Balancing Core Damage Frequency and Containment Robustness in Modern Safety Cases

Joshua A. Tolbert, Ph.D., P.E.

About the Author

Joshua A. Tolbert, Ph.D., P.E., is a seasoned engineer with four engineering degrees, including a doctorate in mechanical engineering. He has extensive experience in power generation, including coal, natural gas, biomass, and nuclear, as well as leadership roles such as CTO and VP of Engineering for a commercial developer of small-scale nuclear power plants. Dr. Tolbert has worked with nuclear regulators domestically and internationally, including the UK ONR, Poland's PAA, Argentina's ARN, Romania's CNCAN, and Turkey's NDK.

In the private sector, Dr. Tolbert has identified the major obstacles to nuclear power development, including regulatory, technical, and economic challenges. He has a deep commitment to participating in the revitalization of the nuclear industry by implementing practical solutions to these obstacles.

His technical expertise includes the full design of a small-scale pressurized water reactor (PWR), encompassing the primary loop, secondary loop, and auxiliary systems, all utilizing advanced modularization techniques to enhance constructability and reduce costs. This hands-on experience positions Dr. Tolbert to offer unique insights into the challenges and opportunities facing the nuclear industry today.

Introduction

Modern nuclear safety analysis sits at the intersection of two powerful but often competing philosophies: preventing accidents from occurring and mitigating the consequences when they do. Over the last four decades, the U.S. regulatory framework has leaned heavily toward the first of these goals, treating Core Damage Frequency (CDF) as the premier metric of reactor safety and calibrating licensing expectations accordingly. Meanwhile, containment—the physical system historically designed to protect the public even in the unlikely event of severe accidents—has remained subject to strict, deterministic design requirements that operate largely independent of probabilistic considerations.

This dual structure has created an imbalance. On the one hand, developers must demonstrate extraordinarily low CDF values, often well below 1×10^{-5} per reactor-year for new designs. On the other, they must build and analyze highly robust containment structures capable of withstanding extreme low-probability events. The combined effect is a regulatory burden that inflates cost, extends licensing timelines, and distorts engineering priorities—without necessarily delivering commensurate improvements in public safety.

This paper examines how the nuclear sector arrived at this imbalance, why the emphasis on ultra-low CDF creates diminishing returns, and how a recalibrated approach—one that restores containment to its rightful place in a layered safety strategy—can reduce unnecessary barriers while keeping safety margins strong.

The Historical Rise of Core Damage Frequency as a Safety Metric

The transition to probabilistic risk assessment (PRA) in the 1970s and 1980s was a landmark achievement for nuclear safety. CDF gave regulators and analysts a way to quantify accident prevention by integrating equipment reliability, operator performance, system interactions, and external hazards into a coherent model. As PRA methods matured, CDF became the principal quantitative expression of overall safety.

Initially, the emphasis on lowering CDF made sense. Early fleet reactors benefited from PRA insights that identified vulnerabilities and improved equipment reliability and redundancy. But over time, regulatory expectations tightened far beyond what empirical risk data—or economic practicality—could justify. The nuclear industry now operates under an implicit belief that lower CDF is always better, regardless of whether reductions represent meaningful improvements to public safety or simply reflect diminishing returns that drive up cost.

Compounding this issue is a regulatory structure that simultaneously maintains highly prescriptive containment standards. Instead of allowing containment robustness to offset some of the burden on accident prevention, the two requirements grew in tandem, each becoming more restrictive without being integrated into a unified risk-informed framework.

Why Ultra-Low CDF Targets Are Economically Distorting

Achieving a low CDF is desirable. But achieving *extremely* low CDF—1×10^{-5} or lower—requires increasingly complex systems, extensive redundancy, and costly equipment qualification. These improvements often target scenarios that, even before mitigation, are vanishingly improbable.

Importantly, CDF represents more than a safety metric—it represents catastrophic economic loss. A core damage event destroys the reactor, contaminates equipment, halts power generation, and imposes multi-billion-dollar cleanup costs. From a purely financial perspective, plant operators have overwhelming incentive to avoid core damage. Market forces, therefore, naturally drive CDF downward even without hyper-prescriptive regulation.

When regulators pursue ultra-low CDF values beyond what market forces already discourage, they push designs toward diminishing returns—returns that ultimately manifest as higher capital cost, longer construction timelines, and reduced competitiveness relative to other energy sources. In effect, the system incentivizes engineering solutions far beyond what risk justifies, transferring cost onto developers while creating little additional public benefit.

The Underappreciated Role of Containment

Containment was originally conceived as the final and most important barrier against radiological release, a robust physical structure capable of preventing harm even in the event of serious accidents. Historically, containment performance was judged through deterministic analyses and conservative assumptions. PRA later introduced Level 2 analysis, linking core damage scenarios to release frequencies and exploring containment's performance across different accident classes.

Yet despite containment's proven ability to prevent significant off-site consequences, regulatory emphasis has shifted overwhelmingly toward preventing core damage itself rather than mitigating its consequences. The result is a safety paradigm where containment is expected to handle near-impossible loads while PRA drives CDF targets steadily lower—creating an unnecessary and expensive duplication of safety functions.

A more rational approach recognizes that very low public risk can be achieved through the combination of:

1. reasonably low CDF, driven by sound engineering and economic incentives, and
2. a containment system designed to reliably limit radiological release under credible accidents.

Such a framework aligns more closely with the underlying principles of defense-in-depth and with the global experience of operating reactors safely for decades.

Toward a More Balanced Safety Framework

A recalibrated approach does not abandon PRA or discard the value of CDF. Instead, it integrates prevention and mitigation into a coherent safety philosophy—something the NRC's proposed Part 53 framework moves toward but does not yet fully achieve in practice.

In a balanced model:

1. PRA Level 1 (prevention) and Level 2 (mitigation) are treated as complementary, not competing.
2. Containment design is recognized for its significant contribution to public safety and is allowed to offset some of the burden on accident prevention.
3. Regulatory requirements prioritize outcomes—public health protection and environmental integrity—over rigid numeric thresholds that may not materially affect those outcomes.
4. Designers gain the flexibility to innovate holistically, optimizing plant architectures across both prevention and mitigation measures.

This approach encourages diversity in reactor designs, supports realistic tradeoffs, and maintains a safety margin consistent with historical performance.

The Benefits of Rebalancing Prevention and Mitigation

When containment and CDF are allowed to work together rather than against each other, several benefits emerge:

Licensing becomes more efficient.
A unified framework reduces the regulatory friction caused by treating safety requirements for prevention and mitigation as independent constraints rather than integrated factors.

Design becomes more flexible.
Developers can explore alternative systems architectures, plant sizes, and component

classifications without being forced into prohibitively expensive solutions solely to achieve ultra-low CDF targets.

Costs decrease without compromising safety.
Recognizing containment's value allows designers to redirect resources away from marginal CDF reductions and toward overall plant performance, reliability, and constructability.

Innovation becomes more accessible.
New reactor developers—especially small modular reactors and advanced non-LWR technologies—face fewer artificial barriers and can focus on solutions that genuinely improve public safety.

Conclusion

The nuclear industry's regulatory framework has spent decades tightening requirements for both core damage prevention and containment robustness without fully considering how these two domains interact. This dual escalation has produced designs that are technically impressive but economically prohibitive, limiting the industry's ability to build new reactors.

A recalibrated safety approach recognizes that core damage prevention and containment are not independent silos; they are two essential components of a single safety philosophy. By integrating them into a balanced, risk-informed model, the nuclear industry can reduce unnecessary regulatory burdens, restore economic viability, and support the next era of nuclear deployment—without compromising the high safety standards that define the field.

This balanced perspective is not merely a regulatory improvement. It is a prerequisite for the future of nuclear power and a necessary step toward rebuilding the industry's capacity to deliver clean, reliable energy at the scale the modern world demands.