When the Fix Stops at the Program Boundary: An Inheritance-Graph Framework for Bounding Sibling Exposure After a Confirmed Safety Anomaly in ISO 26262 Field Monitoring

Jherrod Thomas
Independent Researcher — The Lion of Functional Safety

JHERRODTHOMAS.COM · OCTOBER 2026

***Abstract*—When a safety anomaly is confirmed on one vehicle program, ISO 26262 requires the manufacturer to monitor and handle it, but no clause says how far the evaluation must reach across sibling programs, in what order, or how quickly. We propose an inheritance-graph framework for extent of condition. Programs are nodes; typed edges record shared specification, software logic, supplier, validation method, manufacturing line and service part; each edge carries a transfer probability that screening evidence updates. Six relations give the probability that a sibling inherits the anomaly, its posterior after screening, the vehicle-days of exposure accumulated before the sibling is decided, the break-even event rate at which one event is expected within that latency, the maximum decision latency for a stated event budget, and an evaluation order. Applied to the public NHTSA record for the Hyundai Palisade (26V160) and Ioniq 9 (26V588) power-seat recalls, the framework yields about 8.9×10^5 vehicle-days of sibling exposure over 169 days, a break-even rate near 5×10^{-4} per vehicle-year, and a 75 percent reduction had the sibling been decided within 45 days. Applied to the Ram ProMaster 26V621 scope extension, it yields about 2.6×10^7 vehicle-days outside the investigated population. Six requirements are derived.**

***Index Terms*—ISO 26262, Extent of Condition, Field Monitoring, Safety Anomaly, Dependent Failure Analysis, Recall, Inheritance Graph, Platform Sharing**

I. INTRODUCTION

A confirmed safety anomaly is a statement about one vehicle program, but the engineering that produced it was almost never confined to that program. Specifications are inherited, control logic is carried forward and re-calibrated, suppliers serve several platforms, and test methods that missed a defect in one vehicle are the same methods that approved its siblings. ISO 26262 recognizes that field experience must feed back into development: its safety-management part requires field monitoring and the handling of safety anomalies after release for production, and its supporting parts require impact analysis for changes and a dependent-failure analysis for elements that are supposed to be independent [1]. What the standard does not state is how far a confirmed anomaly must be pursued across programs, in what order, and how quickly. That quantity, the *extent of condition*, is left to the manufacturer's judgment and, in practice, to the speed of its investigation process.

The public record shows what that judgment costs. On March 17, 2026, Hyundai reported NHTSA campaign 26V160, covering 61,093 model-year 2026 Palisade vehicles whose rear power seat could entrap an occupant, with an interim software remedy and a permanent remedy still in development [4]. On September 10, 2026, 177 days later, Hyundai reported campaign 26V588 for 5,360 model-year 2026 Ioniq 9 vehicles, with a consequence statement identical to the Palisade's [5]. Production of the affected Ioniq 9 units continued until May 4, 2026, 48 days after the Palisade report [6]. In a second, unrelated case, FCA US recalled 265,512 Ram ProMaster vans for corroded steering harnesses on September 29, 2026 [7], after NHTSA had opened a preliminary evaluation of 136,666 vehicles on March 2, 2026 [8]. The recall population was 94 percent larger than the population first investigated. In both cases the defect was found once and then found again, in a larger or different population, after a delay measured in months.

This paper addresses one research question. *Given a confirmed safety anomaly in one vehicle program, how should the set of sibling programs to be evaluated, the order of their evaluation, and the maximum time to each evaluation decision be derived, so that the exposure accumulated while a sibling's status is unknown is bounded and auditable?* The contributions are:

- An **inheritance graph** whose nodes are programs, variants and service parts and whose typed edges record shared specification, software logic, hardware or supplier, validation method, manufacturing line and service part, each edge carrying a transfer probability with an explicit prior.
- **Six numbered relations**: sibling transfer probability, a screening update, accumulated sibling exposure with a production-continues term, an expected-event and break-even-rate pair, a maximum evaluation latency for a stated event budget, and an evaluation-ordering rule.
- A **mapping to ISO 26262** field-monitoring, change-management, proven-in-use and dependent-failure clauses, and six derived requirements for an extent-of-condition register.
- A **worked application** to the Hyundai Palisade and Ioniq 9 seat recalls, giving about 8.9×10^5 vehicle-days of sibling

exposure over 169 days, and a second application to the Ram ProMaster scope extension, giving about 2.6×10^7 vehicle-days outside the investigated population.

Section II reviews background and the gap. Section III develops the framework. Section IV applies it. Section V discusses limitations and where the standards lens stops. Section VI concludes.

II. BACKGROUND AND RELATED WORK

A. Standards and Regulation

ISO 26262:2018 assigns post-release monitoring to its safety-management part, whose clause 7 covers production, operation, service and decommissioning, including field monitoring and the handling of safety anomalies [1]. Part 8 clause 8 requires an impact analysis for every change, part 8 clause 14 bases a proven-in-use argument on field data and service-period evidence, and part 9 clause 7 requires a dependent-failure analysis to demonstrate that elements claimed to be independent do not share a cause [1]. Each clause assumes a defined scope: one item, one change, one architecture. None says how a finding on one item defines the scope of the next item's review.

The United States regulator applies a clock to a different event. A manufacturer must submit a defect report within five working days after it determines that a defect is safety-related [2]. The clock starts at the manufacturer's own determination, which is the end of the process this paper models, not its beginning. For the function at issue in the first case there is, to our knowledge, no federal standard: FMVSS 118 prescribes requirements for power-operated windows, partitions and roof panels but, to our knowledge, does not address power seats [3]. Where no external requirement exists, the only safety argument is the manufacturer's own, and its reach is whatever the safety plan makes it.

B. Recalls as an Object of Study

Rupp and Taylor studied the initiation of automobile recalls [10], and Bates *et al.* described trends, patterns and emerging issues in motor vehicle recalls [11]. Chi, Sigmund and Astardi built a root-cause and FMEA classification scheme for passenger-vehicle recalls [12]. None of these treats the sibling question directly. Kang, Hong, Huh and Kang analyze failure risk from the perspective of platform sharing in *IEEE Transactions on Engineering Management* [13], which is the closest economic precedent: sharing a platform couples the failure risk of the products built on it. Our framework makes that coupling operational at the level of a safety register.

C. Defect Propagation in Software and Product Lines

Software engineering has measured how defects spread. Fenton and Ohlsson analyzed the distribution of faults and failures across the modules of a complex software system [14]. Zimmermann and Nagappan showed that network measures on dependency graphs help predict defective modules [15]. Juergens *et al.* asked whether code clones matter and examined faults in inconsistently changed clones [16], which is the code-level form of a defect repaired in one copy and not in another. Abal, Brabrand and Wasowski catalogued variability bugs that appear only under particular configurations [17], and Thüm *et al.* surveyed analysis strategies for software product lines, which seek to analyze a family without examining every product separately [18]. A vehicle program family is a product line in which the shared assets include requirements, logic and suppliers. The extent-of-condition question is the product-line question applied after a field failure.

D. Common Cause and Dependent Failures

Common-cause failure analysis was developed for redundant systems in the nuclear industry, where procedures were documented in NUREG/CR-4780 [23]. Troffaes, Walter and Kelly treat the epistemic uncertainty of common-cause parameters with a robust Bayesian model [22], and Picoco and Rychkov model common-cause failure in dynamic probabilistic risk assessment [25]. These methods quantify dependence among redundant channels at one moment in one system. The dependence studied here runs between systems and across calendar time: the "channels" are separate vehicle programs and the "mission time" is the interval during which a sibling's status is unknown.

E. Propagation on Networks

The independent-cascade model of Kempe, Kleinberg and Tardos activates each edge of a graph independently with a stated probability and asks which nodes are reached [20]. Pastor-Satorras *et al.* review epidemic processes on complex networks, in which the same edge-transmission structure determines how a contagion spreads [21]. We borrow the edge-transmission model and discard the rest: the number of sibling programs is small, the graph is known to the manufacturer, and the interesting quantity is not the final reach but the time-integrated exposure before the reach is known.

F. Assurance Reuse and Compliance Evidence

Kokaly *et al.* manage the reuse of assurance cases when a system evolves, tracing which claims survive a change [19]. Deka and Sarkar catalogue thirteen public automotive compliance-documentation failures and conclude that none was caught by the supply chain's routine verification layer [26]. Ihshaish *et al.* show that the choice among matched records describing the same operational case, including NHTSA recalls, can change the result of text-classification studies more than the choice of model [27]. The last two matter here as a caution: the public recall record is a manufacturer-chosen narrative, and any quantity derived from it inherits that selection.

G. The Gap

The literature gives us the economics of recalls, the propagation of defects through software and product lines, the quantification of common cause within one system, and the reuse of assurance arguments under change. It does not give a field-monitoring artifact that, on the day an anomaly is confirmed, lists the siblings, prices the exposure of leaving each one unevaluated, and sets a deadline. That artifact is what we propose.

III. APPROACH: THE EXTENT-OF-CONDITION FRAMEWORK

A. Anchor, Descriptor and Inheritance Graph

The framework starts from an *anchor*: a program a with a confirmed safety anomaly at date t_a, taken as the date of the manufacturer's defect report. The anomaly is reduced to a *defect-class descriptor*, a triple of the function affected, the failure mechanism, and the hazard. For the first case below the triple is power-seat occupant detection, failure to respond to contact during automatic motion, and occupant entrapment.

The *inheritance graph* is a directed graph $G = (V, E)$ with V the set of programs, variants and service parts, and an edge e from a to j for each artifact class shared between them. Table I lists six classes. Each edge has a transfer probability q_e, the probability that the defect class exists in the sibling given that the artifact is shared, expressed as a Beta(α, β) prior so that evidence can update it. Sharing that cannot be determined is treated as shared until disproved ($q = 1$). Fig. 1 shows the graph for the first case.

Table 1. *Edge classes of the inheritance graph with illustrative Beta priors. The priors are elicited starting points to be replaced by the organization's own calibration; they are not empirical estimates.*

Class	Shared artifact	Prior (α, β)	Mean q	Evidence that raises q
S	Specification or design intent	(2, 2)	0.50	Identical defect or consequence text
L	Software logic or calibration	(3, 2)	0.60	Common code lineage, same function name
H	Hardware, part number or supplier	(7, 3)	0.70	Same part family, same supplier site
V	Validation method or test gap	(2, 3)	0.40	Same release gate, same bench method
M	Manufacturing line or process	(3, 7)	0.30	Same plant, same fixture, same shift pattern
R	Service part or replacement module	(5, 5)	0.50	Same part number across variants

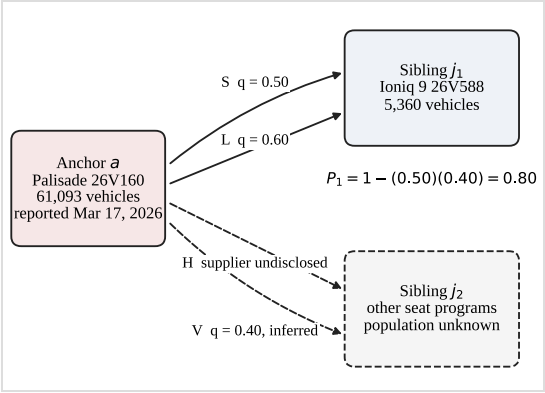


Fig. 1. *Inheritance graph for the first case. Solid edges are supported by the filings: the defect text is identical [4], [5] and the Ioniq 9 filing names the seat logic [6]. Dashed edges are not established by the public record: the supplier is not disclosed, and a shared validation gap is an inference. The sibling probability follows from relation (1).*

B. Relations

Transfer probability. Let Π_{aj} be the set of paths from the anchor to sibling j, and let each edge activate independently with probability q_e. The probability that the defect class reaches j is bounded above by

$$P_j = 1 - \prod_{\pi \in \Pi_{aj}} \left(1 - \prod_{e \in \pi} q_e \right) \tag{1}$$

The expression is exact when the paths are edge-disjoint, as when two sibling edges are parallel, and an upper bound otherwise, because events of the form "this path is live" are positively correlated when paths share edges, a consequence of the Harris–FKG inequality for independent edge activations. An upper bound is the conservative direction for a register.

Screening update. A *screening* applies n test conditions on the sibling, chosen to be diagnostic of the edge class, of which x reproduce the anomaly. The edge's Beta prior updates to

$$q_e' = (\alpha_e + x) / (\alpha_e + \beta_e + n) \tag{2}$$

A single reproduction ($x \geq 1$) converts the sibling from inferred to confirmed and ends the use of (1) for that node. A clean screening ($x = 0$) lowers q_e by an amount that depends on n, so a screening of one condition buys little and a screening of ten buys a lot. For the S edge of Table I, a clean screening of $n = 5$ conditions lowers the mean from 0.50 to 0.22 and $n = 10$ lowers it to 0.14.

Accumulated exposure. Let sibling j have N_j^0 units in service at t_a, a build rate λ_j in units per day, and production that continues until day τ_s after the anchor, whether by the end of the model run or a build hold. Let τ be the latency from the anchor to the sibling decision. The vehicle-days of exposure accumulated before the decision are

$$X_j(\tau) = N_j^0 \, \tau + \lambda_j \, s \, \tau - \tfrac{1}{2} \lambda_j \, s^2, \quad s = \min(\tau, \tau_s) \tag{3}$$

The first term is the standing fleet, the remaining terms are the fleet added by production while the decision is pending. A

manufacturer can reduce the first by deciding faster and the second by deciding faster or by holding the build.

Expected events and break-even rate. Let ρ be the rate of hazardous events attributable to the defect class per exposed vehicle-day, estimated from the anchor's field record or from the HARA exposure and controllability assumptions. The expected number of events in the sibling during the latency, and the rate at which that number reaches one, are

$$E_j = P_j \, \rho \, X_j(\tau), \quad \rho^*_j = 1 / (P_j \, X_j(\tau)) \tag{4}$$

The break-even rate ρ^* is useful because ρ is rarely known with confidence. It turns the question "what is the event rate?" into the question "is the event rate plausibly below ρ^*?", which a safety assessor can answer from the HARA.

Maximum latency. For an event budget ε (default 1), the largest latency that keeps $E_j \le \varepsilon$ solves $P_j \, \rho \, X_j(\tau^*) = \varepsilon$. Writing $c = \varepsilon / (P_j \, \rho)$, the closed forms are

$$\tau^* = [-N^0 + \sqrt{((N^0)^2 + 2\lambda c)}] / \lambda \text{ if } c < X(\tau_s); \text{ otherwise } \tau^* = (c + \tfrac{1}{2}\lambda\tau_s^2) / (N^0 + \lambda\tau_s) \tag{5}$$

Evaluation order. When several siblings await evaluation and evaluation capacity is limited, the cost of leaving sibling j pending accrues at the rate $w_j = P_j \, \rho_j \, N_j(t)$, the derivative of (4) with respect to time. If evaluation j takes d_j days and w_j is approximately constant, minimizing the total accrued cost is a weighted-completion-time scheduling problem whose optimal sequence, due to Smith, is

$$\text{evaluate siblings in descending order of } w_j / d_j \tag{6}$$

with w_j / d_j the exposure-weighted urgency per day of evaluation effort [24]. The rule favors large, probable, quickly evaluated siblings first and defers small, improbable, slow ones.

C. Mapping to ISO 26262 and Derived Requirements

Table II places each framework step under the clause that already gives the manufacturer the authority to perform it. Table III lists the six requirements the framework implies. The requirements are drafts: numeric defaults are illustrative and need justification against the item's HARA.

Table 2. *Framework steps and the ISO 26262:2018 clauses under which each is performed [1].*

Step	Output	Clause
Anchor and descriptor	Defect-class triple, date t_a	Part 2, cl. 7 (field monitoring, anomaly handling)
Inheritance graph	Typed edges and priors	Part 8, cl. 7 (configuration) and cl. 14 (proven in use)
Transfer and screening	P_j, updated q_e	Part 9, cl. 7 (dependent failures); Part 8, cl. 8 (impact analysis)
Exposure and latency	X_j, ρ^*, τ^*	Part 2, cl. 7; defect-report clock of [2]
Sibling re-evaluation	Revised item definition and HARA	Part 3, cl. 5 and 6

Table 3. *Requirements derived from the framework. Numeric defaults are illustrative.*

ID	Requirement	Relation
EOC-R1	Within 10 working days of a confirmed S2 or S3 anomaly, the safety manager shall issue a register listing every program, variant and service part linked to the anchor by at least one typed edge, with owner, population, build window and build rate.	(1), (3)
EOC-R2	Every edge shall be typed by artifact class and carry a prior and its stated basis. Sharing that cannot be determined shall be treated as shared (q = 1) until disproved.	(1)
EOC-R3	Each sibling with $P_j \ge 0.5$ shall be screened with a test derived from the anchor's failure mechanism. The test, n and x shall be recorded and the edge updated.	(2)
EOC-R4	Each sibling shall receive a documented determination (confirmed, cleared with screening evidence, or inapplicable with rationale) within τ^* from (5). Where ρ is unknown, the default ceiling is 45 days.	(4), (5)
EOC-R5	While $P_j \ge 0.5$ and production continues, interim containment, including a build hold, shall be assessed and the decision recorded with its date.	(3)
EOC-R6	Siblings shall be evaluated in the order of (6). A proven-in-use claim shall be suspended for any element with an open register entry, and the register shall be closed only when every node has a determination.	(6)

IV. WORKED EXAMPLES

A. Hyundai Palisade to Ioniq 9

Public record. The anchor is campaign 26V160, reported March 17, 2026 for 61,093 Palisade and Palisade Hybrid vehicles, with the consequence statement that a person, especially a child, may become trapped by a rear powered seat [4]. The sibling is campaign 26V588, reported September 10, 2026 for 5,360 Ioniq 9 vehicles built from March 17, 2025 through May 4, 2026, with

a defect cause stated as seat logic that may not detect contact with an occupant or object [5], [6]. Hyundai's chronology records that its investigation began on March 10, 2026 with benchmarking, that fixture testing on August 4 to 6 measured forces exceeding limits for second-row systems, and that the recall decision came on September 2 [6]. By the decision date Hyundai had five consumer complaints, two reporting minor injuries [6].

Graph and probability. The defect text is identical in the two filings, supporting an S edge, and the Ioniq 9 filing names the seat logic, supporting an L edge [5], [6]. The supplier is not named in the filings, so the H edge is excluded rather than assumed. With the Table I means, relation (1) gives $P_j = 1 - (1 - 0.50)(1 - 0.60) = 0.80$ on the day the anchor was reported. The positive fixture test of August is a screening with $x \geq 1$, which converts the sibling to confirmed by relation (2).

Exposure. The sibling decision came 169 days after the anchor, and the Part 573 report 177 days after it. The build window of 413 days (March 17, 2025 to May 4, 2026) contains 5,360 units, which at a uniform build rate gives $\lambda = 13.0$ units per day. The 48 days from the anchor to the end of the build window then produced about 623 units, and about 4,737 units were in service at t_a. The uniform build rate is our assumption and the filings give no daily figures. Relation (3) with $\tau = 169$ and $\tau_s = 48$ gives $X_j = 8.91 \times 10^5$ vehicle-days, of which about 8.01×10^5 is the standing fleet and the remainder is the fleet added by production. Table IV collects the quantities.

Table 4. *Worked-example quantities for the Palisade to Ioniq 9 case. Published values are from the cited filings; derived values follow from relations (1) to (4) under the stated assumption of a uniform build rate.*

Quantity	Value	Basis
Anchor report date t_a	Mar 17, 2026	[4]
Sibling decision, Part 573 report	Sep 2, Sep 10, 2026	[6], [5]
Latency τ to decision	169 days	Derived
Sibling population N	5,360	[5]
Build window, days	413 (Mar 17, 2025 to May 4, 2026)	[6]
Build rate λ, in service at t_a	13.0 per day, about 4,737	Assumed uniform
Production after anchor, τ_s	48 days, about 623 units	Derived
Transfer probability P_j	0.80	(1), Table I
Exposure $X_j(169)$	8.91×10^5 vehicle-days	(3)
Break-even rate ρ^*	1.40×10^{-6} per vehicle-day (5.1×10^{-4} per vehicle-year)	(4)
Exposure $X_j(45)$	2.26×10^5 vehicle-days, a 75 percent reduction	(3)

Break-even rate. Relation (4) gives $\rho^* = 1.40 \times 10^{-6}$ per vehicle-day, or 5.1×10^{-4} per vehicle-year. If the true rate of entrapment events per exposed vehicle-year is above that figure, one event was expected in the Ioniq 9 population during the 169 days, with probability 0.80 that the defect class was present at all. We do not claim that the true rate exceeds that figure; we claim that the question "is it below 5.1×10^{-4}?" is the one the register forces, for a hazard whose consequence is stated as injury to a trapped child [4]. Fig. 2 shows how X_j grows with latency. Deciding at 45 days instead of 169 would have reduced the accumulated exposure by 75 percent, to 2.26×10^5 vehicle-days, a figure that does not assume a faster test, only a deadline.

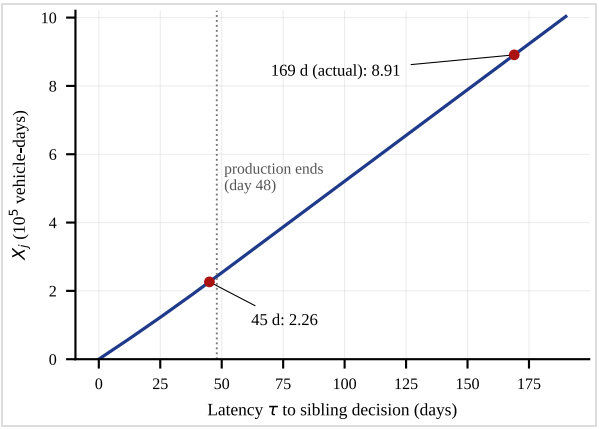


Fig. 2. *Sibling exposure from relation (3) for the Ioniq 9 population under the Table IV assumptions. The curve is steeper until day 48, when production ended, and linear afterwards because only the standing fleet remains exposed. Markers show a 45-day decision and the actual 169-day decision.*

Maximum latency. Table V solves relation (5) for an event budget of one over a range of assumed event rates, with $P_j = 0.80$ and the Table IV fleet. At 5×10^{-4} per vehicle-year the permitted latency is 173 days, close to what happened, which is another way of stating the break-even result. At 10^{-3} it is 88 days and at 5×10^{-3} it is 19 days. A manufacturer that believes its sibling event rate is as high as one in a thousand vehicle-years has, by its own estimate, 88 days. The 45-day default of EOC-R4 corresponds to a rate of about 2×10^{-3} per vehicle-year and is a reasonable ceiling where a child-entrapment hazard cannot be excluded.

Table 5. *Maximum latency τ^* from relation (5) for one expected event, with $P_j = 0.80$ and the Table IV fleet.*

Event rate ρ (per vehicle-year)	10^{-4}	5×10^{-4}	10^{-3}	5×10^{-3}
τ^* for $\varepsilon = 1$ (days)	855	173	88	19
τ^* for $\varepsilon = 0.1$ (days)	88	19	10	2

Screening. The edge classes in Table I are not equally testable. The S and L edges of this case can be screened by the same bench fixture that Hyundai built for the Palisade, so relation (2) applies directly: a clean screening of five conditions would have lowered the S mean from 0.50 to 0.22 and the L

mean from 0.60 to 0.30, giving $P_j = 0.46$ and permitting a lower priority in the order of (6). Hyundai's fixture result of August was not clean, and it converted the sibling to confirmed.

B. Ram ProMaster Scope Extension

The Ram ProMaster case tests the framework on a different edge: not a sibling program but a sibling population defined by model year and variant. NHTSA opened preliminary evaluation PE26-002 on March 2, 2026 on 19 owner questionnaires, covering 136,666 model-year 2022 and 2023 vehicles and water intrusion into a power steering module connector [8]. FCA determined a safety defect on September 22, 2026 and recalled 265,512 vehicles, model years 2022 to 2026, including 1,874 model-year 2024 to 2026 3500 electric vans, estimating that 7.7 percent of the population has the defect [7]. FCA's own records at the time of filing held 2,799 warranty claims, 437 customer assistance records and 191 field reports from April 2022 through June 2026 [7].

The recall population exceeds the investigated population by 128,846 vehicles, a factor of 1.94. These vehicles were in service during the 204 days between the opening of the preliminary evaluation and the defect determination [7], [8]. With production assumed to have ceased contributing ($\lambda = 0$, which understates exposure), relation (3) reduces to $X = N^0\ \tau$, or $128{,}846 \times 204 = 2.63 \times 10^7$ vehicle-days outside the investigated population. If the manufacturer's 7.7 percent prevalence applied uniformly, which is an upper-end assumption because newer vehicles have had less time to corrode, about 9,900 defective vehicles were in that population. The manufacturer-side record of 3,427 warranty, assistance and field records (through June 2026) is about 180 times the 19 owner questionnaires on which the regulator opened its evaluation (March 2026) [7], [8]; the dates differ, so the ratio is indicative. That ratio is a screening opportunity of the kind relation (2) prices: the manufacturer held evidence about the extension population before the regulator's scope included it.

C. Service Parts as a Third Edge Class

A third filing shows why Table I includes class R. Ford's campaigns 26V582 and 26E063 cover 86 vehicles that received a replacement battery energy control module and 83 service modules, respectively [9]. Press accounts of the filings say that some modules were sold over the counter or fitted under warranty without traceability to individual vehicles, which is why the part itself was recalled [28]. When part-to-vehicle traceability is incomplete, the graph's node set is itself incomplete, and the register must carry a node for the part population in addition to the vehicle population. Requirement EOC-R1 therefore includes service parts.

V. DISCUSSION

Limitations. The edge priors of Table I are elicited, not estimated. A calibrated version would fit q by class from a corpus of anchor-and-sibling recall pairs, and no such fit is attempted here. The uniform build rate in Table IV is our assumption, and daily production data would change the production term of (3), though not the standing-fleet term, which is 90 percent of the total in this case. The event rate ρ is unknown for both cases and the framework deliberately returns break-even rates and latency tables instead of a point estimate. Relation (1) assumes independent edge activation. Real edges are correlated, because two programs that share specification probably share logic, and an upper bound on reach can overstate the register for a tightly coupled family.

Threats to validity. The public record is a manufacturer-selected chronology, and a Part 573 report states what the manufacturer chose to state [26], [27]. A study based on such records can measure the latency between filed dates but not the quality of the internal analysis in between. The exposure measure counts calendar vehicle-days, not vehicle-days of use of the affected function, and the two differ by a factor that depends on the function. We have chosen two cases in which the sibling was eventually found, and a selection of that kind says nothing about the false-positive rate of an extent-of-condition process, which is the cost of evaluating and recalling siblings that turn out to be clean. Relation (2) is the framework's answer to that cost, since a clean screening is cheap, but the claim is untested.

Where the standards lens stops. ISO 26262 can require that the register exist, that edges be justified and that determinations be recorded. It cannot determine the legal moment at which a manufacturer's information amounts to a defect determination, which is the moment at which the five-working-day clock of [2] starts. It also cannot supply a supplier's design history, which the development interface agreement must provide under part 8, clause 5 [1]. When a supplier serves several manufacturers, the inheritance graph crosses company boundaries and a single manufacturer's register cannot be complete. Finally, the framework assumes that the anomaly can be reduced to a descriptor. The hardest anomalies, those whose mechanism is not understood, resist the typing that Table I requires, and for them EOC-R2 defaults to $q = 1$, which is correct but expensive.

VI. CONCLUSION AND FUTURE WORK

We have proposed an inheritance-graph framework that turns the extent-of-condition question after a confirmed safety anomaly into six relations and six requirements: a transfer probability, a screening update, an exposure integral with a production term, a break-even event rate, a maximum evaluation latency, and an ordering rule. Applied to the public NHTSA record, the framework prices the 169-day gap between the Palisade and Ioniq 9 seat decisions at about 8.9×10^5 sibling vehicle-days and shows that a 45-day deadline would have removed three quarters of it, and it prices the Ram ProMaster scope extension at about 2.6×10^7 vehicle-days outside the investigated population.

Future work has three directions. The first is calibration: fitting the Table I priors to a corpus of recall pairs with shared components, where the time between anchor and sibling reports is the observable that the framework predicts. The second is relaxing the independence assumption with correlated edge activation, using the cascade and product-line literature as the starting point. The third is integration with the safety case: a

register entry that is open should suspend the proven-in-use and independence claims that depend on the sibling, and an assurance-reuse tool of the kind surveyed in Section II-F could enforce that automatically.

REFERENCES

[1] International Organization for Standardization, *ISO 26262:2018, Road Vehicles — Functional Safety*, Parts 2, 3, 8 and 9, 2nd ed., Geneva, Switzerland, Dec. 2018. [Online]. Available: https://www.iso.org/standard/68384.html

[2] *Electronic Code of Federal Regulations, Title 49, Part 573 — Defect and Noncompliance Responsibility and Reports*. [Online]. Available: https://www.ecfr.gov/current/title-49/part-573

[3] *Electronic Code of Federal Regulations, Title 49 §571.118 — Standard No. 118; Power-operated window, partition, and roof panel systems*. [Online]. Available: https://www.ecfr.gov/current/title-49/section-571.118

[4] Hyundai Motor America, "Part 573 Safety Recall Report acknowledgement, campaign 26V160: Rear power seat entrapment, 2026 Palisade and Palisade Hybrid," National Highway Traffic Safety Administration, Washington, DC, USA, Mar. 17, 2026. [Online]. Available: https://static.nhtsa.gov/odi/rcl/2026/RCAK-26V160-3716.pdf

[5] Hyundai Motor America, "Part 573 Safety Recall Report acknowledgement, campaign 26V588: 2nd row power seat, 2026 Ioniq 9," National Highway Traffic Safety Administration, Washington, DC, USA, Sep. 10, 2026. [Online]. Available: https://static.nhtsa.gov/odi/rcl/2026/RCAK-26V588-0744.pdf

[6] Hyundai Motor America, "Part 573 Safety Recall Report 26V-588: Second-row power seat logic," National Highway Traffic Safety Administration, Washington, DC, USA, Sep. 10, 2026. [Online]. Available: https://static.nhtsa.gov/odi/rcl/2026/RCLRPT-26V588-0403.pdf

[7] FCA US LLC, "Part 573 Safety Recall Report 26V-621: Ram ProMaster electric power steering wiring harness," National Highway Traffic Safety Administration, Washington, DC, USA, Sep. 29, 2026. [Online]. Available: https://static.nhtsa.gov/odi/rcl/2026/RCLRPT-26V621-5191.pdf

[8] National Highway Traffic Safety Administration, Office of Defects Investigation, "Preliminary Evaluation PE26-002 opening resume: Loss of power steering assist, 2022–2023 Ram ProMaster," Washington, DC, USA, Mar. 2, 2026. [Online]. Available: https://static.nhtsa.gov/odi/inv/2026/INOA-PE26002-10013.pdf

[9] National Highway Traffic Safety Administration, "Acknowledgement letters, recall campaigns 26V582 and 26E063: Incorrect BECM software, Ford Mustang Mach-E," Washington, DC, USA, Sep. 14, 2026. [Online]. Available: https://www.autoevolution.com/pdf/news_attachements/ford-mustang-mach-e-with-lfp-battery-recalled-over-potential-loss-of-power-and-fire-risk-275623.pdf

[10] N. G. Rupp and C. R. Taylor, "Who initiates recalls and who cares? Evidence from the automobile industry," *J. Ind. Econ.*, vol. 50, no. 2, pp. 123–149, 2002, doi: 10.1111/1467-6451.00171.

[11] H. Bates, M. Holweg, M. Lewis, and N. Oliver, "Motor vehicle recalls: Trends, patterns and emerging issues," *Omega*, vol. 35, no. 2, pp. 202–210, Apr. 2007, doi: 10.1016/j.omega.2005.05.006.

[12] C.-F. Chi, D. Sigmund, and M. O. Astardi, "Classification scheme for root cause and failure modes and effects analysis (FMEA) of passenger vehicle recalls," *Rel. Eng. Syst. Saf.*, vol. 200, art. no. 106929, Aug. 2020, doi: 10.1016/j.ress.2020.106929.

[13] C. M. Kang, Y. S. Hong, W. T. Huh, and W. Kang, "Risk propagation through a platform: The failure risk perspective on platform sharing," *IEEE Trans. Eng. Manag.*, vol. 62, no. 3, pp. 372–383, Aug. 2015, doi: 10.1109/TEM.2015.2427844.

[14] N. E. Fenton and N. Ohlsson, "Quantitative analysis of faults and failures in a complex software system," *IEEE Trans. Softw. Eng.*, vol. 26, no. 8, pp. 797–814, 2000, doi: 10.1109/32.879815.

[15] T. Zimmermann and N. Nagappan, "Predicting defects using network analysis on dependency graphs," in *Proc. 30th Int. Conf. Softw. Eng. (ICSE)*, 2008, pp. 531–540, doi: 10.1145/1368088.1368161.

[16] E. Juergens, F. Deissenboeck, B. Hummel, and S. Wagner, "Do code clones matter?," in *Proc. 31st Int. Conf. Softw. Eng. (ICSE)*, 2009, pp. 485–495, doi: 10.1109/ICSE.2009.5070547.

[17] I. Abal, C. Brabrand, and A. Wasowski, "42 variability bugs in the Linux kernel: A qualitative analysis," in *Proc. 29th ACM/IEEE Int. Conf. Automated Softw. Eng. (ASE)*, 2014, pp. 421–432, doi: 10.1145/2642937.2642990.

[18] T. Thüm, S. Apel, C. Kästner, I. Schaefer, and G. Saake, "A classification and survey of analysis strategies for software product lines," *ACM Comput. Surv.*, vol. 47, no. 1, pp. 1–45, Jul. 2014, doi: 10.1145/2580950.

[19] S. Kokaly, R. Salay, V. Cassano, T. Maibaum, and M. Chechik, "A model management approach for assurance case reuse due to system evolution," in *Proc. ACM/IEEE 19th Int. Conf. Model Driven Eng. Lang. Syst. (MODELS)*, 2016, pp. 196–206, doi: 10.1145/2976767.2976792.

[20] D. Kempe, J. Kleinberg, and É. Tardos, "Maximizing the spread of influence through a social network," in *Proc. 9th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, Washington, DC, USA, 2003, pp. 137–146, doi: 10.1145/956750.956769.

[21] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, "Epidemic processes in complex networks," *Rev. Mod. Phys.*, vol. 87, no. 3, pp. 925–979, Aug. 2015, doi: 10.1103/RevModPhys.87.925.

[22] M. C. M. Troffaes, G. Walter, and D. Kelly, "A robust Bayesian approach to modeling epistemic uncertainty in common-cause failure models," *Rel. Eng. Syst. Saf.*, vol. 125, pp. 13–21, May 2014, doi: 10.1016/j.ress.2013.05.022.

[23] A. Mosleh *et al.*, *Procedures for Treating Common Cause Failures in Safety and Reliability Studies*, U.S. Nuclear Regulatory Commission, Washington, DC, USA, NUREG/CR-4780, vols. 1–2, 1988–1989. [Online]. Available: https://www.nrc.gov/sr0933/reference/1119.html

[24] W. E. Smith, "Various optimizers for single-stage production," *Nav. Res. Logist. Q.*, vol. 3, no. 1–2, pp. 59–66, 1956, doi: 10.1002/nav.3800030106.

[25] C. Picoco and V. Rychkov, "Modeling common cause failure in dynamic PRA," arXiv:2406.08879, Jun. 2024. [Online]. Available: https://arxiv.org/abs/2406.08879

[26] D. J. Deka and N. Sarkar, "Compliance evidence in the automotive supply chain: A systematisation of the quality-document spine and a taxonomy of documentation failure modes," arXiv:2607.04924, Jul. 2026. [Online]. Available: https://arxiv.org/abs/2607.04924

[27] H. Ihshaish, P. Mayhew, T. M. A. Zayet, and A. Del Amo, "The record is part of the task: Matched-record evaluation of text classifiers across maintenance, safety and recall reporting," arXiv:2609.16267, Sep. 2026. [Online]. Available: https://arxiv.org/abs/2609.16267

[28] "Ford Mustang Mach-E with LFP battery recalled over potential loss of power and fire risk," *autoevolution*, Sep. 15, 2026. [Online]. Available: https://www.autoevolution.com/news/ford-mustang-mach-e-with-lfp-battery-recalled-over-potential-loss-of-power-and-fire-risk-275623.html