



50 Years of Safety “Science”?

A personal view from someone who was there at the beginning

I was teaching chemical engineering in the 1970s. I can therefore confirm from experience that safety was not regarded either as a “science” or even as a central part of a professional engineering education. It was very definitely not included in the approved undergraduate chemical engineering curriculum.

We taught thermodynamics, fluid mechanics, heat transfer, reaction engineering, process design and control. These were the serious scientific and professional subjects. “Safety” was an occupational add-on, often undertaken by the “other ranks” rather than the graduate “officers” who designed, managed and operated the process industries.

The professional engineer designed the plant; somebody else worried about accidents, protective clothing, warning notices and compliance. (Hard Hats and Totectors?). Safety was not yet fully accepted as a design responsibility, a systems responsibility or still less, an intellectually demanding field in its own right.

That is important when considering the fiftieth anniversary of the journal **Safety Science**. Why does its history begin in 1976? Was it because of the Seveso disaster? No. The date was chosen because the **Journal of Occupational Accidents**, later renamed **Safety Science**, began publication that year.

Its original title accurately reflected the status of the subject. Safety was understood largely through occupational accidents—events that injured workers and could be counted, classified and investigated. The change of title to **Safety Science** in 1990 was therefore more than rebranding. It was a genuine attempt to elevate safety into a professional, multidisciplinary and scientifically respected field.

Fifty years later, it is reasonable to ask: **“how far have we actually succeeded?”**

I was there

Safety research did not suddenly appear in 1976. Human factors, industrial hygiene and reliability engineering were already established. Fault trees, failure-mode analysis, HAZOP and quantitative risk assessment were developing across aerospace, nuclear power and chemical processing. The Robens Report and the Health and Safety at Work etc. Act 1974 were moving Britain from detailed prescriptive rules towards broader duties and organisational responsibility.

But the field remained thin and fragmented.

When the Flixborough chemical plant exploded in June 1974, 28 people were killed, the plant was largely destroyed and surrounding properties were extensively damaged. The accident demonstrated that occupational injury statistics were a very poor measure of catastrophic process risk. A company could have an apparently respectable record of minor injuries while retaining the potential for wholesale loss of containment, explosion and destruction.

My own involvement in the Flixborough investigation made the weakness of the existing field particularly clear. Apart from Trevor Kletz's group at ICI—who had relevant experience from an incident at a similar plant but were in a potentially conflicted position—there was remarkably little established safety science on which the investigation could draw.

I was drafted in as the “safety expert” to lead the investigation team, at least partly because I was an academic from the scientifically respectable Combustion Science group at Imperial College. In truth, I was very green.

That experience stimulated a career-long commitment to finding practical ways of understanding and improving safety. What seemed to be needed was not simply more analysis, or reams of computer-generated fault trees, but methods that could help people understand how the plant and its organisation actually worked, identify where control could be lost and make better decisions before another accident occurred.

Two years later, the Seveso accident extended the problem beyond the factory boundary. A runaway reaction released TCDD dioxin over the surrounding population, creating questions of community exposure, environmental contamination, evacuation, remediation and long-term public protection.

I was also directly involved in this second defining event, leading the team that oversaw the scientific basis of the Seveso clean-up. The challenge was no longer simply to explain why an industrial accident had occurred. Science had to support practical decisions about remediation, demonstrate whether the clean-up was adequate and provide a defensible basis for protecting communities and the environment.

The two experiences exposed complementary weaknesses. Flixborough showed how little established safety science was available to investigate catastrophic plant failure. Seveso showed the need for science that continued beyond the accident itself—through containment, recovery, remediation and the justification that a damaged environment could again be considered acceptably safe.

Seveso occurred in July 1976, the same month that the **Journal of Occupational Accidents** had already begun publication. But Flixborough was probably the more immediate stimulus to British

process-safety thinking and the journal's early content, (a Kletz article), while Seveso became the defining European catalyst for major-accident legislation.

The best description of 1976 is therefore not the birth of safety science, but a moment of “crystallisation” when technical, human, organisational and regulatory traditions began to acquire a shared identity.

What has fifty years actually contributed?

It would be wrong to suggest that safety science has achieved nothing. There have been important advances.

We now recognise **process safety as distinct from personal safety**. Flixborough, Seveso, Bhopal, Piper Alpha, Texas City and Deepwater Horizon repeatedly demonstrated that low injury rates do not establish control of major hazards. HAZOP, quantitative risk assessment, layers of protection, safety cases, management of change and inherently safer design have provided practical methods for identifying and controlling catastrophic risk.

We have developed a much richer understanding of **human performance**. Human factors, ergonomics, Crew Resource Management and cognitive systems engineering showed that error cannot be understood independently of equipment, interfaces, procedures, workload, communication and context.

We have recognised **organisational causation**. Turner's disaster incubation, Zohar's safety climate, Reason's active failures and latent conditions, and the findings of major accident inquiries showed that accidents have histories. Production pressures, resource decisions, communication failures and regulatory weaknesses can shape events long before the final loss of control.

We have also acquired **systems perspectives**. Perrow drew attention to complexity and tight coupling. High Reliability Organisation research examined how hazardous organisations sustain reliable performance. Rasmussen represented risk management across interacting social and organisational levels. Leveson's STAMP treated safety as a control problem. Resilience Engineering, FRAM and Safety-II shifted attention towards adaptation, Work-As-Done and the conditions that allow everyday performance to succeed.

These are substantial contributions. They have changed the language of safety and, in many places, its practice.

But have they accumulated into a coherent science? That is much less certain.

Activity is not necessarily scientific progress

A mature science should define its objects clearly, develop testable propositions, accumulate evidence, compare alternative explanations and establish which methods work, under what conditions and with what limitations.

Safety science has often struggled to meet that standard.

Many models are powerful retrospective explanations but much weaker prospectively. After an accident, almost any framework can organise what is already known. It is considerably harder to

demonstrate that the same method would have identified the critical conditions beforehand, separated important signals from background noise and supported a decision that prevented harm.

Concepts such as safety culture, resilience, normal work and complexity have generated insight, but also multiple definitions, overlapping constructs and inconsistent measures. Different traditions use different units of analysis and different understandings of causation. The result resembles Le Coze's "Tower of Babel": a field rich in perspectives but poor in agreed integration.

There has also been a temptation to brand and commercialise each contribution. A model acquires a memorable name, diagram, training course, consultancy package and community of advocates. Commercial translation is not inherently wrong. The problem arises when differentiation becomes more important than cumulative scientific comparison. Methods are promoted alongside one another rather than tested against one another.

We acquire another product, another vocabulary and another conference stream—but still lack professional consensus about which approach should be used for which problem.

That is not good enough. Safety is not an academic fashion. It is a serious professional responsibility with consequences for lives, communities and the environment. It deserves serious science and practical agreement across the professions about the best available approaches.

Would Trevor Kletz be impressed by AI-HAZOP?

HAZOP is one of the great practical achievements of process safety. It is systematic, multidisciplinary and connected to the actual design and operation of a plant. Trevor Kletz helped explain and promote it while insisting that the preferred solution was not merely to add more protection, but wherever possible to remove or reduce the hazard through inherently safer design.

Artificial intelligence can assist a HAZOP. It can review documentation, check consistency, retrieve previous deviations, suggest guideword combinations and maintain an auditable record. Used carefully, it could improve coverage and help teams learn from a much larger body of experience.

But AI-HAZOP is the application of a new technology to a method developed more than half a century ago. It is not, by itself, a new science of safety.

I suspect Kletz would ask some characteristically practical questions. Did it identify a hazard that a competent team would otherwise have missed? Did it lead to a simpler or inherently safer design? Did it connect the present design with previous operating experience? Did it improve engineering judgement—or merely produce a longer worksheet more quickly?

AI may provide more analytical muscle, but it does not answer the underlying questions. What should the system model contain? How should different forms of evidence be combined? How do we know that an analysis is sufficient? What establishes that one intervention is better than another?

What must come next?

Safety science needs a clearer shared object of study. Safety is not a property of a component, person, procedure or culture in isolation. It is an outcome of interacting technical, human, organisational, regulatory and environmental functions within a bounded system.

It needs comparative validation. HAZOP, fault trees, bow ties, LOPA, STPA, FRAM and other methods should not exist merely as parallel choices. We need studies applying them to the same systems and evidence, showing what each identifies, what it misses and whether the resulting decisions improve outcomes.

It needs agreed evidence standards. What constitutes a sufficiently complete analysis? How should assumptions, uncertainty, data quality and model limitations be declared? What evidence justifies a claim that a system is acceptably safe, resilient or under control?

It must also connect theory with Work-As-Done. Models should be tested against operational data, observed behaviour, near misses, successful adaptations and actual outcomes. When evidence disagrees with the model, the model must change.

Functional System Modelling may offer one practical route towards this convergence—not as another branded replacement, but as a common operational framework. It represents a system as a bounded configuration of necessary, interacting and stateful functions. Technical equipment, human interpretation, organisational coordination, regulatory control, resources and environmental conditions can be represented together.

By making these functions executable, stateful and traceable, the realised system becomes the evolving record of functions coupling, acting, adapting and producing outcomes. Work-As-Imagined can be compared with Work-As-Done, while operational evidence can challenge and correct the model. This could move safety analysis beyond retrospective storytelling towards bounded simulation, comparison, testing and learning.

A science must prove itself

Safety science did not begin in 1976. I was there—teaching chemical engineering, leading an investigation team at Flixborough and later leading the team overseeing the scientific basis of the Seveso clean-up. Safety was still fighting to be recognised as part of serious professional engineering.

Fifty years later, the field has earned a place at the professional table. It has distinguished process from personal safety, improved our understanding of human performance, exposed organisational causation and developed powerful systems perspectives.

But gaining professional status is not the same as achieving scientific maturity. Too much remains retrospective, weakly validated, divided into competing schools and vulnerable to commercial repackaging. We still lack sufficient agreement about which methods work best, how they should be combined and what evidence demonstrates that they have made a system safer.

The next fifty years should not be measured by the number of new models, slogans, publications or consultancy products. They should be measured by whether safety science can make explicit predictions, compare methods, learn from operational evidence and support demonstrably better decisions before harm occurs.

In 1976, the challenge was to persuade the professions that safety deserved serious scientific attention. In 2026, the challenge is to prove that the science can deliver practical results.

The task now is not to make safety sound more scientific. It is to make safety science work.

David Slater

Selected sources

- Jean-Christophe Le Coze (2026), [“Half a century of Safety Science (1976–2026): Divergence and convergence in the field”]

(<https://www.sciencedirect.com/science/article/pii/S0925753526002341>).

- José M. Merigó et al. (2019), [“Forty years of Safety Science: A bibliometric overview”]

(<https://www.sciencedirect.com/science/article/abs/pii/S0925753517317435>) .

- Health and Safety Executive, [Health and Safety at Work etc. Act 1974]

(<https://www.hse.gov.uk/legislation/hswa.htm>) and [history of major-hazard land-use planning](<https://www.hse.gov.uk/landuseplanning/methodology.htm>) .

- Trevor A. Kletz (1976), [“Accident data—the need for a new look at the sort of data that are collected and analysed”]

(<https://www.sciencedirect.com/science/article/abs/pii/0376634976900109>) .

- European Commission, [Industrial accidents and the Seveso Directives]

(https://environment.ec.europa.eu/topics/industrial-emissions-and-safety/industrial-accidents_en) .

- Jens Rasmussen (1997), [“Risk management in a dynamic society: A modelling problem”]

(<https://www.sciencedirect.com/science/article/pii/S0925753597000520>) .

- Nancy Leveson (2004), [“A new accident model for engineering safer systems”]

(<https://www.sciencedirect.com/science/article/abs/pii/S092575350300047X>) .