Bounding the Latent-Failure Exposure of Defeated Monitors: A Quantitative Certification-Maintenance-Requirement Framework for Dormant Protective Functions in Transport-Category Aircraft

Jherrod Thomas

Independent Researcher — The Lion of Functional Safety

jherrodthomas.com · Vol. 1, No. 6 · July 1, 2026

Abstract—Two 2026 United States Federal Aviation Administration airworthiness directives expose a recurring blind spot in civil-aircraft safety assessment: the dormant failure of a protective monitor whose unavailability is assumed to be governed only by hardware wear and a periodic inspection interval. In the Boeing 737NG thrust-reverser case (FAA AD-2025-00364-T, 91 FR 15566), a maintenance-manual rigging procedure can extinguish the reverser lock indication by grinding its proximity-sensor target, disabling the monitor without re-proving that the lock still holds. In the ATR 42 flap-asymmetry case (FAA AD 2026-09-13, 91 FR 27187), the detector's splines wore against the very interconnection shaft it watches, silently ending protection until a mechanic found it on the bench. Both monitors participate in an AND-structured catastrophic top event whose 14 CFR 25.1309 budget of 1×10^{-9} per flight hour leans on the monitor's availability, and in both the exposure-time model behind the governing Certification Maintenance Requirement (CMR) failed to capture a defeat path. We propose a six-step framework that ingests airworthiness-directive evidence, locates the monitor in the ARP4761A Functional Hazard Assessment, models its latent unavailability, and enumerates maintenance and test procedures as common-cause defeat paths carrying a probability term a shorter interval cannot buy down. We derive a maximum CMR interval and a set of independence and non-perturbation requirements on the verification itself, demonstrated on the 737NG reverser with the ATR 42 detector as a cross-check. Illustrative computations show that when the defeat term dominates, no achievable interval restores the 25.1309 budget; only an independent verification of both the monitored function and the indication path closes the gap.

Keywords—Certification Maintenance Requirement, Latent Failure, Dormant Monitor, ARP4761A, 14 CFR 25.1309, Thrust Reverser, Flap Asymmetry, Common Cause Analysis, Exposure Time.

I. Introduction

A protective monitor does no useful work in normal operation. It sits on the function it watches, waits for a failure that may never come, and earns its certification basis exactly once — on the day it converts a would-be catastrophe into an annunciated, dispatch-blocking condition. Because such a monitor is *dormant*, its failures are *latent*: they do not announce themselves, and they accumulate silently until either the monitor is exercised by a periodic test or the monitored function actually fails and finds the monitor already dead. The entire quantitative safety argument for a large class of transport-category systems rests on bounding how long a monitor can sit failed before it is found. That bound is the *exposure time*, and the instrument that enforces it is the Certification Maintenance Requirement (CMR): a periodic, mandatory task derived from the safety assessment specifically to detect a safety-significant latent failure before it can combine with a second event to produce a hazardous or catastrophic outcome [8].

Two airworthiness directives issued in the first half of 2026 show, on two unrelated aircraft, that the exposure-time model behind the CMR can be quietly invalidated — not by an unforeseen physical failure mode, but by a maintenance action. The first is the FAA's 30 March 2026 notice of proposed rulemaking for the Boeing 737-600/-700/-800/-900 family (Project Identifier AD-2025-00364-T, Docket FAA-2026-2726, 91 FR 15566), which addresses a thrust-reverser lock-indication defeat: an aircraft maintenance manual (AMM) procedure intended to rig a newly installed upper locking actuator, when applied to a worn actuator, grinds down the proximity-sensor target until the flight-deck REVERSER caption stops illuminating — extinguishing the monitor without any check that the lock still works or that the indication still functions [1]. The FAA had already chased this exact defeat once, in AD 2019-18-03 [2]; the 2026 action is the third revision of the same fix. The second is the FAA's 14 May 2026 final rule for the ATR 42-200/-300/-320 (AD 2026-09-13, 91 FR 27187), adopting EASA AD 2025-0087, after a flap-asymmetry detector was found with worn splines, "not engaging mechanically" to the flap interconnection shaft it exists to monitor — the monitor and the monitored function had worn against the same piece of metal, and the protection was silently gone until it was discovered during unrelated maintenance [3], [4].

The two events are technologically distant: a hydraulic reverser lock on a narrow-body jet, a mechanical asymmetry brake on a regional turboprop. But they share a precise structural signature. In each, a dormant monitor participates in an AND-decomposed catastrophic or hazardous top event whose 14 CFR 25.1309 probability budget depends on the monitor's availability [9]; in each, the SAE ARP4761A safety assessment bounded the monitor's exposure time under a *wear-only* model and set an inspection interval accordingly [5]; and in each, the actual loss of the monitor arrived through a path — a maintenance procedure, or wear shared with the monitored member — that the Common Cause Analysis never enumerated as a source of monitor loss. The certification budget was intact on paper and defeated in service.

This paper argues that the missing artifact is not a better failure-rate estimate or a shorter interval, but an explicit model of the *defeat path* and a corresponding requirement on the verification itself. The contributions are as follows. First, we propose a six-step framework that ingests airworthiness-directive evidence and produces a derived CMR interval together with independence, defeat-path, and non-perturbation requirements on the verification task. Second, we formalize the argument with four numbered relations: a latent-unavailability model for a periodically inspected monitor, the AND-structured top-event probability, an *effective* unavailability that adds a maintenance-defeat term, and the resulting feasibility inequality for the CMR interval. Third, we demonstrate the framework end-to-end on the 737NG reverser lock indication, using the ATR 42 flap-asymmetry detector as a cross-domain check, and provide a framework figure, a fault tree, and two tables including five derived requirements. We are explicit about the framework's boundary: it converts airworthiness evidence into requirements and interval bounds, but it does not itself supply validated component failure rates, and it cannot certify the human-factors reliability of the maintenance actions it constrains.

The remainder of the paper is organized as follows. Section II reviews prior art in latent-failure and standby-unavailability modeling, CMR derivation, and fault-tree treatment of dependent and common-cause events. Section III presents the six-step framework and its four relations. Section IV applies the framework to the worked 737NG example with the ATR 42 cross-check. Section V discusses limitations, threats to validity, and where the standards lens stops. Section VI concludes.

II. Background and Related Work

A. Latent Failures and Standby Unavailability

The mean unavailability of a periodically tested, otherwise-dormant component is one of the oldest quantitative results in reliability engineering. For a component with constant failure rate λ, tested at interval T with perfect and instantaneous test and repair, the time-average probability that it sits failed is well approximated by λT/2 for $\lambda T \ll 1$ — the linear-in-interval "half the test period" rule that underlies proof-test scheduling across process, nuclear, and aviation domains [12], [18]. Vaurio developed the more general treatment, deriving time-dependent and average unavailability for periodically tested aging components under alternative test-and-repair renewal policies, and cautioned specifically that the asymptotic average must be used with care when the hazard rate is increasing, because aging breaks the constant-λ assumption on which the simple rule rests [12], [13], [14]. This caution is directly relevant to both 2026 cases, whose monitors failed by *wear*, an explicitly age-dependent mechanism.

The functional-safety community reached the same relations from the safety-instrumented-system side. Torres-Echeverría, Martorell, and Thompson modeled and optimized proof-testing policies for multi-channel safety-instrumented functions, evaluating simultaneous, sequential, and staggered test strategies against the time-dependent probability of failure on demand [16]. Rielly examined how to assure a target probability of failure on demand when full proof testing is impractical, quantifying the residual contribution of the untested fraction [17] — the SIS analog of a CMR that nulls an indication without exercising the function beneath it. The common thread is that the interval controls only the failures the test can actually reveal; anything the test does not exercise, or anything that fails the test's own assumptions, escapes the λT/2 accounting.

B. Certification Maintenance Requirements

In civil aviation, the CMR is the formal mechanism that converts an exposure-time assumption into a mandatory, trackable maintenance task. FAA Advisory Circular 25-19A defines a CMR as a required periodic task, established during type certification, to detect a safety-significant latent failure that would, in combination with one or more further failures or events, produce a hazardous or catastrophic failure condition [8]. Crucially, the AC ties the interval to the safety analysis: because exposure time is a direct multiplier in the latent-failure probability, limiting exposure time has a first-order effect on the computed failure probability, and where the detection method is a test, the applicant must show that the test actually detects the latent failures of concern [8]. The safety assessment that generates CMR candidates is conducted per SAE ARP4761A in conjunction with ARP4754A, the accepted means of compliance with 25.1309 [5], [6], [7]. The two 2026 ADs are, in effect, the regulator retrofitting CMRs that the original assessment either omitted or under-specified: the 737NG NPRM writes CMRs 78-CMR-01 through 78-CMR-07 into the maintenance program explicitly [1], and the ATR 42 rule mandates a special detailed inspection where none was scheduled [3].

C. Fault Trees, Dependence, and Common-Cause Failure

The AND-gate that combines a primary failure with a latent monitor failure is the canonical fault-tree pattern for satisfying the 25.1309 "no catastrophic failure condition from a single failure" clause [9], [19]. Classical fault-tree analysis, however, assumes independence among basic events and invariant failure and repair rates — assumptions that fail precisely when a maintenance action or a shared physical member couples the monitor to what it monitors. Andrews and Tolo's Dynamic and Dependent Tree Theory (D2T2) was developed to integrate general dependency and non-constant rates into the fault-and-event-tree framework by coupling binary decision diagrams with Markov and stochastic-Petri-net solvers, and its importance-measure extension quantifies how a dependent basic event drives the top-event probability [15], [25]. Common-cause failure modeling — most simply the beta-factor apportionment used in IEC 61508 — captures the fraction of failures that defeat redundancy through a shared root [20]. The 2026 cases are common-cause events of an unusual kind: the shared root is not an environmental stress but a *procedure* (the AMM grind) or a *shared component* (the worn spline), and neither was enumerated in the Common Cause Analysis as a monitor-

loss path.

D. Maintenance-Induced Failure and Reliability-Centered Maintenance

That maintenance is itself a failure source is not new. Reliability-centered maintenance formalized the "failure-finding task" for hidden functions and warned that protective devices with undetected failures expose the system to multiple failure, but it treats the failure-finding interval as the control and does not model the finding task as a potential defeat path [21]. Recent work on aircraft repairable units with hidden functions has refined the interval using imperfect-restoration (Kijima-type) renewal models and the mean-fractional-dead-time metric [22], acknowledging that a restoration action may leave the item worse than the constant-λ model assumes. Human-factors studies attribute a persistent 15–20% share of aviation occurrences to maintenance error [23]. What the literature has not done, and what the 2026 ADs demand, is to fold the maintenance-procedure-as-defeat-path directly into the top-event quantification and then derive a requirement that constrains the procedure rather than merely the interval. That is the gap this paper addresses.

III. Approach

A. Overview

The framework (Fig. 1) takes as input field evidence — airworthiness directives, service bulletins, and in-service or shop findings — that identifies a dormant protective monitor and, ideally, the mechanism by which it was lost. It produces as output a derived CMR interval and a set of verification requirements. The six stages are: (S1) evidence ingestion; (S2) monitor identification and FHA classification; (S3) latent-exposure modeling; (S4) defeat-path enumeration; (S5) requirement derivation; and (S6) verification and closure. Stages S1–S3 are conventional ARP4761A practice reorganized around the monitor; the contribution concentrates in S4 and S5.

B. S1–S2: Evidence Ingestion and Monitor Classification

Stage S1 extracts, from the airworthiness record, the monitored function, the dormant element that watches it, the observed loss mechanism, and the time-in-service over which the loss went undetected. Stage S2 locates the monitor in the Functional Hazard Assessment and confirms the severity of the top event it guards. For a thrust reverser, 14 CFR 25.933 requires that no single failure or probable combination cause unwanted in-flight deployment, or that the aircraft remain controllable if it occurs [10]; the accepted discharge is to show inadvertent in-flight deployment is *extremely improbable*, routing into the 25.1309 Catastrophic budget of $1\times10-9$ per flight hour with the added constraint that it not result from a single failure [9]. That constraint is exactly what forces the AND architecture: the reverser restraint and its monitor must both fail for the hazard to manifest.

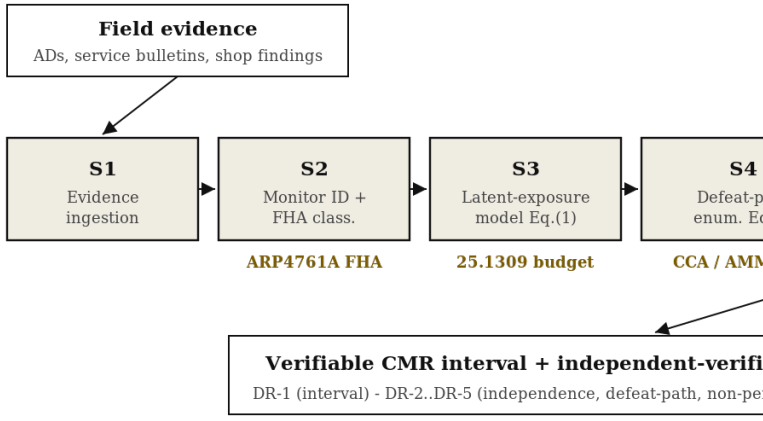

Figure 1: Cross-standards framework. Field evidence (top) flows through six numbered stages (S1–S6), each annotated with its governing clause, into a derived CMR interval and independent-verification requirements (bottom). The novelty is S4, which treats maintenance and test procedures as common-cause defeat paths.

C. S3: Latent-Exposure Model

Let the monitor have latent-failure rate λM and be exercised by a CMR at interval T. Under the standard assumptions of a perfect, non-perturbing test and negligible repair time, the time-average unavailability of the monitor is

$$\bar{Q}_M = \frac{1}{T}\int_0^T \left(1 - e^{-\lambda_M t}\right) dt \approx \frac{\lambda_M T}{2}, \quad \lambda_M T \ll 1 \quad (1)$$

The catastrophic top event requires the coincidence of a demand — loss of the primary restraint, occurring at rate λR per flight hour — with the monitor sitting unavailable. Treating the two as independent, the per-flight-hour top-event probability is the AND product

$$P_{\text{TOP}} \approx \lambda_R \, \bar{Q}_M \quad (2)$$

Equation (2) is the quantitative face of the "no single failure" clause: neither λR alone nor $\bar{Q}$M alone produces the hazard, and the design intent is that their product sit at or below the 25.1309 budget Ptgt = $1\times10-9$ per flight hour.

D. S4: Defeat-Path Enumeration

Stage S4 is the departure from conventional practice. We assert that for a dormant monitor, the set of ways the monitor can be unavailable is not exhausted by wear-driven latent failure. A maintenance or rigging procedure that alters the monitor's sensing target, or a test procedure that perturbs the monitored function, is a *defeat path*: it renders the monitor unavailable through an action the exposure-time model does not see. We represent the aggregate probability that a given monitor

instance has been placed in a defeated state, and not independently re-verified, by a term pdef. The *effective* monitor unavailability becomes the OR-combination of the wear term and the defeat term,

$$\bar{Q}_{M,\text{eff}} \approx \frac{\lambda_M T}{2} + p_{\text{def}} \tag{3}$$

The decisive property of (3) is that pdef does not decrease with T. Reducing the inspection interval buys down only the wear term; the defeat term is set by the maintenance process, not the calendar. Worse, when the CMR test *is* the perturbing action — as in the 737NG history, where the 2019 integrity test could wind residual torque into the reverser flex shafts and cause a failure to deploy on the first post-test command [1] — shortening T can *increase* the effective hazard by increasing the frequency of the perturbation. This is the formal statement of a result that is invisible to a wear-only assessment: exposure-time reduction and defeat-path reduction are different controls acting on different branches of the OR gate beneath the monitor (Fig. 2).

E. S5: Requirement Derivation

Combining (2) and (3) and imposing the budget PTOP $\leq$ Ptgt yields the feasibility condition on the interval:

$$T \leq T_{\text{max}} = \frac{2}{\lambda_M}\left(\frac{P_{\text{tgt}}}{\lambda_R} - p_{\text{def}}\right) \tag{4}$$

Equation (4) carries the paper's central engineering consequence. A real, non-negative CMR interval exists only if

$$p_{\text{def}} < \frac{P_{\text{tgt}}}{\lambda_R} \tag{5}$$

If the defeat probability exceeds the ratio of the target budget to the primary-restraint demand rate, then Tmax is negative: no achievable inspection interval, however short, restores the 25.1309 budget, because the interval cannot reach the term that violates it. The requirement that follows is therefore not "inspect more often" but "close the defeat path." Five derived requirements (Table II) express this: DR-1 sets the interval from (4) using the wear term only, *after* the defeat term has been driven to zero; DR-2 requires the verification to exercise both the monitored function and the indication path independently, so that no procedure can extinguish the annunciation without proving the function; DR-3 requires the Common Cause Analysis to classify any procedure that can alter the monitor's sensing target as a monitor-defeat path with a mandatory post-task functional verification; DR-4 requires the verification action to be shown non-perturbing to the monitored function, or to carry a corrective step; and DR-5 requires that a discovery-only latent finding trigger re-derivation of exposure time using the observed in-service time as an empirical upper bound.

F. S6: Verification and Closure

Stage S6 confirms that the verification specified by DR-2 through DR-4 is implementable and independent, that it does not itself perturb the system, and that the resulting CMRs are written into the operator's maintenance program under the type-certificate holder's control per AC 25-19A [8], with development traced through the MSG-3 scheduled-maintenance logic [21] where applicable. Closure is reached when every OR-branch beneath the monitor in the fault tree (Fig. 2) is either bounded by an interval or eliminated by an independence requirement.

IV. Worked Example and Cross-Check

A. Boeing 737NG Thrust-Reverser Lock Indication

We instantiate the framework on the 737NG case. In S1, the evidence is the 2026 NPRM and its 2019 predecessor: the monitored function is the upper-locking-hydraulic-actuator restraint that keeps the reverser stowed; the dormant monitor is the proximity sensor and its flight-deck REVERSER indication; the loss mechanism is the AMM target-grinding procedure applied to a worn actuator; and the FAA's own words establish that "a locking mechanism failure … could remain undetected for thousands of flights" [1], [2]. In S2, the top event — uncommanded in-flight deployment — is Catastrophic under 25.933/25.1309, with the "not from a single failure" constraint forcing the AND structure [9], [10]. The severity is not in dispute; it is the failure condition behind the 1991 Lauda Air Boeing 767 in-flight reverser deployment, lost with all 223 aboard [24].

In S3 and S4 we quantify. The values below are illustrative order-of-magnitude figures chosen to exercise the relations, *not* certification data; a real assessment would substitute validated rates. Take the budget Ptgt = $1\times10{-}9$/fh, a primary-restraint demand rate λR = $1\times10{-}6$/fh, and a monitor wear-failure rate λM = $1\times10{-}6$/fh. The budget on monitor unavailability implied by (2) is $\bar{Q}$M $\leq$ Ptgt/λR = $1\times10{-}3$. Under the wear-only model (1), the 2019 rule's 750-flight-hour interval gives $\bar{Q}$M = $(1\times10{-}6)(750)/2 = 3.75\times10{-}4$, so PTOP = $3.75\times10{-}10$/fh — comfortably within budget, and by (4) the wear-only Tmax = $2(1\times10{-}3)/(1\times10{-}6) = 2000$ fh. On the wear model alone, the interval is not the problem.

Now introduce the defeat path. Suppose a defeat probability pdef = $1\times10{-}3$ — that roughly one monitor instance in a thousand has been subjected to the target-grinding procedure and not independently re-verified. By (3), $\bar{Q}$M,eff = $3.75\times10{-}4 + 1\times10{-}3 = 1.375\times10{-}3$, so PTOP = $1.375\times10{-}9$/fh — over budget, dominated entirely by the defeat term. By (5), feasibility requires pdef < $1\times10{-}3$; at pdef = $1\times10{-}3$ exactly, (4) gives Tmax = 0. Push the defeat probability to a still-plausible pdef = $1\times10{-}2$ and PTOP = $(1\times10{-}6)(3.75\times10{-}4 + 1\times10{-}2) \approx 1.04\times10{-}8$/fh — an order of magnitude over budget, and Tmax is negative. No inspection interval recovers the budget. This is the arithmetic of the recurrence: the 2019 AD shortened and formalized the interval, and the hazard came back, because the interval was

never the binding control. The binding control is DR-2 and DR-4 — an independent verification of both the lock function and the indication path, and a demonstration that the verification does not itself perturb the flex-shaft torque state.

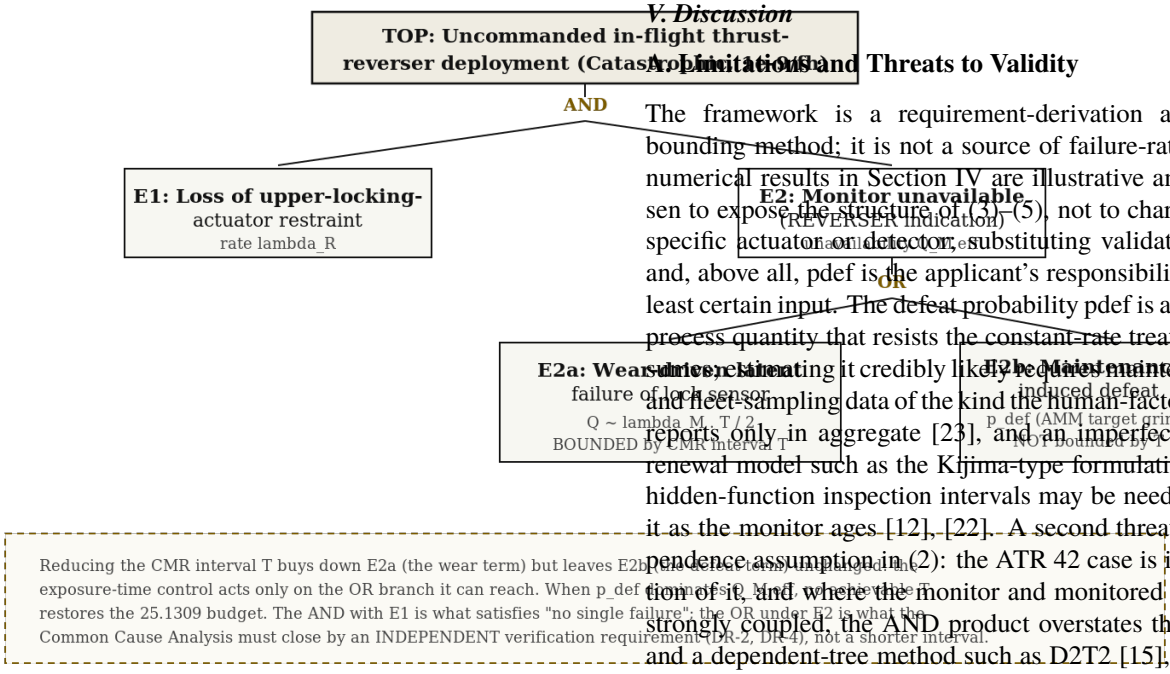

Figure 2: Fault tree for the top event "Uncommanded in-flight thrust-reverser deployment." The AND gate over E1 (restraint loss) and E2 (monitor unavailable) enforces the 25.1309 "no single failure" clause. Beneath E2, an OR gate separates the wear-driven latent failure E2a — bounded by the CMR interval T per (1) — from the maintenance-induced defeat E2b, which the interval cannot reach and which the Common Cause Analysis must close by independence.

B. Cross-Check: ATR 42 Flap-Asymmetry Detector

The ATR 42 case exercises the same relations with a different defeat mechanism, testing whether the framework generalizes. Here the monitored function is flap synchronization on the interconnection shaft required by 14 CFR 25.701 [11]; the dormant monitor is the flap-asymmetry detector that commands the wing-flap brake; and the loss mechanism is spline wear that disengages the detector from the very shaft it reads [3], [4]. The top event — an unarrested flap asymmetry producing an untrimmable rolling moment — is a loss-of-control condition; per the AND structure, it requires an asymmetry demand *and* a dead detector. There is no AMM grind here, so pdef in the sense of a procedural defeat is near zero; instead the failure is a *shared-member common cause*: the detector and the monitored shaft wear together, coupling λM to the demand and violating the independence that (2) assumes. The framework flags this at S4 as a common-cause defeat path of the shared-component type, and the requirement that follows is DR-3 (classify the shared spline as a coupling and inspect it as such) and DR-5 (the detector was found *on the bench*, so its observed time-in-service is an empirical upper bound on exposure that must feed the interval re-derivation, since no CMR had scheduled the finding at all). Table I contrasts the two cases; the structural lesson is identical across a hydraulic jet reverser and a mechanical turboprop flap system.

V. Discussion

A. Limitations and Threats to Validity

The framework is a requirement-derivation and interval-bounding method; it is not a source of failure-rate data. The numerical results in Section IV are illustrative and were chosen to expose the structure of (3)–(5), not to characterize any specific actuator or detector, substituting validated λR, λM, and, above all, pdef is the applicant's responsibility and is the least certain input. The defeat probability pdef is a human-and-process quantity that resists the constant-rate treatment (1) assumes, meaning it credibly likely requires maintenance-error and fleet-sampling data of the kind the human-factors literature reports only in aggregate [23], and an imperfect-restoration renewal model such as the Kijima-type formulations used for hidden-function inspection intervals may be needed to bound it as the monitor ages [12], [22]. A second threat is the independence assumption in (2): the ATR 42 case is itself a violation of it, and where the monitor and monitored function are strongly coupled, the AND product overstates the protection and a dependent-tree method such as D2T2 [15], [25] should replace the product form. Third, the framework inherits the FHA's severity classification; if the top event is misclassified, the budget and hence Tmax are wrong at the root.

B. Where the Standards Lens Stops

ARP4761A, AC 25-19A, and 25.1309 together give a rigorous account of latent failure as a *hardware* phenomenon bounded by an interval, and the CMR is a genuinely effective instrument for the wear term. What the standards lens does not natively model is the maintenance procedure as an element of the safety architecture. The AMM step that grinds the sensor target is, formally, a control-flow edge into the monitor's availability that no failure rate captures; the standards treat it as outside the system boundary even though it can zero the system's most important protective claim. This paper's contribution is to pull that edge inside the boundary as pdef and to make the verification's independence a derived requirement rather than an assumed property. But the lens still stops short of certifying the reliability of the humans and the documents that execute DR-2 through DR-4: the framework can require an independent, non-perturbing verification, but it cannot guarantee the maintainer performs it, and that residual is where the recurrence risk ultimately lives.

VI. Conclusion and Future Work

Two unrelated 2026 airworthiness directives — a Boeing 737NG thrust-reverser lock indication and an ATR 42 flap-asymmetry detector — share a precise and instructive signature: a dormant monitor, sitting in an AND-decomposed catastrophic or loss-of-control top event, whose 25.1309 budget was

Table 1: Two 2026 airworthiness directives sharing the defeated-dormant-monitor signature.

Attribute	Boeing 737NG [1],[2]	ATR 42 [3],[4]
Monitored function	Upper-locking-actuator reverser restraint	Flap interconnection-shaft synchronization
Dormant monitor	Proximity sensor + REVERSER indication	Flap-asymmetry detector + brake command
Top-event severity / clause	Catastrophic; 25.933 / 25.1309	Loss of control; 25.671 / 25.701
Loss mechanism	AMM target-grind defeat (procedural)	Spline wear shared with monitored shaft
Defeat-path type (S4)	Procedure-induced, $p_{\text{def}} > 0$	Shared-member common cause
How discovered	Repeat in-service reports; prior AD	Found during unrelated maintenance
Latent exposure	Thousands of flights (FAA)	Until bench discovery (unbounded)
Regulatory response	CMRs 78-CMR-01..07; measured tolerances	One-time special detailed inspection
Binding control (this work)	DR-2, DR-4 (independent, non-perturbing)	DR-3, DR-5 (coupling + empirical bound)

Table 2: Derived requirements. DR-1 sets the interval only after the defeat term is closed; DR-2 to DR-5 close the OR-branch the interval cannot reach.

ID	Requirement	Term controlled	Clause / case
DR-1	Set CMR interval $T \leq T_{\text{max}}$ from (4) using the wear term, validated λ_M	$\lambda_M T/2$	AC 25-19A; both
DR-2	Verification shall independently exercise the monitored function and the indication path	p_{def}	25.1309; 737NG
DR-3	CCA shall classify any target-altering or shared-member path as a monitor-defeat path with mandatory post-task check	p_{def}	ARP4761A; both
DR-4	Verification shall be shown non-perturbing to the monitored function, or carry a corrective step	p_{def}	25.933; 737NG
DR-5	Discovery-only latent findings shall re-derive exposure time using observed in-service time as an empirical bound	T	AC 25-19A; ATR 42

intact under a wear-only exposure model and defeated in service by a path the Common Cause Analysis never enumerated. We have shown that the governing Certification Maintenance Requirement interval is the wrong lever when the defeat term dominates: the exposure-time control acts only on the wear branch of the OR gate beneath the monitor, and when the procedural or shared-member defeat probability exceeds the ratio of the target budget to the primary-demand rate, no achievable interval restores compliance. The remedy is an explicit defeat-path term in the top-event quantification and a corresponding independence-and-non-perturbation requirement on the verification itself, expressed here as five derived requirements and demonstrated on the 737NG reverser with the ATR 42 detector as a cross-domain check.

Future work should pursue three directions. First, empirical estimation of pdef from fleet maintenance records and shop-finding data, so that the feasibility inequality (5) can be evaluated with defensible numbers rather than illustrative ones. Second, replacement of the independent AND product (2) with a dependent-tree formulation for cases, like the ATR 42, where the monitor and monitored function share a physical member. Third, extension of the defeat-path construct beyond aviation — to the proof-test-defeat and bypass paths of IEC 61508 safety-instrumented functions and to the equivalent maintenance-induced common-cause paths in other domains — since the structural result, that a shorter test interval cannot buy down a defeat the test does not exercise, is not specific to aircraft.

References

1. Federal Aviation Administration, "Airworthiness Directives; The Boeing Company Airplanes," Notice of Proposed Rulemaking, Project Identifier AD-2025-00364-T, Docket No. FAA-2026-2726, *Federal Register*, vol. 91, p. 15566, Mar. 30, 2026. [Online]. Available: https://www.federalregister.gov/documents/2026/03/30/2026-06067/airworthiness-directives-the-boeing-company-airplanes
2. Federal Aviation Administration, "Airworthiness Directives; The Boeing Company Airplanes," Final Rule, AD 2019-18-03, *Federal Register*, vol. 84, p. 49005, Sep. 18, 2019. [Online]. Available: https://www.federalregister.gov/documents/2019/09/18/2019-20184/airworthiness-directives-the-boeing-company-airplanes
3. Federal Aviation Administration, "Airworthiness Directives; ATR-GIE Avions de Transport Régional Airplanes," Final Rule, AD 2026-09-13, Amdt. 39-23335, *Federal Register*, vol. 91, p. 27187, May 14, 2026. [Online]. Available: https://www.federalregister.gov/documents/2026/05/14/2026-09658/airworthiness-directives-atr-gie-avions-de-transport-rgional-airplanes
4. European Union Aviation Safety Agency, "Airworthiness Directive AD 2025-0087: ATR 42 — Flight Controls — Flap Asymmetry Detection," Apr. 16, 2025. [Online]. Available: https://ad.easa.europa.eu/ad/2025-0087
5. SAE International, *ARP4761A: Guidelines for Conduct-*

ing the Safety Assessment Process on Civil Aircraft, Systems, and Equipment, Warrendale, PA, USA, Dec. 2023, doi: 10.4271/ARP4761A.
6. SAE International, *ARP4754A: Guidelines for Development of Civil Aircraft and Systems*, Warrendale, PA, USA, Dec. 2010, doi: 10.4271/ARP4754A.
7. Federal Aviation Administration, *Advisory Circular 25.1309-1B: System Design and Analysis*, Washington, DC, USA. [Online]. Available: https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_25.1309-1B.pdf
8. Federal Aviation Administration, *Advisory Circular 25-19A: Certification Maintenance Requirements*, Washington, DC, USA, Oct. 3, 2011. [Online]. Available: https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_25-19A.pdf
9. *Electronic Code of Federal Regulations, Title 14 §25.1309 — Equipment, systems, and installations*. [Online]. Available: https://www.ecfr.gov/current/title-14/part-25/section-25.1309
10. *Electronic Code of Federal Regulations, Title 14 §25.933 — Reversing systems*. [Online]. Available: https://www.ecfr.gov/current/title-14/part-25/section-25.933
11. *Electronic Code of Federal Regulations, Title 14 §25.701 — Flap and slat interconnection*. [Online]. Available: https://www.ecfr.gov/current/title-14/part-25/section-25.701
12. J. K. Vaurio, "On time-dependent availability and maintenance optimization of standby units under various maintenance policies," *Rel. Eng. Syst. Saf.*, vol. 56, no. 1, pp. 79–89, 1997, doi: 10.1016/S0951-8320(96)00132-9.
13. J. K. Vaurio, "The modelling of degraded and critical failures for components with dormant failures," *Rel. Eng. Syst. Saf.*, vol. 50, no. 3, pp. 315–322, 1995, doi: 10.1016/0951-8320(95)00116-6.
14. J. K. Vaurio, "Time-dependent unavailability analysis of standby components incorporating arbitrary failure distributions and three inspection/maintenance policies," *Rel. Eng. Syst. Saf.*, 1993, doi: 10.1016/0951-8320(93)90146-P.
15. J. D. Andrews and S. Tolo, "Dynamic and dependent tree theory (D2T2): A framework for the analysis of fault trees with dependent basic events," *Rel. Eng. Syst. Saf.*, vol. 230, art. 108959, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0951832022005749
16. A. C. Torres-Echeverría, S. Martorell, and H. A. Thompson, "Modelling and optimization of proof testing policies for safety instrumented systems," *Rel. Eng. Syst. Saf.*, vol. 94, no. 4, pp. 838–854, 2009. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0951832008002287
17. J. Rielly, "Assuring the probability of failure on demand of a safety instrumented system without full proof testing," *Measurement and Control*, vol. 49, no. 8, pp. 251–256, 2016, doi: 10.1177/0020294016663975.
18. M. Kijima, "Some results for repairable systems with general repair," *J. Appl. Probab.*, vol. 26, no. 1, pp. 89–102, 1989, doi: 10.2307/3214319.
19. W. E. Vesely, F. F. Goldberg, N. H. Roberts, and D. F. Haasl, *Fault Tree Handbook*, NUREG-0492, U.S. Nuclear Regulatory Commission, Washington, DC, USA, 1981. [Online]. Available: https://www.nrc.gov/docs/ML1007/ML100780465.pdf
20. International Electrotechnical Commission, *IEC 61508: Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems*, Parts 1–7, Geneva, Switzerland, 2010.
21. Air Transport Association of America, *MSG-3: Operator/Manufacturer Scheduled Maintenance Development*, Washington, DC, USA.
22. S. Tolo and J. D. Andrews, "Calculation of the system unavailability measures of component importance using the D2T2 methodology of fault tree analysis," *Mathematics*, vol. 12, no. 2, art. 292, 2024, doi: 10.3390/math12020292.
23. T. T. Takahashi et al., "Improved field performance through regulatory changes to thrust reverser certification," in *Proc. AIAA SciTech Forum*, AIAA 2016-1280, 2016, doi: 10.2514/6.2016-1280.
24. Federal Aviation Administration, *Aviation Maintenance Technician Handbook — General*, FAA-H-8083-30, Ch. 14: Human Factors, Washington, DC, USA. [Online]. Available: https://www.faa.gov/regulations_policies/handbooks_manuals/aircraf
25. Aviation Safety Network, "Accident record: Lauda Air Flight 004, Boeing 767-3Z9ER, 26 May 1991." [Online]. Available: https://aviation-safety.net/database/record.php?id=19910526-0