An Expected Value Approach to “As Safe as Reasonably Practicable”

Chris Everett
Idaho National Laboratory

18th International Conference on Probabilistic Safety Assessment and Management
PSAM-18

July 2026

Introduction

- **Being as safe as reasonably practicable (ASARP) is a fundamental principle of adequate safety at NASA**
  - A system is ASARP if an incremental improvement in safety would require a disproportionate deterioration of system performance in other areas
  - ASARP is based on the idea that life is precious, and therefore safety should not be traded away in pursuit of other objectives unless absolutely necessary
- **ASARP is at odds with standard decision analysis, which assigns a definite value to each possible outcome of a decision, and identifies the alternative that maximizes expected value**
  - Being ASARP doesn't enter into the analysis
- **The method presented here charts a middle path by defining ASARP in a way that:**
  - **Is conceptually faithful to the ASARP principle**
  - **Works within an expected-value framework**

Being ASARP is a Fundamental Principle of Adequate Safety at NASA

- **Adequate safety: 1) Meeting minimum tolerable levels of safety, and 2) Being ASARP**
- **ASARP entails weighing safety performance against the sacrifice needed to further improve it**
- **A system is ASARP if an incremental improvement in safety would require a disproportionate deterioration of system performance in other areas**

System meets minimum tolerable level of safety?

System is ASARP?	Yes	No
Yes	System is adequately safe	System is inherently unsafe
No	System is sub-optimally safe as designed	System is unsafe as designed

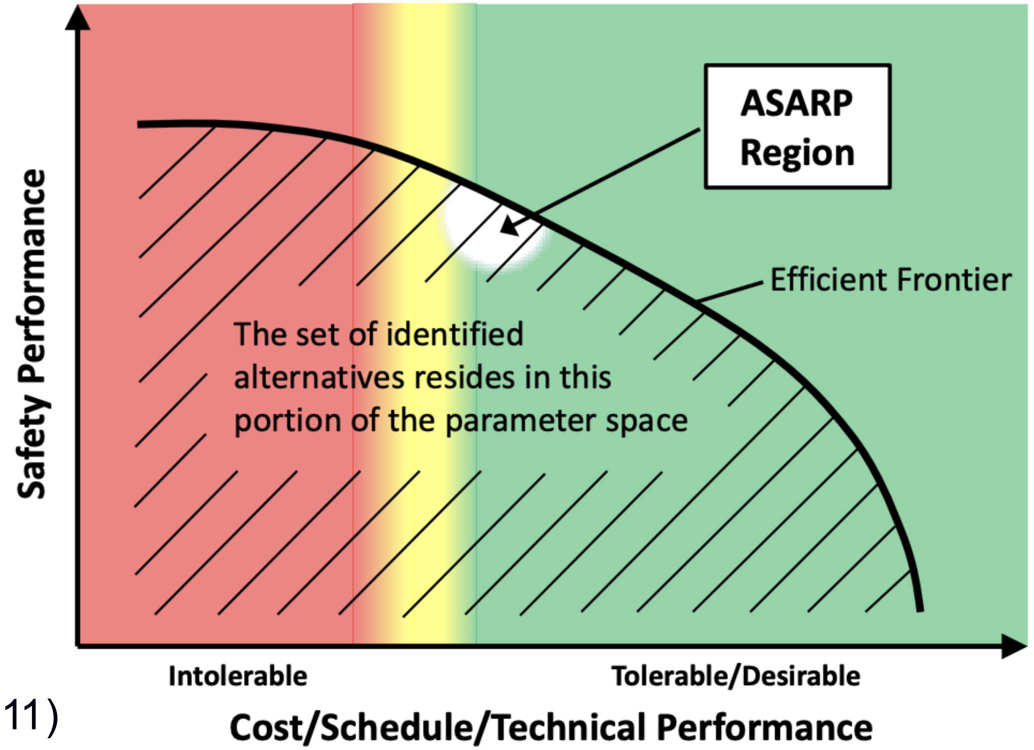

Figures are from the NASA System Safety Handbook, Volume 1 (2011)

Pushback from the Decision Analysis Crowd

- **Standard decision analysis does not involve the concept of ASARP (or even limits of tolerable safety)**
- **Instead, given the “dollarization” of all relevant outcomes and the determination of all relevant probabilities, the best course of action is simply the one that maximizes the expected value E[V]**

For a simple system where failure results in loss of crew and mission:

$$\mathbf{E[V] = P_S R - C - (1 - P_S) L}$$

Where:

E[V] = Expected value
P_S = Probability of success
R = Value of the mission return
C = System cost
L = Value of lives put at risk

So:

$$\mathbf{P_S = [1/(R + L)]C + [(E[V] + L)/(R + L)]}$$

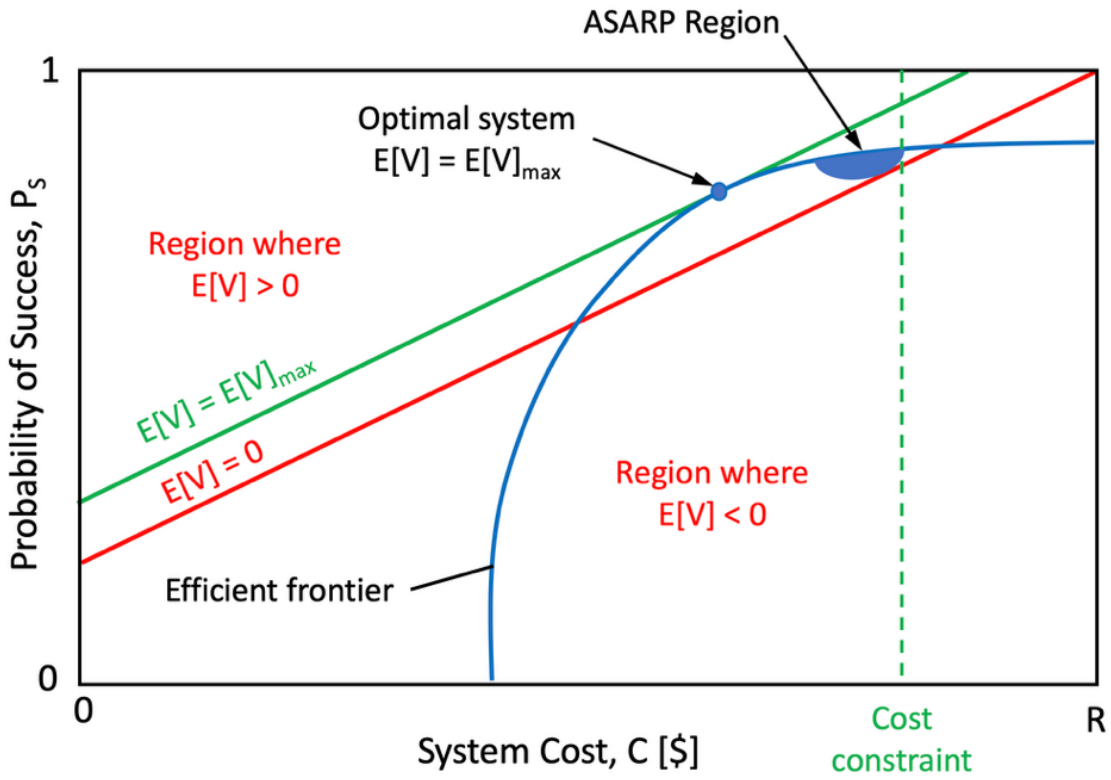

Issues with Setting a Value for Lives Put at Risk

- **The expected value E[V] of each alternative is sensitive to the value L placed on the lives put at risk**
  - Higher values of L shift the optimal solution up and to the right on the efficient frontier
  - Lower values of L shift the optimal solution down and to the left on the efficient frontier
- **This renders the optimal alternative vulnerable to disagreements about the value of L**
  - For reference, the value per statistical life used by HHS is between \$6.2 million and \$20.3 million, with a central estimate of \$13.4 million
- **Assigning *any* specific value L to the lives put at risk is antithetical to NASA's philosophy of putting astronaut safety first** (within the context of actually conducting the mission)

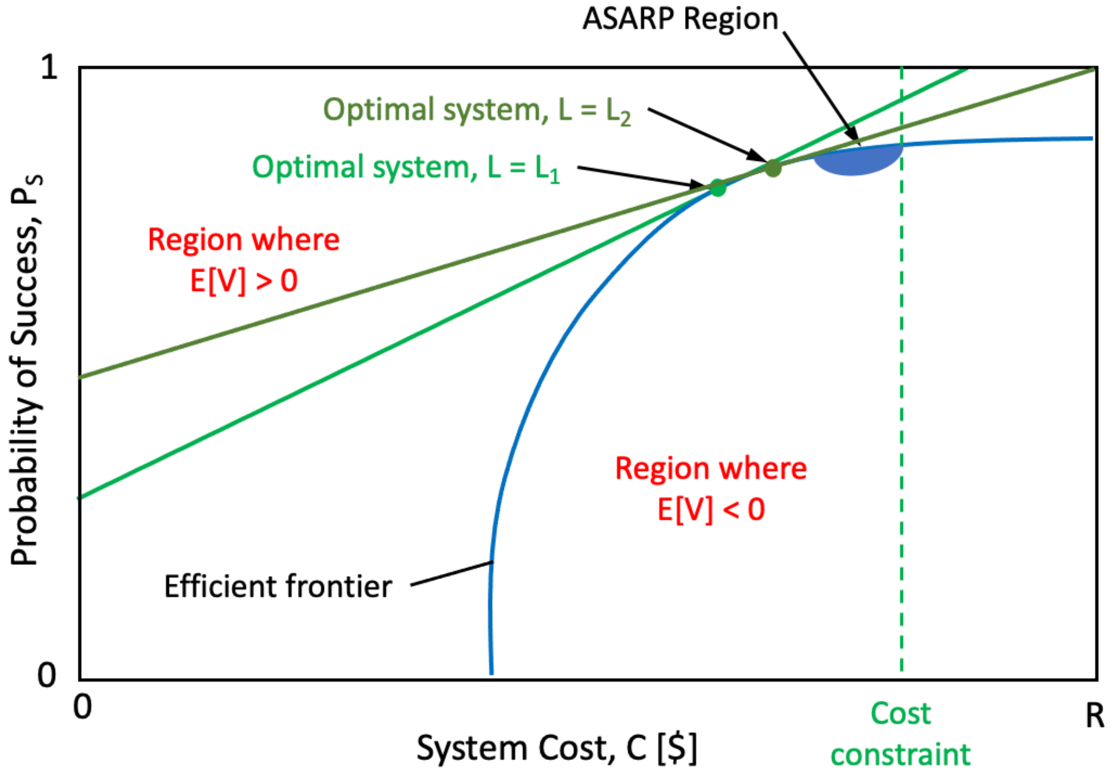

Squaring the Circle: Expected-Value-Based ASARP

- **For each alternative *i* there is a value L_i above which the risk of fatalities outweighs the potential benefit of the mission**
  - L_i is found by setting $E[V]_i = 0$ and solving for L_i: $\mathbf{L_i = (P_{S\_i}R - C_i)/(1 - P_{S\_i})}$
- **Given this, it is possible to find the alternative with the highest such valuation $L_{Max} \geq L_i$ for all i**
- **This alternative can be considered ASARP**
  - It's the only alterantive that can be rationally entrusted with lives of value L_{Max}
    - $E[V] < 0$ for all other alternatives
  - It's also the safest alternative that can be rationally justified at all
- **The value L_{Max} is an artifact of the method, not an assertion of the actual value of L**
- **The method is silent on whether the risk is tolerable (that's a separate issue)**

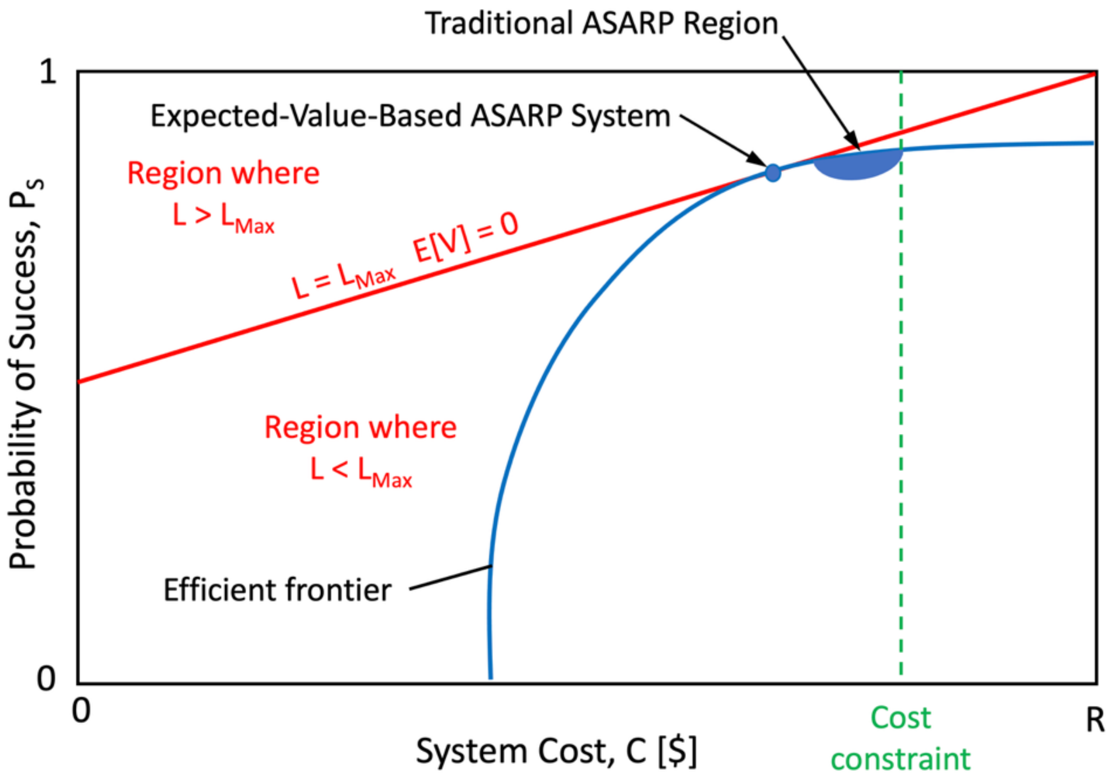

Generalization of the Method

- **When multiple entities of different types are put at risk, it is necessary to assign *relative* values to each type**
  - This allows L_i and L_{Max} to be unambiguously determined as the sums of the values of the underlying entities
- **The analysis can be formulated in terms of expected utility E[U] rather than expected value E[V]**
  - The tipping point between conducting vs. not conducting the mission is the utility U_0 of the "no mission" alternative, rather than zero
  - L_i is found by setting $E[U]_i = U_0$ and solving for L_i
- **Parametric analysis can provide insight where values are uncertain/subjective**
  - Mapping the ASARP alternative over the space of reasonable mission returns
  - Mapping the ASARP alternative over the space of reasonable relative values for the entities put at risk

Conclusion

- **This is a simple and straightforward approach to the ASARP principle that is compatible with standard decision analysis while respecting the philosophy of putting safety first**
- **The method is conceptually true to the meaning of ASARP**
  - The ASARP alternative is the only one that can be entrusted with lives of value L_{Max}
  - The method provides a rigorous mathematical basis for “reasonably practicable,” namely: “can be justified on expected value grounds”
  - In the example, the ASARP alternative is the *safest* alternative that can be so justified
  - Risk tolerability is a separate issue (per the NASA System Safety Handbook)
- **The method is notable for what it doesn’t entail**
  - It doesn’t entail pursuing safety to the edge of tolerable performance (and risk) in other areas
  - It’s non-prescriptive, e.g., it doesn’t mandate the use of best available technology or established best practice (although those should be included in the space of alternatives)
- **The method is timely**
  - ALARA is out and ALARP has been criticized for having “the effect of stepping up measures until they are far beyond the ‘reasonable’ definition,” leading to “excessive nugatory effort and the stifling of innovation” (“What Future for the ALARA Principle?”, NEI. February 17, 2026)

Acknowledgements

The author would like to thank the NASA Office of Safety and Mission Assurance (OSMA) for its sponsorship of this paper through Strategic Partnership Project (SPP) #01717 - Development and Implementation of Risk Management and System Safety Framework at NASA. The views and opinions of the author expressed herein do not necessarily state or reflect those of the U.S. Government or any agency thereof.