **Specified; empirical confirmation pending:** v0.3 theoretical contributions specified at protocol-level, awaiting empirical execution. Treat as binding for v0.3 publication; binding-with-evidence pending v0.4.

---

## Appendix B — Reading List

v0.3 expands the v0.2 reading list to absorb the lineage-integration findings and the v0.4 work targets. Major additions:

**Time-based competition tradition (added in v0.3 — addresses skill M-1):** - Stalk, G. (1988). Time — the next source of competitive advantage. *Harvard Business Review*, 66(4), 41–51. - Stalk, G. & Hout, T. M. (1990). *Competing Against Time: How Time-Based Competition Is Reshaping Global Markets*. Free Press.

**Speed of strategic decision-making tradition (added in v0.3 — addresses skill M-1):** - Eisenhardt, K. M. (1989). Making fast strategic decisions in high-velocity environments. *Academy of Management Journal*, 32(3), 543–576. - Judge, W. Q. & Miller, A. (1991). Antecedents and outcomes of decision speed in different environmental contexts. *Academy of Management Journal*, 34(2), 449–463. - Baum, J. R. & Wally, S. (2003). Strategic decision speed and firm performance. *Strategic Management Journal*, 24(11), 1107–1129.

**Absorptive capacity tradition (added in v0.3 — addresses skill M-1):** - Cohen, W. M. & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128–152. - Zahra, S. A. & George, G. (2002). Absorptive capacity: A review, reconceptualization, and extension. *Academy of Management Review*, 27(2), 185–203.

**Information processing theory (engaged substantively in v0.3 — was name-checked in v0.2):** - Galbraith, J. R. (1977). *Organization Design*. Addison-Wesley.

**Process research methodology (added in v0.3 — addresses skill M-1):** - Langley, A. (1999). Strategies for theorizing from process data. *Academy of Management Review*, 24(4), 691–710. - Pentland, B. T. (1999). Building process theory with narrative: From description to explanation. *Academy of Management Review*, 24(4), 711–724. - Pentland, B. T. & Feldman, M. S. (2007). Narrative networks: Patterns of technology and organization. *Organization Science*, 18(5), 781–795.

**Dynamic capabilities measurement literature (added in v0.3 — addresses manual M3):** - Helfat, C. E., Finkelstein, S., Mitchell, W., Peteraf, M. A., Singh, H., Teece, D. J. & Winter, S. G. (2007). *Dynamic Capabilities: Understanding Strategic Change in Organizations*. Blackwell. - Pavlou, P. A. & El Sawy, O. A. (2011). Understanding the elusive black box of dynamic capabilities. *Decision Sciences*, 42(1), 239–273. - Drnevich, P. L. & Kriauciunas, A. P. (2011). Clarifying the conditions and limits of the contributions of ordinary and dynamic capabilities to relative firm performance. *Strategic Management*

*Journal*, 32(3), 254–279. - Wilden, R., Gudergan, S. P., Nielsen, B. B. & Lings, I. (2013). Dynamic capabilities and performance: Strategy, structure and environment. *Long Range Planning*, 46(1–2), 72–96.

**Boyd contemporary scholarship (added in v0.3 — addresses manual m4):** - Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge.

**Real options reasoning (engaged substantively in v0.3):** - Bowman, E. H. & Hurry, D. (1993). Strategy through the option lens: An integrated view of resource investments and the incremental-choice process. *Academy of Management Review*, 18(4), 760–782. - McGrath, R. G. (1999). Falling forward: Real options reasoning and entrepreneurial failure. *Academy of Management Review*, 24(1), 13–30.

**Counterfactual measurement methodology (added in v0.3 — addresses skill M-6):** - Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*. (2nd ed.) Cambridge University Press. - Morgan, S. L. & Winship, C. (2014). *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. (2nd ed.) Cambridge University Press.

**Inter-rater reliability methodology (added for v0.4 work — addresses skill M-9):** - Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1), 37–46. - Krippendorff, K. (2018). *Content Analysis: An Introduction to Its Methodology*. (4th ed.) Sage. - Shrout, P. E. & Fleiss, J. L. (1979). Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin*, 86(2), 420–428.

**Strategic ambiguity / non-detection (NEW for v0.4 work — addresses manual M6):** - Eisenberg, E. M. (1984). Ambiguity as strategy in organizational communication. *Communication Monographs*, 51(3), 227–242. - McGoey, L. (2012). The logic of strategic ignorance. *British Journal of Sociology*, 63(3), 553–576.

**Crisis-response / high-reliability / mindfulness (NEW for v0.4 work):** - Weick, K. E. & Sutcliffe, K. M. (2007). *Managing the Unexpected: Resilient Performance in an Age of Uncertainty*. (2nd ed.) Jossey-Bass. - LaPorte, T. R. & Consolini, P. M. (1991). Working in practice but not in theory: Theoretical challenges of “high-reliability organizations.” *Journal of Public Administration Research and Theory*, 1(1), 19–48. - Weick, K. E. & Roberts, K. H. (1993). Collective mind in organizations: Heedful interrelating on flight decks. *Administrative Science Quarterly*, 38(3), 357–381.

**Federated organisations / ambidexterity (carried from v0.2):** - Argyris, C. & Schön, D. A. (1978). *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley. - Galbraith, J. R. (2014). *Designing Organizations: Strategy, Structure, and Process at the Business Unit and Enterprise Levels*. Jossey-Bass. - O’Reilly, C. A. & Tushman, M. L. (2008). Ambidexterity as a dynamic capability: Resolving the innovator’s dilemma. *Research in Organizational Behavior*.

The v0.1 / v0.2 reading list (sense-and-respond, OODA, Cynefin, decision latency, design science methodology) remains operative for v0.3 and forward.

## Appendix C — Document control

- **v0.3** (this version) — 2026-05-10. Review-driven reformulation absorbing one critical and eleven major findings from structured peer review of v0.2 (manual + skill-driven; reviews in *Reviews/*). Strengthened anti-confirmation gate; signal-identification protocol;  $t_0$  anchoring protocol; threshold operationalisation harmonised with Playbook calibration anchor; quadrant renamings; substantive generalisability section; reinstated open question 5; comparative ranking framework for worst-signature; inter-analyst reproducibility protocol. Eighteen minor findings also addressed.
- **v0.2** — 2026-05-05. Empirical-driven reformulation absorbing eight findings from the industrial conglomerate and the bank applications. SRD-as-diagnostic identity committed. Methodological note on gate-count-not-pass-count added.
- **v0.1.1** — 2026-05-05. Amendment summarising three industrial-conglomerate-derived findings; superseded by v0.2.
- **v0.1** — early-mid 2026. Initial working paper. All v0.1 claims that survive are restated in v0.2's vocabulary (and v0.3's, where v0.3 changes them); v0.3 supersedes v0.2 as canonical.

**Status:** hypothesis with structure, two empirical applications, a strengthened falsifiability standard, and a substantively-revised lineage and operational-protocol set. Empirical validation against a third materially different organisational shape pending; inter-analyst reproducibility test pending; post-engagement consequentiality follow-up scheduled.

**Productisation track:** SRD diagnostic specified per v0.3; SRD-as-instrumentation remains out of scope for the working paper per §5.5. Specification of the instrumentation track requires its own paper.

**Versioning relationship to the Playbook:** Playbook §D.5 will be updated to v0.3 reference in lockstep (Playbook v0.5 release coordinated with this v0.3 paper release). Quadrant renamings (§4.1), threshold operationalisation (§2.3), strengthened gate (§6.1), and new sections (§2.0, §6.7) propagate into the Playbook's §D.5 summary.

**Author:** Andre Olivier, Soaring Wings Consulting. Drafted with structured AI assistance. Empirical grounding: two engagements (the industrial conglomerate, the bank) with combined twelve-plus measured t-values, eight grounded framework findings absorbed in v0.2, one critical and eleven major review findings absorbed in v0.3, two engagement-actionable interventions delivered (the industrial conglomerate D3 decoupling memo; the bank three-prescription memo with Step 0 transition-readiness review precondition).

**Next revision target:** v0.4 after at least one application against a materially different organisational profile (size, regulation, jurisdiction, or stack), and after the §6.7 inter-analyst reproducibility test has been run.

## Appendix D — Capability Stack derivation rule (promoted from Playbook §D.5)

The Playbook's §D.5 specifies the operational translation of SRD measurement into board-grade capability language via a six-bar Capability Stack. v0.3 promotes the derivation rule to the source paper for completeness (the Playbook is the canonical operational spec; the source paper now matches the spec's depth on this commitment).

### Six capabilities, derived deterministically from SRD measurement:

Capability	Derivation	Theoretical anchor
<b>Sensing</b>	Inverse of $t_0 \rightarrow t_1$ (detection lag)	Teece — sensing primitive
<b>Seizing</b>	Inverse of $t_1 \rightarrow t_2$ (interpretation lag)	Teece — seizing primitive
<b>Reconfiguring</b>	Inverse of $t_2 \rightarrow t_3$ (implementation lag)	Teece — reconfiguring primitive
<b>Absorptive · acquisition</b>	Inverse of $t_0 \rightarrow t_1$ + signal-class filter (whether external knowledge)	Cohen-Levinthal · Zahra-George (PACAP — potential absorptive capacity)
<b>Absorptive · assimilation</b>	Inverse of $t_1 \rightarrow t_3$ (interpretation + implementation latency for external knowledge)	Cohen-Levinthal · Zahra-George (RACAP — realised absorptive capacity)
<b>Reflexive measurement</b>	Engagement-level binary: instrumented capture vs observer capture	SRD-native

**Aggregation across measured domains:** worst-case (max lag). Strategic responsiveness is bounded by the slowest channel; the capability profile must reflect the structural ceiling honestly. Per-domain variance and outlier identification belong in the operational companion report. Average and median may be reported there for completeness, but the strategic report's bars are worst-case by design.

**Status thresholds** (calibration anchor: `SRDDiagnostic.defaultLongSRDThresholdDays = 30`):

Capability score	Status
$\geq 65\%$	<b>Functional</b>
45–65%	<b>Partial</b>
15–45%	<b>Constrained</b>
$< 15\%$	<b>Absent</b>
insufficient data	<b>Unknown</b> (do not render bar; flag for measurement extension)

**Reflexive measurement is meta.** Not derivable from t-intervals — requires a separate engagement-level input recording whether the organisation captured its own SRD via instrumentation, or whether the consulting team captured it via observation. Observer-capture engagements show “Absent” by definition. Future engagements with platform-instrumentation move the score toward “Functional”.

**Distinction from §2.4 worst-signature claim.** The Capability Stack’s worst-case aggregation rule (max lag across measured domains for capability scoring) is *operational* — it produces the strategic-report bars from per-domain measurements. The §2.4 worst-signature pathology claim (continuous  $t_1$  + uncommitted  $t_2$ ) is *categorical* — it names a particular t-state combination as a high-priority pathology. The two are independent contributions; v0.3 is explicit about the distinction (manual M1 + skill M-4 fix).

---

## Appendix E — Construct Frame Library (promoted from Playbook §D.5)

Construct-tagged findings translate SRD-derived capability claims into established theoretical frames the board recognises. Three canonical frames; extensible.

### Required frames for B.5 strategic reports:

Frame	Source	Fires when
<b>Cohen &amp; Levinthal — Absorptive Capacity</b>	Cohen & Levinthal (1990); Zahra & George (2002)	Assimilation lag dominates AND signal class is external knowledge (regulatory, competitive, technological)
<b>David Teece — Dynamic Capabilities</b>	Teece (1997, 2007); Helfat et al. (2007)	Capability profile shows differential strength across sense / seize / reconfigure (not uniformly weak)
<b>Stephan Haeckel — Sense-and-Respond</b>	Haeckel (1999)	Detection latency is the dominant pathology; organisation learns of response only after escalation

**Overlap arbitration.** A finding can plausibly fire multiple frames. The strategic report uses the *best frame per finding* — chosen by which frame’s lens most cleanly explains the binding constraint *for that domain*. The the industrial conglomerate mockup demonstrates: domains 1, 2, and 3 each tagged with a different frame because each domain’s pathology is most legibly explained by a different theoretical lens. There is no “one frame wins”; there is “best frame per finding” with rationale.

**Pattern Library link.** Construct-tagged findings produce Pattern Library captures tagged with both the originating archetype (B.5) and the construct frame. Future engagements retrieve patterns by either dimension.

**Extension candidates for v0.4+:**

Frame	Source	Fires when
Stalk — Time-Based Competition	Stalk (1988); Stalk & Hout (1990)	Total cycle time is the binding constraint and competitive advantage depends on speed; firm-vs-firm competitive frame applies
Eisenhardt — Speed of Strategic Decision-Making	Eisenhardt (1989)	TMT decision-speed is the binding constraint; $t_1 \rightarrow t_2$ interval dominates

These are flagged as candidates; v0.4+ should determine whether they meet the firing-criterion specificity required for the canonical library.

---

## Appendix F — Status thresholds and calibration anchor

The Playbook's `SRDDiagnostic.defaultLongSRDThresholdDays = 30` calibration anchor is the operational baseline for SRD threshold determination across non-categorised signals. This appendix records the calibration anchor's role and provides the v0.3 per-signal-class threshold table (also in §2.3).

**Calibration anchor.** 30 days is the default LONG threshold for signals that do not have an explicit per-signal-class threshold and where no per-engagement override has been set. The choice of 30 days is calibrated against the management-reporting-cadence convention in tier-1-to-2 banking (monthly board reporting; quarterly strategic review with monthly sub-cadence). The anchor is not theoretically derived; it is the empirical baseline against which the diagnostic has been calibrated for the v0.2 / v0.3 case base.

### Per-signal-class default thresholds (v0.3 §2.3, repeated for reference):

Signal class	Default LONG threshold	Provenance
Transaction-time signals	seconds–minutes	industry standard (card auth latency benchmarks)
Incident-response signals	hours	banking ops convention; ITIL incident SLAs
Operational signals	days	management-reporting cadence; weekly ops review benchmark
Capability-decision signals	weeks	management-team decision-cycle benchmark
Strategic-repositioning signals	months	board-cycle benchmark
<b>Non-categorised default</b>	<b>30 days</b>	<b>Playbook calibration anchor</b>

**Per-engagement override mechanism.** The analyst may override the default per-signal-class threshold using either:

- **Peer-organisation benchmark.** Quartile-relative thresholds within a named peer-organisation set.
- **Longitudinal series.** Standard-deviations relative to a longitudinal series within the client.

The override must be named in the engagement record alongside its provenance. v0.4+ should test cross-engagement comparability of overrides as part of the §6.7 inter-analyst reproducibility test.

**v0.4+ calibration target.** The 30-day non-categorised default is calibrated against tier-1-to-2 banking. Application against materially different organisational profiles (per §7.4 generalisability boundary cases) may require a different default or a discipline-extended set of per-signal-class defaults. v0.4 should report whether the default holds across at least one materially-different case.