Taking Advantage of Recent Laws: Efficient Licensing Pathways for Advanced Reactors

Ernie Kee · preliminary draft, July 2026

About this document (background added November 2026)

Between July 2024 and April 2026 the rules for licensing advanced reactors in the United States changed more than in the preceding four decades. The ADVANCE Act (July 2024) redirected the Nuclear Regulatory Commission toward enabling deployment; executive orders in May 2025 imposed an eighteen-month deadline for licensing decisions; and the final Part 53 rule, effective 29 April 2026, accepts "systematic risk evaluations" alongside probabilistic risk assessment as a licensing basis.

This draft lays out a licensing roadmap that takes advantage of those changes, built on the Advanced Reactor Improved Safety Evaluation (ARISE) framework, which Ernie Kee first set out in a May 2025 white paper and which echoes the containment-first idea of his 2020 note to congressional staff. ARISE scopes regulation to containment, bounds the radioactive source term, judges containment performance by an energy balance and by simulation of its penetrations, and replaces contested accident-frequency estimates with a safety-margin criterion that can be monitored through performance-based oversight. It is posted here as an unpublished preliminary draft.

The draft that follows is reproduced as written in July 2026, with the author's e-mail address removed.

Taking advantage of recent laws: efficient licensing pathways for advanced reactors

Preliminary draft R1, July 2026

Ernie Kee

YKrisk, LLC, Sagle, ID, USA

Abstract

Between July 2024 and April 2026 the statutory and regulatory environment for licensing advanced commercial nuclear fission reactors changed more than in the preceding four decades: the ADVANCE Act redirected the regulator's mission toward enabling deployment, executive action imposed an eighteen-month licensing decision deadline, and the final 10 CFR Part 53 rule admits "systematic risk evaluations" alongside probabilistic risk assessment. This paper presents a licensing roadmap that exploits these changes. The technical core is the ARISE framework: regulation is scoped to containment; the source term is bounded by the equilibrium inventory of 100% power operation with infinite operating history; containment expected performance is judged by a first-law total-energy balance and by simulation of the penetration set that every containment must possess in order to deliver energy to the process it serves. A marked-point-process discrete-event simulation, calibrated with industry-average valve performance data and leak-rate-test experience, demonstrates the safety-margin acceptance criterion $P(T > T_L) \geq \gamma$: with Appendix J-class surveillance, the probability of no release over a 60-year license exceeds 0.99, against a requirement of $\gamma = 0.95$. The roadmap couples this quantitative demonstration to the new statutory hooks—the Part 53 systematic-risk-evaluation pathway, fee reductions, licensing prizes, and staff-secondment efficiency measures—into a phased plan for a generic, technology-inclusive containment-based license.

Keywords: advanced reactors, licensing, containment, safety margin, marked point process, ADVANCE Act, 10 CFR Part 53

Nomenclature

Acronyms

A-CNFR Advanced Commercial Nuclear Fission Reactor

ADVANCE the Accelerating Deployment of Versatile Advanced Nuclear for Clean Energy Act of 2024

AEA the Atomic Energy Act of 1954

ARISE the Advanced Reactor Improved Safety Evaluation

CDF Core Damage Frequency

CHR Containment Heat Removal

DES Discrete Event Simulation

DID Defense in Depth

ETtA the Expected Time to Accident

IPE Individual Plant Examination

LERF Large Early Release Frequency

LLRT Local Leak Rate Test

MPP Marked Point Process

NEI the Nuclear Energy Institute

NEIMA the Nuclear Energy Innovation and Modernization Act

PRA Probabilistic Risk Assessment

ROP the Reactor Oversight Process

SER Safety Evaluation Report

SM Safety Margin

SRE Systematic Risk Evaluation

TS Technical Specifications

USNRC the U.S. Nuclear Regulatory Commission

Symbols

T	random time to release of contained radioactivity
T_L	license expiration time
γ	required confidence level (0.95 here)

t^*	passive grace period: first time the total-energy margin reaches zero
$M(t)$	energy margin, $E_{\text{cap}}(t) - E_{\text{chal}}(t)$
E_{chal}	bounding challenge energy (stored + decay + chemical + feed)
E_{cap}	containment energy-absorption capacity
P_0	rated thermal power
$Q_t,\ Q_t^C$	initiating-event and accident counting processes
A_t	protection availability process
λ_t	stochastic intensity of initiating events
p	per-demand failure probability (valve fail-to-close)
q	steady-state latent-failure (pre-existing leak) probability

1. Introduction and motivation

Alvin Weinberg described nuclear energy as a Faustian bargain: an effectively inexhaustible energy source whose side effects, if uncontrolled, could spell disaster [42]. Writing two decades later, he added the prescription that a rebirth of nuclear energy would require "a deterministic, non-probabilistic cap on the worst consequences of even the unlikeliest event" [43]. The framework presented here takes that prescription literally. Its thesis is that the cap Weinberg asked for is the containment; that the containment can be judged by a universally defensible thermodynamic bound; and that the statutes and rules adopted between 2024 and 2026 make a containment-scoped license not merely conceivable but procedurally available.

The regulatory arc that leads here is familiar. the Atomic Energy Act of 1954 (AEA) requires "adequate protection" of public health and safety. The WASH-1400 study introduced quantitative accident-sequence analysis [23], and after Three Mile Island the plant-specific examinations institutionalized Probabilistic Risk Assessment (PRA) in regulatory practice; a critical history of that institutionalization, and of what its critics said at each stage, is given in Appendix Appendix B. the Nuclear Energy Innovation and Modernization Act (NEIMA) [26] and the Accelerating Deployment of Versatile Advanced Nuclear for Clean Energy Act of 2024 (ADVANCE) [27] then added words that do not appear in the AEA: *efficiency*, *risk-informed*, and *performance-based*. What has been missing is a method that satisfies those words for reactor designs that lack the operating history on which PRA calibration depends. This paper supplies one and maps it onto the licensing process.

2. The 2024–2026 regulatory landscape

Five developments define the opportunity; each is a hook the roadmap of Section 9 attaches to.

The ADVANCE Act. Public Law 118-67 [27] revised the the U.S. Nuclear Regulatory Commission (USNRC) mission statement toward *enabling* the safe and secure use of nuclear technology, reduced advanced-reactor applicant fees to roughly the direct-cost rate (effective October 2025), created prizes that refund licensing fees for first movers in five categories, and directed expedited combined-license procedures and micro-reactor licensing guidance.

The final Part 53 rule. Approved March 25, 2026 and effective April 29, 2026, the risk-informed, technology-inclusive framework requires applicants to use "a PRA, other Systematic Risk Evaluations (SREs), or a combination thereof" to identify failures and demonstrate the Subpart B safety criteria [41]. The SRE clause—added in response to industry and public comment—is the direct regulatory pathway for the simulation methodology of this paper. It did not exist in any prior licensing framework.

Executive action. The May 2025 executive orders [6] impose an eighteen-month decision deadline on new reactor applications, direct a revision of the Reactor Oversight Process (ROP), and create an expedited approval pathway for designs tested under Department of Energy authorization.

The DOE pilot population. Eleven advanced designs entered the DOE reactor pilot program in 2025 [28]; the first non-light-water construction permit (March 2026) and multiple pilot criticality milestones create the tested-design population the executive-order pathway is built to receive.

Industry alignment. the Nuclear Energy Institute (NEI) guidance on technology-inclusive development of the licensing basis [21] and the industry regulatory-modernization effort provide the consensus infrastructure into which a generic containment methodology can be introduced as a topical report. NEI's February 2025 white paper on TI-RIPB approaches under Part 53 [22] is read against the present framework in Appendix Appendix A, where backchecks at Three Mile Island and Fukushima Daiichi are applied to its qualitative likelihood classifications.

3. The ARISE framework in brief

the Advanced Reactor Improved Safety Evaluation (ARISE) models reactor operations as a Marked Point Process (MPP): the plant state process $\mathbf{Z}$ evolves on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, P)$, initiating events arrive as a counting process $\mathbf{Q}$ with stochastic intensity λ_t, and accidents are the thinned process

$$Q_t^C = \int_{[0,t]} (1 - A_u)\, dQ_u, \tag{1}$$

where A_t is the protection availability process [3, 46]. The classical PRA metrics arise from (1) as limiting special cases under a hierarchy of increasingly strong assumptions—no unknown-unknowns, no absorbing states, Poisson arrivals seeing time averages [44], stationarity, ergodicity, and independence of protections from initiating events—each of which biases the estimate optimistically when violated. The full development, including the martingale decomposition of (1) and the assumption hierarchy, is given in the companion problem statement and white paper; it is not repeated here.

Two consequences matter for licensing. First, because (1) is solved by Discrete Event Simulation (DES) over a finite horizon rather than by stationary limits, the framework produces the *distribution* of the time to accident, so the acceptance criterion can be posed as a chance constraint,

$$P(T > T_L) \geq \gamma, \qquad \gamma = 0.95, \tag{2}$$

with T the random release time and T_L the license term—Safety Margin (SM) expressed in the engineering unit of time. Second, because the method relaxes rather than contradicts the PRA assumptions, it qualifies as an "other SRE" under the final Part 53 rule: it reproduces Core Damage Frequency (CDF) and Large Early Release Frequency (LERF) where their assumptions hold and adds information where they do not. The same property lets the method serve as the safety-evaluation instrument inside licensing-basis *selection* frameworks (Section 10).

4. Regulatory scope reduction to containment

The empirical record of reactor accidents shows declining frequency and increasing severity concentration: improved protections have made events rarer, while the events that do occur cluster at high International Nuclear Event Scale levels [7]. Both trends point the same direction: the regulatory center of gravity belongs on the last barrier, the containment. All Advanced Commercial Nuclear Fission Reactors (A-CNFRs), regardless of coolant, fuel, or spectrum, share exactly one regulatory obligation—to prevent the release of radioactive material. We therefore propose that the regulated ("safety-related") scope be limited to the containment function:

> *Containment is the Defense in Depth (DID) element that, given failure of the first barrier to radioactive release, must retain the reactor's radioactive inventory under the physical conditions presenting the greatest potential challenge to confinement.*

The bounding source term makes the scope generic. Under ARISE, containment performance is evaluated assuming the *complete release of the equilibrium radionuclide inventory corresponding to 100% power operation with infinite operating history.* No credit is taken for power history,

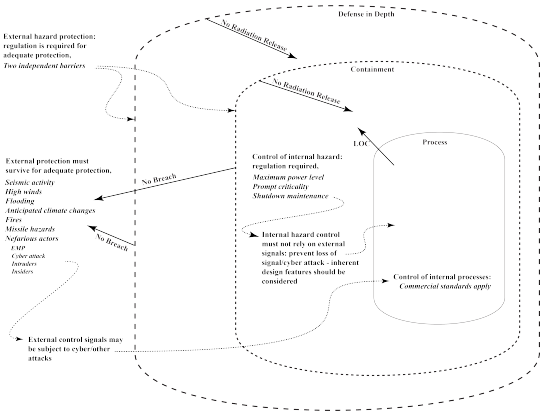


Figure 1: Nested regulated scope: the production process (subject to commercial controls) sits inside the containment boundary ("no radiation release"), which sits inside the defense-in-depth envelope ("no breach"). Only the outer two rings, and the process inputs that penetrate them, fall under regulation.

burnup, or partial release. Because the bound depends only on rated power, it removes the source term—the most design-specific and most contested element of severe-accident analysis—from the licensing critical path. Power *measurement* remains under regulation, since calorimetric calibration is what bounds the equilibrium inventory itself; power measurement physics is common across designs and is a natural subject for consensus standards.

5. A universally defensible performance basis: total energy

What can defeat a containment is energy. The first law provides a design-agnostic bound: the challenge energy transferred to the contained volume by time t after the bounding release,

$$E_{\text{chal}}(t) = E_{\text{stored}} + \int_0^t P_d(u)\, du + E_{\text{chem}} + E_{\text{feed}}(t), \tag{3}$$

comprises the stored thermal energy of the working fluid, the decay power integral for infinite prior operation at P_0, the full chemical energy of the reactive inventory, and—critically—any energy fed through an unisolated energy-transfer penetration. The capacity side aggregates free-volume pressurization to the ultimate capacity, passive heat sinks, and credited active heat removal. The Wortman all-states margin condition [46] requires

$$M(t) = E_{\text{cap}}(t) - E_{\text{chal}}(t) > 0 \quad \text{for all } t, \tag{4}$$

for every reachable protection state; its first violation time defines the passive grace period t^*.

For a generic 3000 MWt plant the passive grace period evaluates to $t^* \approx 24$ hours—the correct physical order for large dry containments—and a single 30 MW heat-removal train maintains positive margin indefinitely, since decay power falls below removal capacity within hours. The

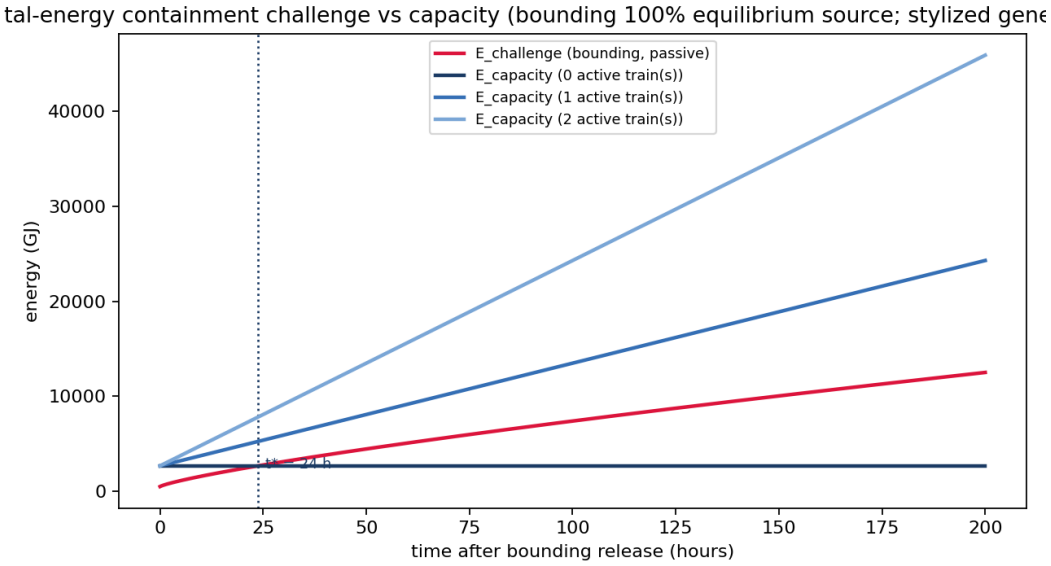


Figure 2: Total-energy challenge versus capacity for a generic 3000 MWt plant (stylized capacity values). Passive grace period $t^* \approx 24$ h; a single 30 MW active heat-removal train holds the margin indefinitely.

decisive result concerns the energy-transfer penetration: with that path unisolated at power, the feed term of (3) grows the challenge at P_0 rather than at decay power, and the grace period collapses from 24 hours to roughly 12 minutes. The penetration that exists *because* the plant is useful is the only one that can defeat the containment on the load side as well as the release side. Containment regulation is therefore, to first order, penetration regulation.

The energy balance also fixes the regulated scope on the capacity side: *everything credited in* E_{cap} *of* (4) *is, by that act, a regulated component of containment.* A credited active heat-removal train—the 30 MW figure above is a stylized representative of the containment fan-cooler and spray class, whose only requirement is to exceed decay power at the time the passive heat sinks are exhausted—enters the technical specifications together with its power supply and heat-sink path, and its availability is what the t^* race of Section 6 exercises. The corollary is a design incentive the framework makes explicit: a design that declines active credit and demonstrates $t^* = \infty$ on passive capacity alone purchases a strictly smaller regulated scope.

6. Simulating the penetration set: expected time to release

Every containment is penetrated: by the energy-transfer path, by process lines, by purge and vent lines, by personnel and equipment access, by instrumentation and electrical canisters, and—in prestressed designs—by the post-tensioning system whose degradation erodes structural capacity. Each path was modeled as an isolation system (series isolation valves succeed if at least one closes; passive seals are single boundaries) in a MPP/DES over a 60-year license, with isolation demands arriving at a stylized rate of 0.1 per year and the bounding inventory assumed present at every demand.

The model is calibrated with industry-average performance data. Per-demand fail-to-close probabilities are the 2006–2020 industry means [10, 4]: 2.28×10^{-4} for motor-operated valves, 1.89×10^{-4} for air-operated valves (energy-path class), and 2.89×10^{-4} for butterfly valves (purge class). Latent, release-significant degradation of the passive boundaries is calibrated so that the containment-wide steady-state pre-existing large-leak probability matches the value used in the leak-rate-testing risk methodology, approximately 2.7×10^{-3} [5, 34], apportioned across the hatch, airlock, electrical, and tendon paths with discovery delays set by their actual surveillance intervals. The Local Leak Rate Test (LLRT) record—148 of 2,055 extended-interval valves leaking above administrative limits, dominated by small leaks—supports a per-valve release-significant latent rate near 10^{-3} per year, which the series-pair geometry renders negligible at the line level. The two-train Containment Heat Removal (CHR) system couples the energy model of Section 5: if both trains are unavailable at a demand, the trajectory enters the t^* race defined by (4), and release follows unless a train is restored within the grace period.

The results demonstrate the acceptance criterion (2) directly: the probability of release within 60 years is 0.009 ± 0.005, the Expected Time to Accident (ETtA) conditional on release is 27 years, and

$$P(T > 40\,\mathrm{yr}) = 0.993, \qquad P(T > 60\,\mathrm{yr}) = 0.991,$$

both exceeding $\gamma = 0.95$ (Figure 3). Three structural findings survive any refinement of the input numbers. First, the active, frequently-tested penetrations—including the energy path—contribute a small fraction of release probability; series redundancy and two-year LLRT intervals

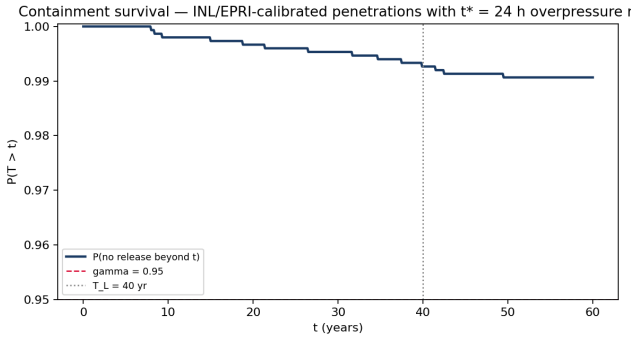


Figure 3: Containment survival $P(T > t)$ over a 60-year license, penetration model calibrated with INL industry-average valve data and the EPRI/NEI pre-existing-leak probability, with the $t^* = 24$ h overpressure race. The chance constraint (2) is met with margin.

control them. Second, the residual risk is dominated by the passive, long-surveillance-interval boundaries (equipment hatch, airlock, electrical canisters, tendons), whose governing variable is the *latent-failure discovery interval*, not the demand-failure probability. This is the empirical face of a general law: the consequential contribution of latent dependence scales linearly with the surveillance interval, while independent coincidence scales only quadratically [14]. Surveillance intervals are precisely what performance-based regulation controls; the framework therefore regulates the quantity that matters. Third, the t^* race is unforgiving of slow restoration: with a median repair time of about two days, ninety percent of both-trains-down demands still lose a 24-hour race, so grace-period credit must be earned with demonstrated restoration times rather than assumed.

7. Performance-based oversight

The license application under this framework asserts, for each penetration path, an assumed performance level: the fail-to-close probabilities, latent-degradation rates, discovery intervals, and restoration times of Section 6. Each assertion is directly verifiable in service. Observed cumulative failures $N_{\text{obs}}(t)$ are compared with the allowed count $N_{\text{allowed}}(t)$ implied by the assumed distributions; leak-rate-test findings update the latent-degradation calibration; restoration times are demonstrable in drills. When observed performance erodes the demonstrated margin—when the recomputed $P(T > T_L)$ approaches γ—a backfit restoring the margin is the regulatory response. This is the ordinary logic of warranties, applied to the containment: the Technical Specifications (TS) become a small set of measurable, penetration-level commitments, and ROP (now under executive direction to be revised [6]) monitors exactly the quantities the safety case uses. The 95/95 convention familiar from criticality practice carries over as the statistical form of the verification.

8. Process efficiency: secondment and fees

The technical scope reduction enables a procedural one. Because the regulated design is the containment and its penetration set, the license application and Safety Evaluation Report (SER) concern a bounded, design-generic engineering package. We propose seconding the licensee's containment engineering staff into the USNRC organization from application submittal through SER completion: requests for additional information become working sessions rather than correspondence cycles, which is the single largest schedule lever available under the eighteen-month statutory deadline [6]. The secondment must be designed against regulatory capture—seconded staff work under USNRC procedures and must adopt the regulator's utility on risk, a training objective in its own right; the Fukushima and recent aviation experiences illustrate the failure mode. The ADVANCE fee structure (direct-cost rates from October 2025) and the Section 202 prize categories reduce the financial exposure of the first applicant to carry this framework through review.

9. The roadmap

Table 1 assembles the elements. The sequence is deliberately incremental: the methodology is introduced as an SRE *complementing* PRA in Phase 1, where the final Part 53 rule explicitly invites it, before carrying a containment-scoped application in Phase 2. Each phase has a statutory hook that did not exist before 2024, and each produces artifacts (topical report, validated code, penetration data) that survive even if a later phase stalls.

10. Relation to licensing-basis selection frameworks

A complementary framework, developed independently and shared with the author,[1] operates at the decision layer above this paper: it structures licensing as *selection among candidate licensing bases*. Candidate bases (the applicant's proposal plus alternative requirement sets) are each evaluated for safety—individual-plant and *cumulative regional* (multi-unit interactions, common-cause vulnerabilities, shared infrastructure, deployment scenarios)—then subjected to a societal benefit–cost evaluation; the license issues on the benefit-justified basis offering the greatest overall safety, with no fixed numerical risk threshold serving as the sole decision rule.

The two frameworks compose rather than compete. The selection framework's safety-evaluation step calls for "the best available methods"; the ARISE SRE of Sections 3 and 6 is such a method, supplying the full $P(T > T_L)$ profile

[1] "Proposed Framework: Selecting the Licensing Basis Under the ADVANCE Act—Safety First, Benefit Justification, and Cumulative Regional Safety," personal communication, July 2026; to be cited upon publication.

Table 1: Licensing roadmap: phases, statutory hooks, and status.

Phase	Content	Statutory hook	Status
0. Pre-application (now)	Consortium (NEI, EPRI, national laboratories); topical report on the containment-scoped SRE; pre-application engagement	Part 53 SRE clause; NEI 18-04 infrastructure	Open since April 2026
1. Methodology pilot	ARISE DES as SRE on a micro-reactor or DOE-pilot graduate; validation against RG 1.247-class expectations	Part 53 Subpart B; DOE-tested-design expedited pathway	DOE pilot population exists
2. Containment-scoped application	Bounding-inventory containment case (Sections 5 and 6); penetration-level TS; second-ment agreement; 18-month clock	ADVANCE fees and §202 prizes; executive-order deadline	First mover sought
3. In-service oversight	Penetration performance indicators feeding the revised ROP; margin-erosion backfit criterion	Executive-order ROP revision	Follows license

for each candidate basis where point-value CDF/LERF estimates would otherwise stand. Conversely, the containment-scoped basis proposed here enters the selection tournament as one candidate—one this paper argues occupies the favorable region of the decision space, combining a demonstrated time-domain margin with the cost structure of a reduced regulated scope. The selection framework also formalizes the regulator's utility on risk, which Section 8 invokes only qualitatively.

Two tensions deserve candor. First, the selection framework demonstrates adequate protection *through* the benefit-justified choice, making the finding relative to the candidate set and entangled with cost; the long-standing reading of AEA treats adequate protection as an absolute finding made without regard to cost. The division of labor proposed here is the conservative one: the chance constraint (2) and the margin condition (4) supply an absolute safety floor, and benefit–cost selection operates on candidates above that floor. Second, the selection framework rejects fixed numerical thresholds as the sole acceptance rule, while this paper retains $\gamma = 0.95$; the reconciliation is that γ bounds a time-domain margin rather than a frequency, and within the composed framework it functions as the floor just described, not as the decision rule. The cumulative regional safety dimension decomposes on inspection. Its dependence content—common-cause coupling across the units of a site, and fleet-level coupling through the corrective-action and 10 CFR Part 21 notification channel—is already treated in time-gated terms in the author's companion work [14], where the fleet behaves as a very large multi-unit site with discovery lag. What remains genuinely open is the *regional aggregation* layer: shared infrastructure, dense-deployment emergency-planning effects, and simultaneous regional demands. That remainder is taken up in Section 11.

11. Challenges

Candor about the open items strengthens the case. The calibration argument cuts both ways: the same scarcity of data that limits PRA for novel designs limits the simulation's inputs, and the framework's answer—that penetration hardware is common across designs, so industry-average valve and seal data apply where design-specific data do not exist—must be defended explicitly. Rare-event computation becomes binding when per-penetration resolution at 10^{-4}–10^{-6} probabilities is required for design comparison; importance sampling or trajectory splitting, not brute replication, is the appropriate machinery, and it reduces rather than increases the computational burden relative to naive Monte Carlo at those levels. Verification and validation of the simulation to consensus-standard expectations (the SRE analog of RG 1.247 [40]) is a prerequisite for Phase 1. The engine's treatment of latent failures should be extended so that undiscovered conditions are invisible to the modeled operator until surveillance reveals them, a distinction the present calibration approximates through discovery-delay repair times. Of the cumulative regional safety dimension of Section 10, the dependence content (multi-unit and fleet common-cause coupling) is treated in [14]; the remaining gap is the regional aggregation layer—shared infrastructure, dense-deployment emergency-planning effects, and simultaneous regional demands—which the MPP construction reaches naturally through units coupled to a common environment process, and which is programmed future work. Finally, no quantification method escapes unknown-unknowns; the framework's honest claim is not that it forecasts them but that its filtration structure incorporates each discovery gracefully, and that the total-energy bound of Section 5 caps their consequences—which is exactly Weinberg's deterministic cap.

12. Conclusions

The 2024–2026 statutes and rules opened a pathway that the nuclear field has requested for decades: a lawful route to license a reactor on the strength of its last barrier, judged by measurable engineering quantities. The ARISE framework supplies the technical content for that route. Regulation scoped to containment is generic because the bounding source term depends only on rated power; containment expected performance is defensible

because it rests on the first law; and the penetration-set simulation, calibrated entirely with public industry-average data, already demonstrates the safety-margin chance constraint with margin for a representative configuration. The residual risk concentrates where performance-based regulation has its instruments—surveillance intervals on passive boundaries—and the grace period bought by containment thermal capacity is quantified rather than assumed. The roadmap of Table 1 converts these results into a phased licensing strategy whose every step attaches to a statutory hook enacted since 2024.

AI Use Acknowledgment

Artificial intelligence tools assisted in the research, drafting, computation, and editorial preparation of this document, including the simulation runs and figures reported in Sections 5 and 6. The author specified, reviewed, and verified all technical content, data selections, and conclusions.

Verification note

Regulatory facts (ADVANCE Act provisions, Part 53 final-rule status and SRE language, executive-order deadlines, DOE pilot milestones) were verified against primary and trade sources in July 2026. The INL 2020-update valve probabilities and the leak-rate-test statistics were taken from the cited reports; the MSV-specific fail-to-close table of Idaho National Laboratory [10] remains to be transcribed (energy-path valves presently carry the air-operated-valve class value). All containment-model capacity parameters are stylized pending design data. The Three Mile Island facts of Appendix Appendix A (initiator, relief-valve failure record, Davis–Besse precursor) were verified against the Kemeny Commission report and NRC historical accounts in July 2026. The Fukushima facts (design-basis tsunami values, the 2008 operator calculation and its March 7, 2011 reporting date, the Jōgan recurrence estimate, the absence of a pre-2011 backfit obligation) rest on the NAIIC report and the cited paleotsunami literature and should be re-verified against the NAIIC English edition before submission. The ensemble figures (15,080 reactor-years, the 13-accident power-reactor data set, the five-year window probabilities) were transcribed directly from Ha-Duong and Journé [7]; the cumulative B&W fleet-exposure figure is the author's estimate from commercial-operation dates and should be firmed against NUREG operating-experience data before submission. The provenance of the qualitative likelihood classification table discussed in Appendix Appendix A remains to be confirmed. Appendix Appendix B is adapted from the addendum to the companion preventive-maintenance review [13]; its Federal Register citations, quoted passages (NUREG-0880, Generic Letter 88-20), and document numbers were carried over from that review and should be re-verified against the primary sources before submission.

Appendix A. The NEI TI-RIPB white paper read against ARISE, with a backcheck at Three Mile Island

The document and its claims. NEI's February 2025 white paper [22], submitted as Enclosure 1 to its comments on the proposed Part 53 rule, argues that the rule must accommodate a spectrum of TI-RIPB methodologies and illustrates the point with four worked examples of decreasing PRA reliance: Example A, the full NEI 18-04/LMP process; Example B, deterministic analysis with confirmatory PRA; Example C, an IAEA-guidance approach with PRA insights used only to confirm deterministic assessments; and Example D, a risk-informed bounding approach for microreactors in which a maximum hypothetical accident (MHA) bounds all events below 10^{-4} per year. Accompanying the copy shared with the author is a qualitative likelihood classification table[2] that assigns event classes by annual frequency: *Anticipated* above 10^{-2}, *Unlikely* between 10^{-2} and 10^{-4}, *Extremely Unlikely* between 10^{-4} and 10^{-6} ("events that will probably not occur during the lifetime of the facility"), *Beyond Extremely Unlikely* below 10^{-6}, and an MHA/MCA row bounding all credible events. Notably, the white paper's own summary concedes that Examples B and C "work well for designs like LWRs where the licensing basis accidents are well understood thanks to operating experience," and that even Example A requires "some OE. . . as well as enough data to support a PRA analysis."

Where the frameworks agree. The white paper and this article push on the same door. Both hold that an all-hazards, all-modes PRA cannot be the sole admissible instrument for designs without operating history—the position the final rule's SRE clause now codifies. And Example D's bounding MHA is recognizably Weinberg's deterministic cap, the same prescription Section 5 makes precise through the bounding equilibrium inventory and the all-states margin condition. On regulatory scope, Example D's concentration of analysis on the bounding release barrier is congenial to the containment-scoped proposal of Section 4.

Where they diverge. The divergence is in what fills the space the PRA vacates. In Examples B through D, the replacement is engineering judgment exercised at licensing time and expressed through static qualitative likelihood classes of the Table 12-1 type: an event sequence is assigned a frequency bin, the bin determines its treatment, and nothing in the proposed process measures, in service, whether the assignment was right. The frequency claims are asserted once and never verified. ARISE fills the same space differently: the asserted quantity is a time-domain margin,

[2] "Table 12-1, Example of Qualitative Likelihood Classification," received as an image accompanying the white paper; its bins and examples follow the DOE-STD-3009 style of qualitative hazard classification. Original source to be confirmed.

(2), produced by simulation whose inputs are named and calibrated; the oversight program of Section 7 verifies the asserted penetration-level performance against in-service evidence; and the filtration structure of the MPP construction incorporates each new observation into the current state of knowledge as a matter of course. The difference is not conservatism but verifiability—one framework's likelihood statements are licensing-time claims, the other's are standing predictions exposed to data.

A backcheck at Three Mile Island. A classification scheme of the Table 12-1 type can be backchecked: applied retrospectively to a fleet whose outcomes are known. TMI-2 is the natural test because it is the case where operating experience was thinnest—precisely the regime for which the white paper proposes judgment-based classification. The initiating event of March 28, 1979 was a blockage during resin transfer in a condensate polisher, a routine and recurring maintenance problem; the resulting loss-of-feedwater transient sits squarely in the *Anticipated* bin. The escalation mechanism, a power-operated relief valve that failed to reseat, had a documented record: the valve type had failed eleven times previously at B&W plants, nine times in the open position, and the full sequence—stuck-open relief valve masked from operators during a feedwater transient—had occurred at Davis–Besse eighteen months earlier, on September 24, 1977, at 9% power, where operators identified the failed valve in twenty minutes [16]. A qualitative classification exercised in 1978 would have placed the composite sequence (anticipated transient, relief-valve failure, failure to recover, core damage) in the *Extremely Unlikely* bin or below: an event that "will probably not occur during the lifetime of the facility." What the record shows instead is a core-damage accident within roughly thirty cumulative reactor-years of B&W commercial operation—nine units, the accident unit itself not ninety days into commercial service. If the sequence class truly carried a frequency of 10^{-4} per reactor-year, the probability of observing it within thirty reactor-years is about 3×10^{-3}; the observation rejects the bin assignment at the 99.7% level, and the demand record of the valve (nine fail-open events in eleven challenges) contradicted the assignment *before the accident occurred.* The failure was therefore structural, not numerical: evidence sufficient to overturn the classification existed inside the fleet's own operating record, and the classification framework had no channel through which that evidence could act—the Davis–Besse precursor was analyzed by B&W and never propagated to TMI's operators. A small fleet is exactly where judgment-based bins are most fragile, and exactly where the white paper proposes to rely on them most.

A second backcheck at Fukushima Daiichi. The Fukushima Daiichi accident of March 11, 2011 backchecks the external-hazard tail of the same classification style, and it exhibits a failure mode TMI does not. The design-basis tsunami at the site was 3.1 m at licensing, revised to 5.7 m in 2002; the March 2011 run-up was 13–15 m. The contradicting evidence was again in hand, and this time quantitative: sediment deposits of the 869 Jōgan tsunami, published a decade before the accident, established comparable inundation of the same coast with a recurrence interval of roughly 800–1,100 years [19]; the government's 2002 long-term seismic assessment allowed a magnitude-8-class tsunamigenic earthquake anywhere along the Japan Trench; and the operator's own 2008 calculation from that source produced 15.7 m at the site—reported to the regulator on March 7, 2011, four days before the wave [20]. The correction channel was again missing: the pre-2011 Japanese framework carried no backfit obligation, and the Diet commission's central finding was regulatory capture—the Kemeny finding at the scale of a national system [20]. The unpropagated precursor was again available: external flooding disabled safety trains at Blayais in December 1999 and flooded the Madras pump house in 2004, with no channel forcing the lesson into the Japanese licensing basis. The new failure mode is inferential. The millennium elapsed since Jōgan was read as *confidence*—no tsunami of that size in a thousand years—when for a strain-accumulation process the conditional intensity increases with elapsed time: by 2011, 1,142 years after Jōgan, the conditional probability was near its historical maximum, and the instrumental record was far too short to support any claim about a thousand-year tail (no tsunami in four hundred years of records is consistent with a one-in-a-thousand-year event about two-thirds of the time). Static frequency bins invite exactly this misreading because elapsed time appears nowhere in them; retaining it is precisely what distinguishes the MPP conditional intensity of Section 3 from the stationarity assumption in the PRA hierarchy, which is the formal point at which the elapsed-time information is discarded. Fukushima also backchecks the independence assumption directly: the initiating event disabled the protections across three operating units of a shared site, the correlation structure treated in [14]. Candor requires the last entry. At TMI the containment held; at Fukushima the Mark I containments did not. The lesson runs toward rather than against the present framework: the accident was a loss-of-containment-heat-removal race lost on a clock—the grace-period quantity t^* of Section 5—and an all-states margin condition under a bounding inventory is the instrument that would have exposed the thin Mark I margin before the event rather than after.

An ensemble backcheck: the empirical accident record. The two case backchecks generalize. Ha-Duong and Journé [7], cited in Section 4 for the severity trend, reconstructed the historical record of accidents at commercial power reactors—reconstruction was necessary because, as they document, no authoritative, comprehensive, public accident data set exists—and confronted it with the PRA-based objectives directly. Regulators' core-damage frequency objectives of 10^{-4}–10^{-6} per reactor-year imply between 0.015 and 1.5 accidents over the industry's 15,080 cumulative reactor-years

through 2012; the reconstructed record contains thirteen INES level-4-or-above accidents at power reactors (fifteen if the Fukushima units are counted separately), a realized rate of order 9×10^{-4} per reactor-year—between one and three orders of magnitude above the asserted range, depending on which end of it is claimed. Sampled from their record, the probability of at least one INES ≥ 4 accident somewhere in the fleet within any five-year window is two-thirds. Three features of this ensemble result bear on the white paper. First, the miscalibration is systematic rather than anecdotal: each accident in the set occurred at a plant whose safety case had classified the event as not credible, so the discrepancy indicts the *generator* of the likelihood assignments, not any individual assignment—and Examples B through D propose to lean more heavily on that generator precisely where it is least constrained by data. Second, the accidents span eight design families—sodium fast breeder, Magnox, UNGG, RBMK, HWGCR, PWR, VVER, BWR—so the premise that a particular design's characteristics remove the accident class, which is the load-bearing premise of Example D's bounding argument for unproven designs, has failed in every family that accumulated enough exposure to test it. Third, the correction-channel failure recurs at the ensemble level: their Biblis example—a serious 1987 precursor at a West German plant that became known internationally only through a trade-press article a year later—is Davis–Besse again, and their concluding recommendation is a mandatory, harmonized, public international event record. A correction channel needs a ledger; the performance-verification infrastructure of Section 7 is that ledger, kept as a condition of the license rather than reconstructed by researchers afterward. The caveats run in both directions: the record is small, INES boundaries are fuzzy and pre-1991 events are back-rated, and the process is non-stationary—the same declining-frequency trend of Section 4 that argues the modern fleet is better than its history. But non-stationarity is an argument against static bins, not for them; it is the reason a likelihood statement must be a standing prediction under a filtration rather than a licensing-time constant.

What the backchecks spare. The backchecks discriminate within the NEI proposal rather than condemning it wholesale. The static frequency bins failed at both TMI and Fukushima; the bounding row fared differently at the two sites, and the difference is instructive. At TMI the containment—a deterministic barrier whose capacity owed nothing to a frequency argument—held, and offsite dose was small. At Fukushima the nominally bounding protection, the seawall, was sized by the same misread frequency tail it was meant to bound, and failed with the bins. A cap is Weinberg's cap only if it is deterministic; a "bounding" event derived from a frequency argument is a bin by another name. The element of Table 12-1 that survives its own backcheck, the MHA/MCA cap, survives exactly where the cap rests on a deterministic bound—which is what this article's total-energy formulation of Section 5 supplies. The implication for Examples B through D runs opposite to the white paper's ordering: it is precisely where operating experience is scarce that licensing-time judgment must be backed by *more* in-service measurement, not less, and by a working channel through which fleet evidence corrects the licensing basis—the channel whose absence both TMI and Fukushima demonstrated, whose statutory form is the Part 21/corrective-action coupling treated in [14], and whose quantitative form is the performance verification of Section 7. Read this way, the white paper's Examples B through D are acceptable as licensing-time structure if and only if they are paired with in-service verification of the asserted likelihood classes. That pairing is what the ARISE SRE supplies, and it is the natural content of the topical-report engagement contemplated in Section 9.

Appendix B. PRA practice and its critics, WASH-1400 to the present: a history read through the IPE turn

This appendix is adapted from the addendum to the author's companion literature review of single-item preventive maintenance for passive devices [13], revised to the framing of this article. The body of this article argues that quantification should be trusted where the physics can support it and that decision authority should rest with the margin (Section 5); Appendix Appendix A backchecks the practice's likelihood claims against the accident record. This appendix supplies the history behind both positions: how the practice of PRA reached its present shape, what its critics said at each stage, and why—in the author's reading—the decisive event was not WASH-1400 itself but the Individual Plant Examination (IPE) program that followed Three Mile Island by a decade.

The Rasmussen report and its review. WASH-1400 [23] assembled, for the first time at full scope, the event-tree/fault-tree architecture of modern PRA, and its enduring engineering contribution was a redirection of attention: severe-accident risk was dominated not by the large-break design-basis event but by small breaks and transients compounded by system and operator failures. Criticism arrived almost immediately, and its target is worth stating precisely. The American Physical Society study group questioned the consequence modeling [17]; the Risk Assessment Review Group (the Lewis Committee) found the framework sound but the uncertainty statements indefensible—error bounds understated, the executive summary unsupported by the body of the report [18]—and in January 1979 the Commission withdrew its endorsement of that executive summary. The critics, that is, accepted the event logic and rejected the confidence claimed for the numbers: the same division this article maintains, in which the event-logic structure is sound engineering practice and the assumption hierarchy beneath the computed frequencies is what Section 3 relaxes and the backchecks of Appendix Appendix A test.

Codification: the risk triplet. Between the Lewis review and the IPE program, the practice acquired its canonical definition. Kaplan and Garrick codified risk as a set of triplets—what can go wrong, how likely is it, what are the consequences [11]—and the formulation was absorbed into USNRC usage so thoroughly that the staff's present working definition of a PRA ("scenarios that result in undesired consequences... and their frequencies" [39]) is the triplet nearly verbatim, four decades on. Two features of the original deserve the historian's attention, because both were shed in the absorption. First, Kaplan and Garrick were candid about the enumeration problem: the scenario list must be closed by a residual category—an "everything else" whose likelihood the analysis cannot supply from within—so completeness is an assumption of the method, not a result of it. Second, "how likely is it" presupposes that the question has a well-posed answer for the system at hand; for rare events in aging, maintained, decision-coupled systems, that presupposition is exactly what the missing limit theorem of Section 3 denies. The practice adopted the triplet without the candor: the enumerated set became identified with the possible set—Hansson's tuxedo fallacy [8] operationalized—and the second question was answered from fleet averages whatever the condition of the item asked about.

Three Mile Island and the two inquiries. Two months after the withdrawal, the TMI-2 accident realized a sequence of exactly the class WASH-1400 had flagged: a small loss of coolant, equipment misbehavior, and operator action compounding it—the sequence Appendix Appendix A reads as a backcheck against qualitative likelihood classification. The two post-accident inquiries—the Kemeny Commission [16] and the Rogovin Special Inquiry [24]—indicted the compliance-focused, design-basis mindset of both industry and regulator; and it must be recorded accurately that both recommended *greater* use of quantitative risk analysis, Rogovin explicitly so. Rogovin's criticism was aimed at engineering practice, not at PRA; his remedy included PRA. The irony that drives the rest of this history is that the inquiries demanding engineering introspection thereby placed the quantification instrument within reach of policy.

The safety goal and the ethical gap. The quantitative anchor of the structure that followed deserves its own scrutiny. The safety-goal policy statement [31], developed through post-TMI workshops and Commission deliberation and finalized in 1986, set two qualitative goals and two quantitative health objectives: the risks of prompt fatality and of latent cancer fatality to individuals near a plant should not exceed one-tenth of one percent of the corresponding background risks. The 0.1% anchor was not derived—not from data, not from theory, not from any articulated principle of acceptable imposition; it was chosen, in ad hoc deliberation, as a judgment of what would not constitute a "significant addition" to the risks people already bear, and the statement offers it as policy judgment, not as result. Two critiques, one written before the goal and one after, bracket what is missing. Weinberg had already named the class: questions that can be asked of science but cannot be answered by it—his examples include the probabilities of extremely improbable events—are *trans-scientific*, and "how safe is safe enough" is among them [42]. Hansson supplies the ethical half: no accumulation of frequencies and expectations can, by itself, settle what it is permissible to impose on whom—consent, distribution, and agency do not appear in the arithmetic [9]. The 10^{-4} benchmark that now anchors the decision arithmetic has a traceable pedigree and no derivation. It entered through the USNRC advisory committee's quantitative safety-goal deliberations [1]; the Commission then "selected" it in the trial-use policy statement as a design objective—"the likelihood of a nuclear reactor accident that results in a large-scale core melt should normally be less than one in 10,000 per year of reactor operation" [29]; the 1986 final statement dropped the core-melt objective (retaining the health objectives); and the number re-entered through staff practice as the subsidiary core-damage-frequency objective, joined by the 10^{-5} large-early-release surrogate, on which the decision arithmetic of RG 1.174 [39] is built. It is the same number that seeds the frequency boundaries of the qualitative likelihood bins and F–C targets in the licensing approaches read in Appendix Appendix A. The trial-use statement was candid about the ground the number stood on: the Commission conceded "substantial associated uncertainties" and "a serious question whether, for a specific nuclear power plant, the achievement of a regulatory-imposed quantitative risk goal can be verified with a sufficient degree of confidence" [29]. The number survived; the caveat did not.

The IPE turn. The Commission's 1985 Severe Accident Policy Statement concluded that existing plants posed no undue risk and that no generic rulemaking was warranted, while recognizing that "systematic examinations are beneficial in identifying plant-specific vulnerabilities to severe accidents that could be fixed with low cost improvements" [30]. Generic Letter 88-20 then requested, under 10 CFR 50.54(f), that each existing plant perform a systematic examination to identify plant-specific vulnerabilities and report the results [32]. Its stated purposes were an *appreciation* of severe accident behavior, an understanding of the most likely sequences, a *more quantitative* understanding of the probabilities, and—only then, "if necessary"—hardware and procedural fixes. The examination was the point; the quantification was a means. The generic letter offered three acceptable vehicles: a Level 1 (or better) PRA, the industry examination method with staff-identified enhancements, or "other systematic examination methods" subject to case-by-case staff acceptance—a menu that made PRA the path of least regulatory resistance, and nearly every licensee took it. The staff's synthesis of the program [36] records the outcome: the definition of "vulnerability" was left to licensees and varied widely, few licensees declared any, and the program nonetheless

produced hundreds of plant improvements—evidence that the *examination* did useful engineering work even where the quantification declared nothing actionable. The permanent legacy, however, was neither the improvements nor a habit of reinvestigation: it was a fleet-wide inventory of plant-specific PRA models. (It is worth noting that the "other systematic examination methods" clause of 1988 is the direct ancestor of the SRE clause of the final Part 53 rule on which this article's licensing pathway rests.)

Institutionalization, and the inversion. With the models in hand, their employment was institutionalized, and the Federal Register dates the turn precisely: the safety goals (51 FR 30028, 1986) preceded the IPE request; the Maintenance Rule arrived while the IPE submittals were still being written (56 FR 31306, 1991); NUREG-1150 refreshed the technical basis [33]; and then, with the fleet's models complete, the 1995 policy statement directed that PRA methods be used in all regulatory matters "to the extent supported by the state of the art" [35], RG 1.174 gave the change-evaluation framework in 1998 [39], and risk-informed categorization out of special treatment was codified in 10 CFR 50.69 [37]. The application stream that resulted is dominated not by vulnerability discovery but by burden relief—extended allowed outage times, extended surveillance and test intervals, categorization out of special treatment. Here is the inversion this appendix exists to name, stated as the author's position: an instrument commissioned to find unexpected vulnerabilities became the instrument of relief from the provisions that traditional engineering had erected against them, and traditional conservative safety engineering—margins, defense in depth, prescriptive surveillance—was recast in the process as "unquantified conservatism," a cost to be harvested rather than a defense to be maintained. The community registered the tension at the time: the structuralist-versus-rationalist debate over defense in depth [25] is precisely an argument over whether the traditional structure has standing independent of the risk numbers.

The data layer. Institutionalization also fixed where the numbers come from. The basic-event probabilities of the fleet's PRAs are drawn from industry-average parameter estimates maintained for the USNRC at the Idaho National Laboratory [4, 10]: event counts pooled across the operating fleet, divided by an aggregated exposure. As the answer to the triplet's second question for a specific, repairable, maintained item, this is on its face the wrong object, for the reason Ascher and Feingold made canonical [2]: pooling the failures of repairable systems and dividing by fleet exposure produces an average rate of occurrence of failures, not a hazard and not the intensity of any one item, and the ratio's denominator belongs to no item in particular and carries none of the maintenance, state, and environmental history that governs the item in front of the engineer [45, 15]. The estimates are unconditional by construction: the same number before and after an inspection, for the well-maintained item and the neglected one. The distinction that matters for this article—which calibrates its own penetration simulation from the same INL compilations (Section 6)—is between the two uses of an average: supplied as the *final, unconditional answer*, which is the practice criticized here, and supplied as the *prior* inside a filtration that updates on the item's own inspection and performance record, which is what the MPP construction does and what the oversight program of Section 7 verifies in service. The average is where inference starts, not where it ends.

The legislative terminus. The history now runs through Congress. NEIMA [26] directed the USNRC to establish a technology-inclusive, risk-informed, performance-based licensing framework—the mandate behind Part 53—and ADVANCE [27] directed the agency to revise its mission statement itself, so that licensing and regulation be "efficient" and not unnecessarily limit the civilian use of nuclear energy. Read plainly, these are efficiency and fee-discipline statutes, and their supporters present them so. Read against the history above—and this is the author's position—they are its terminus: the relaxation dynamic, having migrated from application practice into guidance and rule, is now installed as a statutory instruction to the regulator's mission, applied in the name of a "modern risk understanding" whose theoretical underpinnings are undeclared [47, 45, 15]. The engineering objection is not to efficiency. It is to the premise underneath: that quantified risk is the senior measure of conservatism, fit to serve as a design target. No properly trained engineer designs *to* a probability, because a probability is not a coordinate of design space: design proceeds in load, geometry, and material variables, where the engineer commands physics, and a probability is a property of an ensemble of modeled outcomes reached from those variables only through a model whose assumptions the design itself cannot secure. Where probabilistic reasoning legitimately informs design—the calibration of load and resistance factors in modern structural codes—it is exercised offline, by code committees, and the instrument handed to the designer is deterministic: factors and margins. This reading sharpens rather than undercuts the use this article makes of the same statutes. What the statutes install is an instruction for efficiency; what they do not settle is the quantification by which safety is demonstrated. The SRE clause of the final rule (Section 2) is the opening through which a quantification with a stated theoretical basis can enter, and this article's claim on efficiency runs through reduced regulated scope and process (Sections 4 and 8), not through relief from protections justified by the frequency machinery this appendix has traced. The margin condition (4) and the chance constraint (2) are the deterministic-and-verifiable instruments that the design-space discipline of factors and margins demands.

Why only reinvestigation finds the unexpected. The redirection mattered because of what a PRA can and cannot see. A PRA quantifies the event space it enumerates; an *unexpected* vulnerability is, by definition, outside the enumeration—the no-unknown-unknowns assumption of the hierarchy in Section 3, operating at the level of program design rather than of any one calculation. Generic Letter 88-20's own ordering of purposes shows the drafters understood this: the appreciation and the understanding come first, the numbers are instrumental, and the fixes are conditional. The record since bears it out: the discoveries that mattered came from engineering reinvestigation and operating experience—the Davis–Besse vessel-head degradation of 2002 from inspection, the GSI-191 debris phenomenology from test programs and physics simulation [12], the refueling-water-storage-tank degradation of IN 2013-18 from walkdowns finding boric-acid residue [38]—not from the risk models, whose enumerated spaces did not contain them. And the backchecks of Appendix Appendix A show the same channel failing in the other direction: evidence in hand (the relief-valve record, the Jōgan deposits, the Biblis precursor) unable to act on the models' claims. None of this is an argument against quantification. It is an argument about *which* quantification and *what for*: quantify the physics of the system in front of you (the MPP of Section 3), let the margin carry the decision (Section 5), and reserve for periodic engineering reinvestigation—in this framework, the performance oversight of Section 7 and the filtration's incorporation of each discovery (Section 11)—the one task no model performs: looking for what the model does not contain.

References

[1] Advisory Committee on Reactor Safeguards. An approach to quantitative safety goals for nuclear power plants. Technical Report NUREG-0739, U.S. Nuclear Regulatory Commission, 1980.

[2] Harold Ascher and Harry Feingold. *Repairable Systems Reliability: Modeling, Inference, Misconceptions and Their Causes*. Marcel Dekker, 1984.

[3] Pierre Brémaud. *Point Processes and Queues: Martingale Dynamics*. Springer-Verlag, 1981.

[4] S. A. Eide, T. E. Wierman, C. D. Gentillon, D. M. Rasmuson, and C. L. Atwood. Industry-average performance for components and initiating events at U.S. commercial nuclear power plants. Technical Report NUREG/CR-6928, U.S. Nuclear Regulatory Commission, 2007.

[5] Electric Power Research Institute. Risk impact assessment of extended integrated leak rate testing intervals. Technical Report TR-1009325, Rev. 2-A, EPRI, 2008. Basis for the NEI 94-01 ILRT interval methodology.

[6] Executive Office of the President. Ordering the reform of the nuclear regulatory commission. Executive Order, 2025. May 23, 2025.

[7] Minh Ha-Duong and V. Journé. Calculating nuclear accident probabilities from empirical frequencies. *Environment Systems and Decisions*, 34:249–258, 2014.

[8] Sven Ove Hansson. From the casino to the jungle: Dealing with uncertainty in technological risk management. *Synthese*, 168(3): 423–432, 2009.

[9] Sven Ove Hansson. *The Ethics of Risk: Ethical Analysis in an Uncertain World*. Palgrave Macmillan, 2013.

[10] Idaho National Laboratory. Industry-average performance for components and initiating events at U.S. commercial nuclear power plants: 2020 update. Technical Report INL/EXT-21-65055, Idaho National Laboratory, 2021.

[11] Stanley Kaplan and B. John Garrick. On the quantitative definition of risk. *Risk Analysis*, 1(1):11–27, 1981.

[12] E. Kee, J. Hasenbein, A. Zolan, P. Grissom, S. Reihani, Z. Mohaghegh, F. Yilmaz, B. Letellier, V. Moiseytseva, R. Vaghetto, D. Imbaratto, and T. Sakurahara. RoverD: Use of test data in GSI-191 risk assessment. *Nuclear Technology*, 196(2):270–291, 2016.

[13] Ernest J. Kee. Single-item preventive maintenance optimization for passive devices: A literature review and a point-process, margin-constrained reframing. Working draft, YKrisk, LLC, 2026.

[14] Ernie Kee. Revisiting common cause failure in probabilistic risk assessment: A time-gated interpretation. Working paper, YKrisk, LLC, 2026.

[15] Ernie Kee and M. A. Wortman. Perspective chapter: PRA and protective system maintenance. In P. Tsvetkov, editor, *Nuclear Fission—From Fundamentals to Applications*. IntechOpen, 2023. doi:10.5772/intechopen.110049.

[16] John G. Kemeny et al. Report of the President's commission on the accident at Three Mile Island: The need for change: The legacy of TMI. Technical report, U.S. Government Printing Office, Washington, DC, 1979.

[17] H. W. Lewis et al. Report to the American Physical Society by the study group on light-water reactor safety. *Reviews of Modern Physics*, 47, 1975. Suppl. 1.

[18] H. W. Lewis et al. Risk assessment review group report to the U.S. nuclear regulatory commission. Technical Report NUREG/CR-0400, U.S. Nuclear Regulatory Commission, 1978.

[19] K. Minoura, F. Imamura, D. Sugawara, Y. Kono, and T. Iwashita. The 869 Jōgan tsunami deposit and recurrence interval of large-scale tsunami on the Pacific coast of northeast Japan. *Journal of Natural Disaster Science*, 23(2):83–88, 2001.

[20] National Diet of Japan. The official report of the Fukushima nuclear accident independent investigation commission (NAIIC). Technical report, National Diet of Japan, Tokyo, 2012. Kurokawa Commission.

[21] Nuclear Energy Institute. Risk-informed performance-based technology inclusive guidance for non-light water reactor licensing basis development. Technical Report NEI 18-04, Rev. 1, NEI, 2019.

[22] Nuclear Energy Institute. Technology-inclusive, risk-informed, performance-based approaches for development of licensing bases under Part 53. White paper, NEI, 2025. Enclosure 1 to NEI comments on the proposed 10 CFR Part 53 rule.

[23] Norman C. Rasmussen et al. Reactor safety study: An assessment of accident risks in U.S. commercial nuclear power plants. Technical Report WASH-1400 (NUREG-75/014), U.S. Nuclear Regulatory Commission, 1975.

[24] Mitchell Rogovin and George T. Frampton. Three mile island: A report to the commissioners and to the public. Technical Report NUREG/CR-1250, NRC Special Inquiry Group, 1980.

[25] J. N. Sorensen, G. E. Apostolakis, T. S. Kress, and D. A. Powers. On the role of defense in depth in risk-informed regulation. In *Proc. Int. Topical Meeting on Probabilistic Safety Assessment (PSA '99), Washington, DC*. American Nuclear Society, 1999.

[26] U.S. Congress. Nuclear energy innovation and modernization act, public law 115-439, 2019.

[27] U.S. Congress. Accelerating deployment of versatile, advanced nuclear for clean energy (ADVANCE) act of 2024, public law 118-67, title ii, 2024. Signed July 9, 2024.

[28] U.S. Department of Energy. DOE reactor pilot program: Initial selections, 2025. Eleven advanced designs selected for the DOE authorization pathway; target of three criticalities by July 4, 2026.

[29] U.S. Nuclear Regulatory Commission. Safety goals for nuclear power plant operation. Technical Report NUREG-0880, Rev. 1 (policy statement adopted for trial use), U.S. Nuclear Regulatory

Commission, 1983.

[30] U.S. Nuclear Regulatory Commission. Policy statement on severe reactor accidents regarding future designs and existing plants. 50 FR 32138, 1985.

[31] U.S. Nuclear Regulatory Commission. Safety goals for the operations of nuclear power plants; policy statement. 51 FR 30028, 1986.

[32] U.S. Nuclear Regulatory Commission. Individual plant examination for severe accident vulnerabilities—10 CFR 50.54(f). Generic Letter 88-20, 1988.

[33] U.S. Nuclear Regulatory Commission. Severe accident risks: An assessment for five U.S. nuclear power plants. Technical Report NUREG-1150, U.S. Nuclear Regulatory Commission, 1990.

[34] U.S. Nuclear Regulatory Commission. Performance-based containment leak-test program. Technical Report NUREG-1493, U.S. Nuclear Regulatory Commission, 1995.

[35] U.S. Nuclear Regulatory Commission. Use of probabilistic risk assessment methods in nuclear regulatory activities; final policy statement. 60 FR 42622, 1995.

[36] U.S. Nuclear Regulatory Commission. Individual plant examination program: Perspectives on reactor safety and plant performance. Technical Report NUREG-1560, U.S. Nuclear Regulatory Commission, 1997.

[37] U.S. Nuclear Regulatory Commission. Risk-informed categorization and treatment of structures, systems and components for nuclear power reactors. 10 CFR 50.69; 69 FR 68008, 2004.

[38] U.S. Nuclear Regulatory Commission. Refueling water storage tank degradation. Information Notice 2013-18, ADAMS ML13128A118, 2013.

[39] U.S. Nuclear Regulatory Commission. An approach for using probabilistic risk assessment in risk-informed decisions on plant-specific changes to the licensing basis. Technical Report Regulatory Guide 1.174, Rev. 3, U.S. Nuclear Regulatory Commission, 2018.

[40] U.S. Nuclear Regulatory Commission. Acceptability of probabilistic risk assessment results for non-light-water reactor risk-informed activities. Technical Report Regulatory Guide 1.247, U.S. Nuclear Regulatory Commission, 2022.

[41] U.S. Nuclear Regulatory Commission. Risk-informed, technology-inclusive regulatory framework for advanced reactors; final rule (10 CFR Part 53). Federal Register, 2026. Approved March 25, 2026; effective April 29, 2026.

[42] Alvin M. Weinberg. Science and trans-science. *Minerva*, 10: 209–222, 1972.

[43] Alvin M. Weinberg. The first nuclear era: The life and times of a technological fixer. *Physics Today*, 48(11), 1995. See also A. M. Weinberg, The First Nuclear Era, AIP Press, 1994.

[44] Ronald W. Wolff. Poisson arrivals see time averages. *Operations Research*, 30(2):223–231, 1982.

[45] M. A. Wortman, E. J. Kee, and P. Kannan. Is core damage frequency an informative risk metric? arXiv:2105.09798, 2021.

[46] Martin Wortman, Ernest Kee, and Pranav Kannan. A note on probability quantification for protective system efficacy analysis: Stochastic dynamics, information flow, and initiating event arrival times. *arXiv preprint arXiv:2203.04315*, 2022.

[47] G. Zheng, P. Nelson, E. Kee, F. Yilmaz, Z. Zhang, and M. Wortman. Whether a theory of nuclear safety is deterministic or probabilistic? A metatheory. *Annals of Nuclear Energy*, 110: 979–988, 2017.