

# The Incremental Safety Manifesto

Erik Hollnagel and David Slater

Safety only has one definition, namely: a state where the number of **unexpected and unacceptable outcomes are kept as low as practicable**. But there are different ways to get there. One is to do it directly by trying to reduce the number of unexpected and unacceptable outcomes. This represents a decremental approach to safety. The other is to get there by trying to increase the number of expected and acceptable outcomes. The argument is that what goes well cannot fail at the same time. This represents an incremental approach to safety.

On that basis, *decremental* and *incremental* safety are not competing definitions of safety. They are alternative strategies for how best to get the same result. A decremental approach aims to reduce unwanted outcomes to a minimum (ideally, but unattainably zero) , so that nothing goes wrong. An incremental approach aims to increase wanted outcomes to a maximum, ideally but unattainably 100 percent, so that work always goes well. The practical difference is what each strategy leads an organisation to focus on, to measure, to reinforce, and to learn from.

In a paper published in 2000 professor James Reason described the paradox: *that safety defined and measured more by its absence than by its presence*. And most institutional definitions confirm that. The American National Standards Institute (ANSI) and the American Society of Safety Engineers (ASSE) both define safety as “freedom from unacceptable risk.” The International Civil Aviation Organization (ICAO) adds more detail and defines safety as “the state in which harm to persons or property damage is reduced to, and maintained at or below, an acceptable level through a continuing process of hazard identification and risk management.” All three are defensible formulations. But they all pay attention to harm, damage, injury, and loss—to the manifestations of the **absence** of safety, rather than to the **presence** of safety

Measurement practices go along. When safety is framed as the absence of unexpected and unacceptable outcomes, organisations count the manifestations of harm: lost-time cases, incident rates, severity categories and their derivatives. Improvement is then pursued by any means imaginable: drive the numbers down and safety will improve. But a paradoxical asymmetry follows from this logic. The better (or safer) you get, the less data you have. Work that goes well is recorded as “nothing happened,” which may be comforting, but a poor empirical basis for learning.

A comparison with health is instructive. The WHO defines health as “a state of complete physical, mental, and social well-being and not merely the absence of disease or infirmity.” Health is not only the elimination of illness; it is also the presence of capability, which is why many people dutifully exercise every day. The traditional focus on the absence of safety also creates a terminological problem: we have abundant language for failure and the **absence** of safety—incident, accident, injury—but practically none for the presence of safety, there are few applicable antonyms for accident and incident. One way to avoid this trap is to dispense with conventional labels and instead describe outcomes by what they plainly are: *unexpected and unacceptable* versus *expected and acceptable*. It may well be cumbersome, but it keeps the logic honest, and the mind clean.

This matters because it suggests a different and more effective kind of addressing safety issues and of learning from them. Unexpected outcomes are **disruptive** because they interrupt planned and prescribed activity. Unacceptable outcomes are potentially **harmful** because they carry negative consequences: material, or financial losses, operational disruption, reputational damage, injury, pain, death. But if safety instead is the reliable production of expected and acceptable outcomes, then it can be strengthened not only by suppressing harm, but by improving the conditions under which work goes well. That is the practical distinction between decremental and incremental safety.

A decremental approach focuses on preventing unexpected and unacceptable outcomes. Mainly by identifying actions and conditions associated with harm and adding constraints: barriers, blocks, interlocks, approvals, rules, checks and barriers (physical, functional, symbolic, or incorporeal).

The learning story is familiar: investigate what went wrong, specify what should not have happened, and create obstacles to prevent a recurrence. This is seen as necessary work. But It is not sufficient.

Adding constraints to the workplace does not remove the need to adapt; it rather reshapes it. Work-as-Done is never and cannot ever be identical to Work-as-Imagined. People at both the sharp and the blunt ends inevitably adjust their work to the local conditions—time pressure, resource shortfalls, (un)available competence, competing demands, incomplete information, awkward tools, and so on. When obstacles make work harder, people never just stop. Instead they route around the perceived hindrance so the task can still be done. A purely decremental intervention may therefore increase performance adjustments (and with them, variability) despite the intention to reduce them.

This is where a judgement asymmetry enters. The everyday inventiveness that keeps work on track is rarely analysed because it produces a non-event. It happens all the time, and is therefore not considered interesting, but treated as normal. The same inventiveness becomes salient in hindsight when it is inferred to have contributed to an unwanted outcome, and is then recast as deviation, non-compliance, or human error. The ability to adjust work to the current conditions is praised when the outcomes are acceptable and condemned when they are not. That is not a moral failing of individuals; it is a dysfunctional property of the safety legacy, and of how many organisations operate.

Incremental takes a different empirical stance: in complex work, the routine achievement of expected and acceptable outcomes is not discarded as trivial. It instead becomes the phenomenon that needs explaining. The question changes from “Why did this not go well, or suddenly fail?” to “How does it usually go well, despite variability?” Incremental safety pays systematic attention to successful Work-as-Done: the adjustments, trade-offs, informal checks, and stabilising routines that allow performance to succeed despite varying and unpredictable internal and external conditions. Incremental safety treats the frequency of work that goes well as data, not as background.

Incremental safety is not an argument against investigating failures, But an argument for doing safety differently and not accept failures as the only worthwhile evidence. Learning that relies mainly on rare, high-severity outcomes will have limited data, be reactive, and only occur exceptionally. Learning that also studies frequent, expected and acceptable outcomes provides richer data for learning it can happen more frequently (nearly continuously) and lead to changes that increase capability and not just add constraints. The organisational consequences are significant. Decremental safety is inherently protective: fewer unacceptable outcomes. Efforts that restrain work to increase the absence of safety will always incur a cost – they may well be necessary, they are often resisted, and only reluctantly complied with. Incremental safety is inherently productive: more acceptable outcomes to be reliably delivered. Incremental safety is an investment in capacity, quality, and operational stability. The point is not to romanticise one and disparage the other. The point is that an exclusive focus on decremental measures quietly makes safety synonymous with restriction.

Safety culture is, in practice, a set of shared assumptions about what counts as evidence, what deserves attention, and what kinds of action are legitimate. If safety is understood as “freedom from unacceptable risk,” the organisation will prioritise counting, compliance, and barriers. If safety is also acknowledged as the routine production of expected and acceptable outcomes, the organisation can institutionalise the study of work that goes well, value adaptive expertise, and treat safety as part of performance rather than as a brake on it. The benefits of doing so may be both modest and delayed, but they matter.

#### INCREMENTAL SAFETY MAKES THE PRESENCE OF SAFETY DISCUSSABLE AND ACTIONABLE.

Two forthcoming books support and further explain incremental safety:

1. Hollnagel, E. & Slater, D. (Eds,) Incremental safety practices. This book presents thirteen studies from around the world in widely different domains that document already existing

incremental safety practices. Abingdon, Oxon: Routledge.

2. Hollnagel, E. *Decremental and incremental safety cultures; Safety-I and Safety-II revisited*. Boca Raton, FL. CRC Press (<https://www.routledge.com/Decremental-and-Incremental-Safety-Cultures-Safety-I-and-Safety-II-Revisited/Hollnagel/p/book/9781041064077>)

Incremental safety does not replace hazard identification, risk management, and protective barriers. It corrects their dominance by insisting that, safety in complex systems cannot be managed only by just trying to eliminate what does not go well. It must also be developed by first understanding— and then strengthening—what reliably goes well.

**Learning should be based on the frequency of what goes well rather than on the severity of what does not!**

For possible updates and revisions to this Manifesto, please consult  
<https://incrementalsafety.com/>

**incremental safety is for those who want to go beyond a decrease in the number of unexpected and unacceptable outcomes, and instead get an increase in the number of expected and acceptable outcomes**