University of Antwerp
Chair of Crisis Governance

ANTWERP CRISIS ISSUE PAPERS
Series editors: Professors Peter Verhezen, Hugo Marynissen, Sascha Albers

Antwerp Crisis Issue Paper 2/2026

HOW MANY WAYS CAN THIS GO WRONG? FROM MESSY WORK TO SAFER CHOICES[1]

By Frank Imans[2]

Abstract. This paper presents the ACE model (Ability of Capacity to compensate for Entropy) as a way to describe and manage the everyday complexity of work in socio-technical systems. Rather than starting from a single underlying theory, the model is built from observations of how work is actually done in settings with shifting plans, interfaces and constraints, and organises this into two core parameters: entropy and capacity. Entropy describes the viable variety in how work can proceed under given conditions, structured into five drivers: who and what is around the job, how tightly activities are linked in time, how well procedures fit the situation, how much slack and redundancy exist, and how far and fast harm can spread if control is lost. Capacity describes how much of that variety the system can absorb without losing control, through three components: structural (what is built into design and layout), operational (how work is organised and run) and adaptive (how quickly people and organisations can notice and adjust). Entropy and capacity are combined in a stress ratio that shows how far current conditions are stretching the system and how this balance moves over time. A closure index further distinguishes between settings that are relatively self-contained and those that are exposed to shifting external influences, indicating which safety responses are likely to be effective. Short case studies of the Chernobyl accident and the COVID-19 pandemic illustrate how ACE reframes familiar events in terms of structural margins, changing entropy and lagging capacity, and how it can support clearer operational and strategic decisions.

Most safety programs still assume work is stable and tidy, with the shift looking like the plan: the right checklist is used, and people follow and flow through the different steps. However, reality is not that linear. Over a single shift, tasks are resequenced; crews rotate;

---

[1] Methodological note. This analysis integrates academic literature and practice-based evidence. Practice-based evidence consists of non-public observations from the author's work as an HSE professional (in staff and consulting roles) from 2002 to the present, across multiple industries (oil and gas operations and projects; steelmaking; construction; logistics; renewable energy) and regions (Europe, the former Soviet Union and Africa). These observations are used for illustrative purposes and to inform interpretation; they do not replace and should not be read as a substitute for the cited scholarly sources.

[2] Frank Imans is a Senior HSE Consultant at Tetrax Consulting BV and the developer of the ACE model, which conceptualises safety as a dynamic balance between entropy (the operational variability and disorder inherent in real work) and capacity (the structural, operational and adaptive resources that sustain resilient performance), shifting the focus from hindsight to foresight. E-Mail: frank.imans@tetrax-consulting.com

permits lag; equipment is shared; weather moves in; and information arrives late or is noisy. A hiccup in one spot ripples across shared schedules: people queue for the same equipment; deliveries miss their slot, and handovers take longer. Most days, we can absorb it all. But, some days, the hits pile up and the system tips. This is the landscape that organisations have to work with every day.

This complexity does make safety management anything but straightforward. Over the last 40-odd years, several influential safety science models have emerged from existing theories across domains such as psychology, sociology, organisational studies, and systems engineering. Examples include Reason's Swiss Cheese model of barriers and latent conditions (Reason, 1997); Hollnagel's FRAM (Functional Resonance Analysis Method) (Hollnagel, 2012); Wood's notion of graceful extensibility (Wood, 2015); Perrow's normal accidents (Perrow, 1984) and Leveson's systems approach (STAMP/STPA) (Leveson, 2011). Rather than starting from a single underlying theory and applying it to safety, our approach takes a different route, aimed at managing this daily reality.

This approach culminated in the ACE model (Ability of Capacity to compensate for Entropy), which describes the everyday complexity of work and how to respond to it using two parameters: entropy and capacity. The ACE model is deliberately rooted in practice; designed to reflect what actually happens in real operations. It rests on a set of statements about how work is carried out in complex settings, paired with equally simple rules for deciding when it is sensible to let a job proceed and when conditions or capacity first need to change.

The point of the ACE model is to improve safety and resilience decisions when conditions are changing, not to classify systems in the abstract. While traditional risk analysis asks the questions "What can go wrong and how likely is it?" and resilience approaches "How do we keep going under stress?", the ACE model will change the way we look at the daily complexity: "How close are we to being stretched beyond our capacity today, and where is that stretch coming from?" For safety management, it provides a structured way to assess whether it is sensible to proceed, which conditions to add, and which barriers or controls need reinforcement. For resilience and crisis work, it shows whether everyday variability (entropy) is outpacing the system's ability to absorb it (capacity), and whether to act by removing avoidable variety or by building specific layers of capacity.

A closure index makes visible whether the setting is relatively closed and self-contained or more open and exposed to changing interfaces, actors, and constraints. In practice, similar safety management approaches are often applied across this spectrum. Based on the practice-based observations from the author's work across projects and operating sites, transferring closed-system control regimes into more open settings can lead to both over-control and under-control: disproportionate procedural burden in some areas, while other, more consequential gaps remain insufficiently addressed. The closure index makes explicit that safety management will not look the same in both cases: more closed settings can lean more on stable barriers and routines, whereas more open settings require stronger emphasis on managing interfaces, SIMOPS, and coordination across organisations, alongside more adaptable operational controls.

Because its structure is the same at every level, the ACE model can be used for a single job or shift as well as for a site, company, or sector. Entropy and capacity are combined into a single stress ratio that can be read in the same way by operations, safety, risk, and crisis teams, creating a shared language across these domains rather than a separate tool for each.

From daily complexity to entropy and capacity

The introduction sketched the daily reality of work: shifting plans, moving interfaces, late information, and several small disturbances sometimes landing at once. If we want to make better safety and resilience decisions in that landscape, we need a way to describe it that is simple enough to use, but rich enough not to mislead. Listing every detail is impossible. Reducing everything to a single "risk score" misses what matters most: whether the system is being stretched faster than it can respond.

To address that question, the ACE model focuses on how far current conditions are stretching the system relative to what it can safely handle. It separates this into two related questions: How much variety the system is currently exposed to, and how much capacity the system currently has to absorb that variety without losing control?

In ACE, the first is called entropy (χ) and the second is capacity (Ω). Entropy is the range of ways work can unfold under current conditions, including how far and how fast problems can spread. Capacity is the combination of structures, routines, and adaptations that allow the system to absorb that variety without it becoming harmful.

This is close to what people in the field already do informally. When they say "there is a lot going on today", they are pointing to entropy: many actors, overlaps, changes, and uncertainties that create several ways the work can unfold. When they say "we're stretched thin" or "we don't have the slack for this", they are pointing to capacity: buffers, attention, supervision, and technical protections. In practice, concern grows when the first outpaces the second. ACE gives these intuitions a shared, systematic language and a way to use them consistently.

Entropy – from Boltzmann & Shannon to field reality

In the previous section, entropy was introduced in an operational sense as the amount of variety a system is exposed to under given conditions. The notion itself originates outside safety science: in thermodynamics and information theory, where it was developed to express how many different configurations a system can take, or how uncertain a source is.

Classical entropy has two familiar faces: Boltzmann’s physical entropy, which counts how many microscopic configurations (microstates) sit behind the same macroscopic state, where more possible configurations result in higher entropy (Boltzmann, 1964); and Shannon's information entropy, which measures the uncertainty of a source, where the more different messages it can produce and the more evenly their probabilities are spread, the higher the entropy (Shannon, 1948).

They share the same root idea: as the number of possible next states increases, more effort is required to maintain order.

In socio-technical systems, this variety mainly comes from several sources: who and what is around the job; how things depend on each other in time; how well the script (procedures and plans) fits the situation; how much slack or backup exists; and how far and how fast harm can spread if control is lost.

In the ACE model, these are captured as five main entropy drivers.

(1) **Actors & Variety.** The range of actors around the job whose presence or actions can affect what is done, when it is done, or how it is done. Examples: a pour where the main contractor, two subcontractors, a crane crew, and a delivery driver are all working in the same area at the same time. A normally quiet pump area that, during a shutdown, fills with operations, maintenance, inspectors, and vendors all needing access to the same equipment.

(2) **Real-Time Coupling.** How strongly the job depends on other activities or systems, so that a change or delay in one place quickly forces changes somewhere else. Examples: several trades share one crane; when the crane is delayed on one lift, the others must stop or reshuffle immediately. A process unit that must shut down within minutes if a single upstream pump or utility fails.

(3) **Procedure-Reality Fit.** The gap between what the procedures assume and what the situation in the field actually looks like. Examples: an excavation permit states "no underground services in the area", but on site, old, unrecorded utilities are found crossing the dig line. The isolation procedure lists valves that have been removed or relocated, so technicians have to work around outdated tags and drawings to make the job safe.

(4) **Buffering & Redundancy.** How much backup and margin is available to absorb changes, delays, or failures before they start to affect the job. Examples: a schedule with no float and only one available excavator, so any breakdown or delay immediately pushes work out of sequence. A cooling system with two parallel pumps, where only one is installed; if that pump trips, there is no standby to take over.

(5) **Hazard Propagation Potential**. How far and how fast harm can spread beyond the immediate work area if control is lost locally. Examples: a high scaffold erected

above a busy walkway, where a collapse or falling objects could injure many people below, not just the crew on the scaffold. A flammable gas line in a congested process area, where a leak can quickly form a large vapour cloud and affect several units and escape routes, not only the point of release.

In practice, each driver can be scored on a simple scale and combined into a single entropy (χ) index for the job, shift, or site. The scores are normalised before they are added, so that no single driver dominates solely because of how it is measured and so that results are comparable across contexts and over time. The main value is not the absolute number, but how entropy moves: watching how the index changes with different configurations of work, and how quickly it rises or falls as conditions change.

Capacity – from safeguards on paper to holding power in the field

Where entropy (χ) describes how much variety is present, capacity (Ω) describes how much of that the system can absorb without losing control. In practice, the system's ability to absorb variety comes from three main sources: what is built into the design and layout of the system; how the work is organised and run from day to day; and how quickly people and organisations can notice changes and adjust.

In the ACE model, these are captured as three components of capacity: structural, operational, and adaptive.

(1) **Structural Capacity.** The fixed features of the plant or site layout, equipment, separations, and engineered safeguards set the basic margins within which work can be done. It covers topics such as how hazards are contained and separated, how easily systems can be isolated, how people can get in and out, and how much technical margin exists in structures, utilities, and storage. Examples: a site where excavation areas are clearly separated from traffic routes, with solid edge protection, dedicated access stairs, and fixed anchor points for lifting. A process unit with double isolation points, properly rated relief devices to a flare, blast-resistant control rooms, and clearly marked, unobstructed escape routes.

(2) **Operational Capacity.** The way the work is organised and run from day to day-planning, supervision, permits, checks and information, so that people, tasks and equipment line up in time and space. It covers how well routines are followed, how clearly responsibilities are shared, and how quickly small problems in the field are identified and addressed. Examples: a site where daily coordination meetings, clear SIMOPS rules and a working permit system keep overlapping jobs from getting in each other's way. A plant where shift handovers are structured, field rounds are done and logged, and permits, isolations and line-ups are routinely checked by a second person.

(3) **Adaptive Capacity.** The ability to notice changes, make sense of them, and adjust plans or actions in time. It covers topics such as whether people can speak up and stop or change the job, how quickly information flows to where decisions are made, and whether there are agreed ways to change course when conditions shift. Examples: a concrete delivery is delayed and will now clash with other high-risk work. The supervisor and crews meet briefly, decide to postpone the clash, and bring forward other tasks that can be done safely in the meantime. Control room and field operators spot small changes in trends, call each other, slow the operation down, and adjust line-up and setpoints before alarms start to go off.

In practice, each capacity component can be scored on a simple scale and combined into a single capacity index for the job, shift, or site. The scores are normalised before they are added, so that no single component dominates solely because of how it is measured and so that results remain comparable across contexts and over time. As with entropy, the main value is not the exact number but what it shows about where capacity is thin, and how capacity in each layer changes as work, organisation or design are altered.

Stress ratio – the balance

Now that entropy and capacity have been described, we still need a way to judge how far the current situation is stretching the system. High entropy is not a problem on its own if capacity is also high; low capacity is less of an issue if entropy is low. What matters for safety and resilience is the balance between the two at a given point in time, and how that balance changes as conditions evolve.

In ACE, this balance is expressed as a stress ratio (ρ):

$$\rho = \frac{\chi}{\Omega}$$

with χ for entropy and Ω for capacity.

Exactly where these bands sit numerically depends on the context, but the interpretation is the same: in the headroom band it is reasonable to proceed; in the thin-margin band it is sensible to add conditions and strengthen weak parts of capacity; in the stretched band the priority is to reduce entropy, increase capacity, or both, before proceeding.

Because entropy and capacity can change over time, the stress ratio is best viewed as a curve rather than a single value. Entropy can rise quickly when work is resequenced, actors are added or buffers are used up. Structural and operational capacity usually move more slowly, because they depend on design decisions, staffing or reorganising how work is run. Adaptive capacity can sometimes be changed faster, for example, by clarifying decision rights or agreeing on new responses. Many serious events arise in the lag between a rise in entropy and the moment capacity catches up. Following the stress ratio over time helps to see that lag and to act before the system is pushed beyond what it can handle.

Even when full scoring is unavailable, the same logic can be used qualitatively. Looking at changes in entropy and capacity side by side (for example, "entropy is rising; capacity is flat") already gives a sense of whether the system is moving towards or away from a stretched state, and whether the next sensible step is to reduce avoidable variety, add capacity, or do both in sequence.

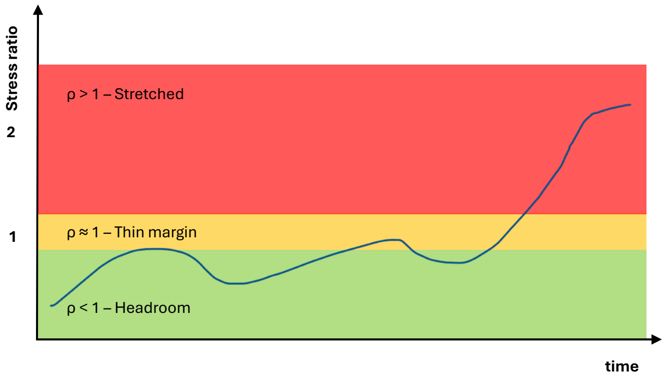

Figure 1 Example of a stress-ratio curve over time, with indicative bands for headroom (ρ<1), thin margin (ρ≈1) and stretched conditions (ρ>1). The exact limits are context-specific; the value is in seeing how ρ moves and how long it spends in each band

Closure Index – openness, exposure, and safety management

Up to this point, entropy and capacity have been described within a setting: the variety around the workplace and how the system can cope. Settings can, however, also differ in how much they are shaped by influences outside their boundary. At one end of the spectrum are settings with a porous boundary: the workfront shifts, different organisations work alongside each other, who is nearby changes frequently, and key conditions are often set or altered by others. At the other end are relatively self-contained settings: physical and organisational boundaries are clear, most important hazards and constraints sit inside those boundaries, work takes place within a defined operating envelope with engineered separations and strict procedures, and changes usually come through formal processes. To take this difference into account in a simple way, the ACE model introduces a closure index.

The closure index places a setting on this spectrum from more open to more closed. It is built from the same five entropy drivers, but instead of scoring today's variability it looks at the typical posture of the setting: how actors, coupling, procedures, buffers, and hazard control usually behave in that domain or project. Lower values represent settings where many important influences sit outside the immediate boundary and interfaces move often ("open"), and higher values represent settings where most relevant influences sit inside a clear boundary and conditions change more slowly and deliberately ("closed").

In practice, the closure index is used in two ways: it puts the stress ratio into context, and it guides which safety response makes sense. The same value of the stress ratio can mean different things in different settings, and one way of managing safety does not work equally well across the whole open-closed spectrum. In more closed settings, it is realistic to rely strongly on built-in safeguards, detailed procedures, and stable routines, because the operating envelope is well defined and most changes can be managed internally. In more open settings, the main stretch on the system comes from moving interfaces, overlapping activities, and decisions taken elsewhere; simply adding more rules and barriers of the same type can increase load without reducing the stress. Used alongside the stress ratio, the closure index makes this distinction explicit and points to different emphases.

Case studies – how ACE frames familiar events

Chernobyl (1986) - beyond safety culture: shrinking structural margins

On 26 April 1986, during a late-night test on Unit 4 of the Chernobyl nuclear power plant, a power excursion led to explosions and the destruction of the reactor. Radioactive material was released over vast areas, causing immediate fatalities and long-term health and environmental effects.

Early investigations and much of the public discussion framed the accident mainly in terms of operator error and safety culture: violations of procedures, experiments conducted outside the approved envelope, inadequate training, and weak oversight. Later reports, notably INSAG-7 (International Nuclear Safety Advisory Group, 1992), and subsequent technical summaries made this picture more balanced by highlighting both the RBMK design features (behaviour at low power, control rod design, and the absence of a full-pressure containment structure) and the organisational environment in which the test was prepared and conducted. The prevailing view today is, therefore, a combination of flawed reactor design and weak safety culture.

Broken down into an informed ACE reading, the following three elements make the balanced picture more explicit.

How the situation became unstable (demand side – entropy)

The test is often described as an “unauthorised” or “unsafe” sequence. ACE reframes this as an analysable progression in which entropy rises across several drivers at once. Actors and variety expanded beyond routine operations as responsibilities and decision rights became diffuse between planning, supervision, and execution. Real-time coupling tightened as the reactor was pushed into a regime where small state changes required rapid, correct responses, leaving little time for recovery. Procedure–reality fit degraded because procedures and assumptions no longer matched the configuration being created in the field. Buffering and redundancy were reduced as safeguards were disabled or bypassed and as the system was brought closer to operating limits. Hazard propagation potential was high because once control was lost, the consequences were not locally containable.

Put plainly: the test did not merely “break rules”; it created a configuration with more ways to go wrong, faster pathways for escalation, and fewer stabilising buffers.

Which capacity layer was missing?

Culture and decisions sit largely in operational and adaptive capacity: how work is authorised, supervised, checked, and whether people can slow down, stop, and reframe when conditions drift. Those layers were clearly thin: procedures did not fit the configuration being created; decision rights to halt were weak in practice; and once the sequence was in motion, there was limited ability to stop, delay, or redesign the work. However, ACE also makes the structural layer central. INSAG-7's technical emphasis maps naturally onto "structural capacity": the reactor's structural margins were sharply state-dependent; outside certain operating conditions (notably low-power regimes relevant to the test), behaviour could deteriorate rapidly. In such a setting, operational excellence and a questioning culture reduce the likelihood of entering the brittle region, but they do not remove the underlying vulnerability of a design that can be driven into a state where margin collapses quickly.

This is the core shift: the accident is best understood as a case where culture and governance were asked to compensate for structural margins that were thin in parts of the reachable operating space.

What would have changed if leaders had an ACE-lens before the event?

The practical implication is not simply "strengthen safety culture". ACE would have focused attention on preventing combinations of entropy increases and capacity reductions that push the stress ratio into a stretched region. Concretely, it would have raised three decision points:

(1) Treat entry into brittle regimes as a capacity problem, not a behavioural one. When the operating configuration moves into conditions where structural margins are known to shrink (e.g. unstable low-power regimes), the default assumption should be that routine operational capacity is no longer sufficient. Proceeding then requires explicit compensating capacity (independent technical review, enhanced monitoring, conservative setpoints, and clear stop authority) or the configuration should be avoided;

(2) Make safeguard inhibitions visible as a reduction in capacity that must be compensated. Disabling or bypassing protections is not just a procedural violation;

it is a measurable reduction in operational capacity. If safeguards are inhibited, the decision must explicitly state what compensating capacity is being put in place (alternate protections, independent verification, strengthened supervision/monitoring) or the activity should not proceed;

(3) Design stop-and-reset capability for non-routine configurations. When procedure-reality fit is degrading (the script no longer matches the configuration being created), adaptive capacity must be built into the work: defined pause points, rapid escalation to competent review, and a practical mechanism to return the system to a stable configuration before continuing.

Seen this way, ACE operationalises INSAG-7's dual message. It provides a structure that connects design limitations (structural capacity), organisational context (operational and adaptive capacity), and the test's evolving conditions (entropy) into one coherent picture of why the system became stretched.

The case suggests a general lesson that carries beyond nuclear power: in high-hazard, high-closure settings, safety performance depends on keeping work inside the envelope where structural capacity is valid, and on preventing routine organisational pressures from trading away operational and adaptive margin. Chernobyl illustrates that when a system can be driven into a configuration where entropy rises rapidly while capacity is simultaneously reduced, the stress can outrun the system's ability to recover, regardless of how the final moments are narrated.

COVID-19 (2020+) – when entropy outruns capacity

By early 2020, the spread of SARS-CoV-2 turned into a global pandemic. Health systems faced recurrent surges, essential supply chains were disrupted, and governments introduced successive measures affecting work, travel, and social life. The impact included large numbers of deaths, long-term health effects, sustained pressure on hospitals, and wide social and economic consequences.

The pandemic is often discussed in terms of specific decisions: lockdowns, border measures, testing strategies, vaccination campaigns; or in terms of preparedness failures. ACE does not dispute these debates, but it reorganises them around a more operational question: how fast did entropy rise, how fast could capacity respond, and how long did the system spend in the gap between them?

Broken down into an ACE-reading:

How the situation became unstable (demand side - entropy)

Entropy increased rapidly and repeatedly. Transmission chains multiplied and shifted geographically; outbreaks emerged in clustered settings and then propagated through broader mobility networks; new variants changed key parameters; and information remained uncertain and delayed, especially early on. This is not “bad behaviour” in itself; it is the practical expansion of the number of plausible futures the system had to manage at once, often with short decision cycles and incomplete data.

Which capacity layer was missing (and how it lagged)

The limiting factor was not a single “missing barrier,” but the time constants of capacity. Structural capacity such as ICU beds, oxygen infrastructure, equipment, trained staff, and vaccine development and production could not be expanded at the speed at which entropy increased. Operational capacity including testing and tracing systems, lab throughput, supply distribution, hospital reconfiguration, staffing models, and governance routines could move faster, but still required lead time, coordination, and repeated learning cycles. Adaptive capacity such as sense-making, decision-making, public communication, and behavioural changes at between structural and operational layers: it could respond quickly in principle, but was constrained by uncertainty, trust, and fatigue. The practical consequence was a prolonged period in which χ rose faster than Ω and the stress ratio remained elevated.

What would have changed if leaders had an ACE lens before the event?

The principal implication is not "reduce entropy," much of which is inherent to an open and mobile society. The leverage lies in shortening the lag: building operational and adaptive capacity in advance so that Ω can accelerate as soon as χ rises. In concrete terms, this points to "capacity playbooks" rather than generic preparedness plans: pre-defined mechanisms to scale testing and tracing, flex staffing and ward models, switch supply chains for critical items, stand up temporary facilities, and activate pre-agreed triggers for distancing, masking, or reorganising work. Vaccines remain crucial structural capacity, but their development is uncertain in timing; the more controllable objective is to ensure that operational and adaptive capacity can be mobilised quickly and repeatedly, keeping the lag as small as possible.

Taken together, COVID-19 illustrates a general ACE lesson for system crises: failure is often produced not by the presence of entropy, but by the duration and magnitude of the period in which entropy outruns capacity. Prepared "how-to-scale" capacity arrangements, especially operational and adaptive, are therefore a primary resilience investment, because they determine how quickly a system can regain margin once conditions begin to change.

Read together, Chernobyl and COVID-19 show why ACE is more than a relabelling exercise. In both cases, deterioration is driven by the relationship between rising entropy (χ) and available capacity (Ω), not by any single "root cause." Chernobyl illustrates a rapid move into a configuration where structural margins were sharply state-dependent and operational and adaptive capacity were simultaneously reduced: the stress rose faster than the system could recover. COVID-19 illustrates the same logic over a different time scale:

entropy increased quickly and repeatedly while structural capacity could only grow slowly: outcomes were shaped by the length of time the system spent in that gap.

Across both, two practical themes recur. First, the decisive danger sits in the lag between changing conditions and the mobilisation of capacity. Second, the most effective interventions are those that make capacity explicit: protecting it from being traded away in high-hazard settings, and building "how-to-scale" operational and adaptive playbooks in system crises. So that capacity can accelerate when entropy rises. ACE therefore adds value by making stretch visible early, identifying which layer of capacity is limiting, and turning complex narratives into repeatable decision levers.

Conclusion – from description to decisions

This paper started from the mismatch between how work is usually described in safety models and how it actually unfolds in complex settings. Much of the last decades' work in safety science has grown out of existing theories in psychology, sociology and organisational studies. The ACE model (Ability of Capacity to compensate for Entropy) takes a different route. It is built empirically, from observations of how work is done in real operations, and organised around two parameters that are directly tied to that reality: entropy and capacity.

Entropy captures the viable variety in how work can proceed under current conditions, structured into five drivers: who and what is around the job, how tightly activities are coupled, how well procedures fit reality, how much slack and redundancy exist, and how far and fast harm can spread if control is lost. Capacity captures how much of that variety the system can absorb without losing control, through three components: what is built into structures and layout, how work is organised and run, and how quickly people and organisations can notice changes and adjust.

These are combined in a stress ratio that expresses how far current conditions are stretching the system. The ratio is not a precise risk number; it is a way to see whether entropy is outpacing capacity, how fast that balance is moving, and where the thinnest layer of capacity lies. The closure index adds a further distinction by showing whether the setting is relatively closed and self-contained or open and exposed to shifting interfaces and external constraints, and thus, which kind of safety management response is likely to be effective.

The practical value of ACE lies in this structure. It offers a common language that operations, safety, risk, and crisis teams can share. It keeps the focus on how entropy and capacity change over time rather than on single scores, and it encourages preparations that build capacity explicitly, rather than assuming that more rules or barriers will always be sufficient. By tying entropy and capacity directly to how work is done, ACE makes visible what most traditional approaches leave implicit: how close the system is to being stretched, which parts of capacity are thinnest, and where to act before rising variety turns into failure. This makes it particularly suited to complex socio-technical settings, where predefined scenarios and static risk matrices struggle to keep pace with changing configurations of work, but where watching the balance between entropy and capacity still allows for timely, concrete decisions.

References

Boltzmann Ludwig (1964). *Lectures on Gas Theory*, translated by Brush, Stephen G., Dover Publications, Inc.

Hollnagel, Erik (2012). *FRAM: the Functional Resonance Analysis Method: modelling complex socio-technical systems*, Ashgate Publishing

International Nuclear Safety Advisory Group (1992). *INSAG-7. The Chernobyl Accident: Updating of INSAG-1*, International Atomic Energy Agency

Leveson, Nancy G. (2011). *Engineering a Safer World: systems thinking applied to safety*, The MIT Press

Perrow, Charles (1984). *Normal Accidents. Living with High-Risk Technologies*, Basic Books, Inc., Publishers

Reason, James (1997). *Managing the risk of organizational accidents,* Ashgate Publishing

Shannon C.E. (1948). A Mathematical Theory of Communication, *The Bell System Technical Journal*, 27, 379-423

Woods, David D. (2015). Four concepts for resilience and the implications for the future of resilience engineering, *Reliability Engineering & System Safet*y, 141, 5-9.

Milward, H. B., & Provan, K. G. (2000). Governing the hollow state. *Journal of Public Administration Research and Theory*, 10(2), 359-380.

The Chair of Crisis Governance at the University of Antwerp has been established with the financial support of the University Fund and these external partners:

- Brandweer Zone Antwerpen
- Brussels Airport
- PM • Risk Crisis Change
- Port of Antwerp-Bruges

For more information, please contact Prof. dr. Sascha Albers (sascha.albers@uantwerp.be) and/or Prof. dr. Hugo Marynissen (hugo.marynissen@uantwerp.be) or check out the Chair’s webpage.