DECISIONAL COLLAPSE IN RAPID-ONSET EMERGENCIES

Neurobiology of Motor Freeze, Pre-Evacuation Latency, and Real Survival Time Margins

Dr. Laura Baccaro

Forensic Psychologist — Honorary Judge, Surveillance Court of Venice

Scientific Report

Prepared for PS 1978 Limited — SafeWay Project

Abstract

This report addresses a question that the emergency management literature habitually treats in two separate streams — which this paper argues constitute a single operational problem. On one side, the technical lead time that warning systems manage to generate before a destructive event reaches its target; on the other, the time a typical person requires to convert that warning into physical movement. The two quantities are normally measured by different scientific communities — operational meteorology and seismology on one hand, survival psychology and stress neuroscience on the other — and are almost never subtracted from each other.

When that subtraction is performed, the residual survival margin is zero or negative across a substantial class of scenarios. Not because the warning is absent or late, but because the recipient is biologically occupied, during the minutes that matter, by a freeze response that his or her will does not control. The argument of this report is that the bottleneck of civilian survival in rapid-onset emergencies is no longer technological or informational, but neuro-cognitive, and that its only practicable mitigation at population scale consists in transferring the decisional load away from the victim's brain.

1. Statement of the Problem

Emergency management rests, almost universally, on an implicit assumption: if we warn people in time and have told them what to do, they will do it. The assumption is reasonable, and it has guided decades of public investment toward shortening detection time and widening the reach of dissemination channels. That investment has produced real, measurable results. The average lead time of U.S. Flash Flood Warnings rose from roughly fifteen minutes in the late 1980s to roughly sixty-five minutes by 2021 [B4]; tornado detection probability grew in parallel with average lead time, which increased from around five to approximately thirteen minutes over two decades [B6][B7]. On the technical plane, the warning system functions better today than it ever has.

The figure that does not fit is that the technical gain translates into reduced fatalities only up to a threshold, and then levels off. Research on tornado warning systems shows that extending lead time beyond fifteen minutes has a modest additional effect on fatality reduction [B5]. If the bottleneck were technical lead time, that curve should not flatten: more warning should monotonically mean fewer deaths. The fact that it does flatten indicates that downstream from the warning there exists a second process with its own characteristic duration — one that absorbs the time gained upstream.

That second process is human behaviour in the first minutes after the signal. In the evacuation literature it is called pre-evacuation delay: the interval separating the receipt of a danger cue from the onset of actual movement. It is an interval that behavioural science has measured with increasing precision and that stress neurobiology has, over the past fifteen years, explained at the molecular level. This report holds both measures together — technical time and human time — and draws the operational consequences.

2. Technical Lead Time: How Much Warning Actually Exists

Before subtracting anything it is necessary to establish the order of magnitude of the minuend. The figures below are drawn from the operational definitions and historical series of NOAA/National Weather Service, USGS, and the Pacific Tsunami Warning Center, and describe net lead time — the interval between warning issuance and physical impact.

2.1 Tsunami

The fundamental operational distinction is between near-field and far-field sources. A near-field tsunami, generated by a submarine fault within roughly one hundred kilometres of the coast, arrives in minutes — often under an hour, in many cases under twenty minutes [B1]. In the best case, when seismic and tide-gauge sensors near the source detect the event within three to five minutes and the warning centre issues immediately, twenty to thirty minutes can be gained [B2]. In the worst case — which is not rare — the wave front reaches the coast before the warning has been disseminated, and effective lead time lies between zero and five minutes [B1]. PTWC/NOAA documentation supplies a figure that frames the entire discussion that follows: more than ninety percent of tsunami fatalities occur precisely where warning time is shortest [B2].

The far-field, or tele-tsunami, generated by a distant oceanic source, behaves oppositely. Warning time is measured in hours — typically at least three, often four to eight or more for large ocean basins [B1][B3] — and in this regime the limiting factor is no longer available time but the organisational quality of the evacuation response.

2.2 Flash Flood

The National Weather Service adopts a performance target of no less than sixty-five minutes lead time for Flash Flood Warnings, and the national average sits consistently near that value [B4][B8]. The average, however, conceals substantial dispersion. Severity-stratified operational analyses show that events classified as considerable severity — those that actually kill — carry an average lead time of around twenty-four minutes, against roughly fifty minutes for base-severity events [B9]. In small, steep catchments hydrological response to critical rainfall can occur in under an hour and sometimes in under thirty minutes, and in a fraction of cases the warning is issued after the event has begun — yielding zero or negative lead time [B9][B4].

2.3 Tornado

Tornado is the phenomenon with the tightest technical margin in absolute terms. The institutional target is a mean lead time of at least thirteen minutes, and recent data cluster between thirteen and fifteen [B5][B10]. Here too the distribution matters more than the mean: the first tornado of an outbreak has a median around thirteen or fourteen minutes, while subsequent tornadoes in the same sequence reach medians of nineteen to twenty-two minutes, because operational alertness is by then maximal [B11]. At the extreme, the first tornado of an intense outbreak may be preceded by four minutes of warning, and unwarned tornadoes count as zero-lead-time events in national averages

[B11][B10].

Table 1 — Technical lead time by phenomenon (best case / worst case).

Phenomenon	Maximum warning	Minimum warning	Frequency of worst case
Tsunami (near-field)	20–30 min	0–5 min	Frequent
Tsunami (far-field)	≥ 3–8 hours	1–2 hours	Rare
Flash flood	90–120+ min	0–15 min	Common in fast-response catchments
Tornado	30–60+ min (exceptional)	0–5 min	Typical for first event of outbreak

3. Human Time: The End of the Will-to-Live Myth

For much of the twentieth century, survival under extreme conditions was attributed to a presumed will to live, or to exceptional personality traits. The implicit corollary of that reading — that those who perished were in some measure deficient in that vital drive — went largely unchallenged in popular narrative and in part of psychiatry. Modern survival psychology, to which John Leach's work in civil and military settings has contributed decisively, has shown that conception to be not only obsolete but devoid of explanatory power: it cannot account for the enormous number of victims who die without any organic cause, despite having had time, means, and available escape routes [A2][A6].

The academic centre of gravity has shifted accordingly from the study of motivation to the investigation of cognitive failures. In emergency, what is missing is not will: what is missing is the mind's capacity to process environmental information. Cognition, which underlies all purposeful human behaviour, undergoes severe deterioration under life threat. In the same movement, the classical fight-or-flight formula has been recognised as an incomplete description of actual disaster behaviour: extreme stress far more frequently produces fight, flight, or freeze [A6][A7].

Analyses of eyewitness accounts, survivor debriefings, and official investigations into maritime, fire, and aviation disasters have identified freezing — motor block — as the most frequently observed and reported response. Leach describes a behavioural distribution among populations caught in a catastrophic impact that has proved notably stable across different events and contexts [A8][A9].

Table 2 — Behavioural distribution of the population at impact (Leach).

Category	Typical share	Psychological and cognitive response
Adaptive / rational	10–20%	Retain composure and the ability to evaluate options. Act appropriately and productively for survival.
Stunned / confused	70–80%	Cognitive paralysis, disorientation, inaction, severe deterioration of logical thought. Appear docile, waiting for directives, or blindly imitate ineffective behaviours.
Maladaptive	10–15%	Irrational behaviour, unproductive hyperkinesia, panic, paralysing anxiety, or severe denial. Frequently endanger themselves and others.

The correct reading of this table is not that eighty percent of people react poorly. It is that eighty percent of people react exactly as they are built to react — and that the population model on which emergency plans are designed, the informed individual who receives a message, interprets it, and acts, describes a minority.

3.1 Dysexecutive Survivor Syndrome

The state in which that majority finds itself has been formalised clinically as Dysexecutive Survivor Syndrome. During the impact phase — characterised by speed, extreme danger, sensory violence, and uncontrollability — survivors exhibit a transient, non-clinical dysexecutive syndrome consisting in the fragmentation of normal, coordinated behaviour [A8][A9].

Manifestations form a recognisable subset of maladaptive behaviours: loss of initiative, hypokinesia, stereotypy, thought and action perseveration, frantic hyperkinesia, and in extreme cases total akinesia. Continuing to search for luggage in a smoke-filled cabin, or remaining seated waiting for an order no one is giving, belong to this phenomenology. The critical point is that these behaviours close a fatal feedback loop: cognitive delay increases environmental exposure, and increased danger in turn intensifies the severity of the paralysis.

4. Pre-Evacuation Latency and Milling Behaviour

In the evacuation sciences the time required to physically initiate escape is one of the most determinant parameters for survival in rapid-onset events. Traditional engineering modelling treated occupants as homogeneous entities — the familiar ball-bearing metaphor — moving instantaneously and rationally toward the nearest exit. Behavioural science has dismantled that model: latencies are extended, heterogeneous, and governed by psychological and social factors, with field values ranging from a few tens of seconds to over seven minutes depending on perceived severity, social affiliation, and cognitive state [A10][A11][A14].

The first response to an alarm signal or an environmental danger cue is almost never flight. It is denial, or a compulsive search for additional information — a cognitive and sociological process known as milling behaviour. Humans are intensely gregarious and constantly seek validation from peers; if others nearby do not react, the observer conforms to their passivity, doubts the severity of the threat, and normalises the ongoing danger. This lethal conformism, combined with the inability to formulate a new action plan, extends latency to the point where environmental conditions no longer permit survival.

4.1 Field Measurements

Studies of real evacuations supply numbers that must be placed alongside those in Section 2. For tsunami, the mean time to initiate evacuation after the earthquake or warning is approximately nineteen minutes among survivors and twenty-one minutes among fatalities; the line separating life from death in those datasets runs at around twenty minutes [B12]. In the 2011 Tohoku event, 54.5 percent of those who delayed evacuated between twenty and sixty minutes after the shaking [B13]. In American Samoa in 2009, many waited fifteen minutes or more before moving, and the reason cited is recurrent: the search for multiple confirmations — shaking, sea withdrawal, neighbours' reactions — before accepting that the threat was real [B14]. This is milling behaviour observed in the open field.

Operational models decompose that latency into two components. Individual Decision Time — the interval to recognise the threat and decide to evacuate — is estimated at five to ten minutes in conservative formulations; Evacuation Preparation Time — the interval to prepare oneself and family before moving — at three to nine minutes for residents and around three minutes for beachgoers [B15][B16]. For near-field tsunami, in the absence of clear directives and regular drills, the operational reference estimate is ten to twenty minutes of mean latency, with values above twenty or thirty minutes where risk culture is not consolidated [B12][B13][B14].

For flash flood, behaviour differs but the latency structure is identical: systematic underestimation of water depth and velocity means many people act only when water is visible in the street or when they receive repeated alerts. In urban scenarios the conservative estimate is five to fifteen minutes of latency for risk-aware residents, rising to fifteen to thirty or more when risk perception is low or the warning is purely textual and lacks operational instructions.

For tornado, the required response is not area evacuation but reaching shelter. With good awareness and a household plan, reaching a safe room takes two to five minutes; in complex structures — offices, schools, hospitals — or in the absence of a plan, it rises to five to ten minutes or more, especially where persons of reduced mobility are present [B5].

4.2 The Manchester Airport Disaster, 1985

A single event illustrates the mechanics of the block more clearly than any statistical series. On 22 August 1985, during the take-off roll of British Airtours flight 28M at Manchester Airport, a combustion-chamber failure in the left engine of a Boeing 737 punctured the fuel tank and ignited a catastrophic fire. Pilots aborted the take-off and stopped the aircraft on the runway twenty seconds before the evacuation order. Fifty-five people died out of 137 on board, the large majority not from burns but from inhalation of toxic fumes produced by cabin materials, in a state of incapacitation [A15][A16].

The Air Accidents Investigation Branch report and subsequent reconstructions document how psychological block and lack of initiative turned a ground incident into a mass fatality. At the sudden appearance of dense black smoke, passengers exhibited a disorienting paralysis; witnesses described a mass of people locked and entangled in the central fuselage section, entirely unable to advance. Rather than evaluating alternative routes, the crowd pressed toward the forward exits while the usable right-wing exit was opened with fatal delay [A17][A18].

Manchester crystallises the central theoretical point of this report. When the learned procedural schema fails — the planned door is blocked, the planned route is impassable — the brain requires time to formulate a new strategy, and the environment does not grant that time. Dozens of people died within metres of safety not for want of exits, but because processing spatial alternatives under acute stress was neurobiologically unavailable.

5. Why It Happens: The Neurobiological Bases of Decisional Collapse

Behavioural description alone remains phenomenology. The reason why the freeze is predictable, reproducible, and substantially independent of the individual's will lies in the conflict between two cerebral systems: higher cortical executive functions and subcortical defensive responses.

5.1 The Vulnerability of the Prefrontal Cortex

The prefrontal cortex is the brain's evolutionarily most recent region and sustains the capacities that define rational agency: working memory, focused attention, cognitive flexibility, impulse inhibition, strategic planning. It is where the recall of evacuation plans, received training, and the capacity to decode complex instructions reside. And it is, as decades of research — and in particular the work of Amy Arnsten at Yale — have established, the region most sensitive to the deleterious effects of stress, bar none [A21][A22][A24].

Under acute, uncontrollable stress the autonomic nervous system immediately mobilises the sympatho-adrenal axis, flooding the brain with catecholamines — dopamine and noradrenaline in the first instance — followed by glucocorticoids. At moderate doses these substances optimise attention and performance. At the supranormal levels secreted during trauma they trigger a biochemical cascade that structurally degrades prefrontal networks.

5.2 The Disconnection Mechanism

At the molecular level, excessive stimulation of noradrenergic alpha-1 and dopaminergic D1 receptors activates feedforward intracellular signalling cascades: intracellular calcium release, activation of protein kinase C and, in parallel, of the cyclic AMP pathway coupled to protein kinase A. The terminal effect of these cascades is mechanical: rapid opening of potassium channels on the dendritic spines of prefrontal pyramidal neurons. Potassium efflux hyperpolarises the membrane and drastically weakens synaptic efficacy [A22][A24][A28].

In functional terms, the prefrontal cortical architecture is temporarily disconnected. Networks lose the capacity for the persistent firing that constitutes the physiological basis of working memory and abstract thought. Simultaneously the same catecholamine concentrations strengthen synaptic transmission in primary sensory cortices, the striatum, and above all the amygdala. Behavioural control therefore shifts, rapidly and involuntarily, from a reflective, flexible modality governed by the prefrontal cortex to a reflexive, rigid modality mediated by amygdala and brainstem circuits. Without prefrontal supervision, cognitive flexibility vanishes: the individual becomes neurobiologically incapable of implementing novel plans, evaluating alternative escape routes, or interpreting complex spatial information.

5.3 Amygdala, Periaqueductal Grey, and Freezing

Once executive control has collapsed, responses are dictated by primitive survival circuitry. The central nucleus of the amygdala functions as the primary output hub for threat responses, projecting massively to the midbrain periaqueductal grey, which coordinates defensive strategies. PAG anatomy reflects distinct responses: dorsal and dorsolateral columns are associated with active defences — fight, flight — with sympathetic tachycardia and hypertension, while lateral and ventrolateral columns are activated when the threat is perceived as inevitable or when the organism must assess the environment without being detected [A29][A31][A32].

Activation of the lateral and ventrolateral PAG triggers freezing. Contrary to the prevailing narrative associating fear exclusively with adrenergic surge, freezing is characterised by massive parasympathetic dominance. Its physiological signature is fear bradycardia — a sudden, severe deceleration of heart rate. Neuronal tracing studies have documented a direct monosynaptic pathway from vlPAG neurons to the nucleus ambiguus in the medulla oblongata, the primary origin of parasympathetic preganglionic neurons innervating the heart [A33][A35][A37].

Confronted with the roar of an earthquake or a tsunami warning, this pathway transmits fear-driven central signals that initiate parasympathetic outflow, which acts as a brake on both the motor and cardiac systems. The individual cannot move because the brainstem has imposed motor inhibition while computing an optimal escape route or attempting environmental concealment. Freezing is, in origin, an alert, preparatory immobility — an evolutionary advantage. It becomes lethal when the ongoing event requires instantaneous muscular action that fear bradycardia precludes [A36][A38].

5.4 Tonic Immobility

Motor block covers a spectrum, and the clinical distinction carries operational consequences. If the threat becomes overwhelming, sustained, or involves unavoidable physical contact — entrapment under rubble, for instance — the nervous system may escalate from freezing to tonic immobility: an extreme, obligatory, submissive defensive response homologous to thanatosis observed in many prey species [A39][A40].

Tonic immobility is not a preparatory state. It is extreme muscular rigidity, total loss of responsiveness, deep analgesia, and a drastic drop in arterial pressure, which distinguishes it sharply from the cardiovascular profile of freezing. Victims report having been fully conscious, registering every detail mentally, while experiencing absolute involuntary paralysis, unable even to produce a sound to call for help. Having experienced peritraumatic tonic immobility is, moreover, one of the most robust clinical predictors for subsequent development of severe post-traumatic stress disorder [A42][A43].

5.5 Neuroception and the Polyvagal Hierarchy

The framework that holds these levels together is the polyvagal theory introduced by Stephen Porges, which redefined the autonomic nervous system's operation by superseding the simple dichotomy of sympathetic excitation versus parasympathetic calm. Threat response follows a phylogenetic hierarchy of three overlapping evolutionary strata [A44][A47][A48].

Table 3 — Polyvagal hierarchy of threat response.

Evolutionary stage	Dominant system	Behavioural and physiological response
Most recent (mammals)	Ventral vagal complex (myelinated)	Regulates social engagement and safety. Acts as a vagal brake modulating heart rate; enables communication and active listening.
Intermediate	Sympathetic nervous system	When social engagement fails, the vagal brake is withdrawn: massive mobilisation, tachycardia, adrenaline cascade, fight or flight.
Most ancient (primitive reptiles/amphibians)	Dorsal vagal complex (unmyelinated)	Intervenes when threat is overwhelming or inescapable: extreme bradycardia, hypotension, apnoea, dissociation, tonic immobility.

The wholly subconscious process by which the nervous system discriminates between safety, danger, and lethal threat is what Porges terms neuroception, and it operates constantly below cognitive awareness. This detail is decisive for the scenarios considered here. Confronted with a twenty-metre tsunami wave or an expanding pyroclastic cloud, neuroception detects a threat of such magnitude that the probability of successful active escape is evaluated — subcortically, not rationally — as approaching zero. The system then bypasses sympathetic fight-or-flight activation and plunges directly into dorsal vagal engagement. The body enters extreme conservation mode and the individual is left frozen in place. This is not a character defect: it is the correct response of an archaic system to a threat it classifies as unescapable.

6. Why Training Does Not Solve the Problem

The most uncomfortable conclusion, from the standpoint of civil protection policy, is that theoretical training, preventive information, and rational knowledge of procedures do not significantly mitigate this reaction. The axiom underpinning much emergency management — if we educate people on what to do during an earthquake, they will do it — presupposes that the brain retains intact access to explicit long-term memory and prefrontal computational resources during the crisis. This is exactly what the preceding section excludes.

6.1 Collapse of the Supervisory Attentional System

The Norman and Shallice model of the Supervisory Attentional System describes the mechanics of the failure. In ordinary life, behaviour is guided by contention scheduling, which executes learned schemata automatically and without conscious effort. Faced with a novel situation — a floor giving way, a valley flooding — established schemata become lethal, and the individual should activate the SAS to suppress the automatism, evaluate new environmental variables, and synthesise an adaptive response. The critical constraint is that the SAS, to function, requires two resources that emergencies systematically destroy: time and intact working memory [A52].

Because the disaster evolves at unnatural speed and saturates sensory processing, the SAS's capacity to act as a buffer for building a survival response collapses. The person who has read the instructions, but has not been militarily trained, regresses into stereotypy and paralysis precisely because the supervisory system cannot assemble an action plan in real time [A6][A52].

6.2 Experimental Evidence

Prefrontal disconnection is a neural hardware response, not a software gap correctable by supplying more information. Having memorised an evacuation route or attended a fire drill engages prefrontal synaptic networks; the moment the ground shakes, the biochemical cascade temporarily severs the dendritic connections that hold that explicit memory.

The experimental record is direct. A study measuring cognitive performance in volunteers who had just completed a basic fire-safety training course found that exposure to a realistic emergency simulation massively elevated salivary cortisol and simultaneously produced a statistically significant impairment of visual attention and working memory — in subjects who had finished the procedural training minutes earlier [A53]. On-field cognitive deterioration nullifies the efficacy of formal training, explaining why even informed individuals exhibit freezing or maladaptive responses during real fires.

A single documented neurobiological exception exists: overlearning. Repeating a gesture or procedure tens of thousands of times, until the neural pathway migrates from the fragile prefrontal cortex to the basal ganglia and cerebellum, where it becomes a subcortical reflex that survives panic. This is the method of special forces and test pilots. It is also, for structural and cost reasons, impossible to apply to a civilian population for every conceivable natural disaster scenario.

The paradox thus remains unresolved within current terms: authorities ask citizens to act rationally in the face of catastrophe, while human biology, confronted with the same catastrophe, suspends precisely the analytical faculties that would be required to do so.

7. The Time Balance: Subtracting Human Time from Technical Time

The two measures can now be placed side by side. The operational question is not how much warning the alert system manages to produce, but how much remains after the recipient's nervous system has drawn its own share.

Table 4 — Residual survival margin: technical lead time minus behavioural latency (population without explicit directives).

Phenomenon	Best-case lead time	Worst-case lead time	Mean pre-evacuation latency	Worst-case margin
Tsunami (near-field)	20–30 min	0–5 min	10–20 min (often >20 min)	Negative
Tsunami (far-field)	≥ 3–8 hours	1–2 hours	10–30 min	Ample
Flash flood	90–120+ min	0–15 min	5–15 min; 15–30+ if risk perception low	Negative
Tornado	30–60+ min	0–5 min	3–7 min if trained; 7–15+ in complex settings	Negative

Three rows out of four close in deficit. For near-field tsunami, technical warning in the worst case is zero to five minutes while mean behavioural latency is ten to twenty: the person begins to move when the wave has already arrived. For flash flood in a fast catchment, lead time may be null or negative while decision delay remains positive. For tornado, four minutes of warning against five to ten minutes needed to organise movement in a school or hospital produces the same result.

This arithmetic explains the two apparently disconnected data points cited at the outset. First: more than ninety percent of tsunami fatalities concentrate where lead time is shortest [B2] — precisely where the residual margin is structurally negative. Second: extending tornado warning beyond fifteen minutes does not further reduce fatalities in any appreciable measure [B5]. When the dominant term of the subtraction is no longer the minuend but the subtrahend, acting on the minuend ceases to produce effects. Every additional resource invested in advancing detection yields diminishing marginal returns as long as the behavioural delay stays where it is.

The operational literature implicitly acknowledges this when it recommends, for near-field tsunami, driving latency below ten minutes through clear signage, short evacuation paths, and regular drills, or when it recommends for tornado massive investment in institutional drills to bring shelter-seeking time below five minutes. These recommendations are correct in direction and insufficient in means: they seek to reduce the delay by acting on training and information — the two levers that Section 6 demonstrates to be neurologically blunt outside overlearning regimes.

It follows that pre-evacuation latency must be attacked elsewhere: not by reducing the time the person takes to decide, but by eliminating the decision.

8. Cognitive Load and Decision Support Architecture

Cognitive Load Theory provides the technical framework for that elimination. Its postulate is that the mind disposes of a finite information-processing capacity. In a survival scenario intrinsic load explodes — terror, death threat, acoustic and visual sensory saturation — and the residual resources available to process additional information or evaluate articulated directives approach zero.

The cleanest evidence comes from emergency medicine. Studies on dispatcher-to-layperson communications show that complex or multiple-choice scripts generate overload, delay action, and increase errors, while eliminating verbal redundancy and supplying a single sequential instruction — the principle of segmenting — bypasses paralysis-by-overload and enables action even in a state of shock [A1][A55]. Clinical Decision Support Systems in emergency departments operate on the same principle: they automate information acquisition and analysis, limit cognitive workload, and supply selective recommendations, reducing diagnostic latency and procedural error [A56][A57].

Transferring this principle from medical emergency to territorial emergency is conceptually linear and technically demanding, because the computation to be automated is not clinical but physical and spatial. Asking a person in dorsal vagal activation whether an upper floor will withstand a wave impact or whether running toward a hill of unknown distance is preferable is asking for an operation that their prefrontal cortex, at that moment, cannot perform. Each disaster scenario additionally imposes a different — and often counter-intuitive — escape topology.

Table 5 — Computational load required from the victim by scenario.

Scenario	Physical variables to assess	Correct response	Load
Flood	Terrain elevation, water velocity, proximity of high ground.	Reach verified high ground before water level precludes it.	High

Scenario	Physical variables to assess	Correct response	Load
Tsunami	Coastal morphology, tall structures versus natural ridges.	Vertical evacuation to structure or inland flight, depending on coastline.	Critical
Earthquake	Surrounding structural integrity, collapse potential.	Open spaces away from collapsible structures and projectile objects.	Extreme
Tornado	Wind trajectory, solid masses, underground shelters.	Descent below ground level or shelter behind immovable mass.	High
Eruption	Wind dynamics for ash, valley orography for lava.	Vectorial movement toward areas the hazard will not cross.	Maximum

The right-hand column measures, in effect, the distance between what the scenario demands and what the nervous system can supply in the first minutes. None of these assessments is within reach of a disconnected prefrontal cortex, and all must be completed within time windows that Table 4 has shown to be frequently already exhausted.

9. Applied Implications: Decision Support as a Neuro-Cognitive Countermeasure

From the convergence of these findings a precise functional specification emerges — not a generic technological recommendation. A system designed to reduce civilian mortality in rapid-onset emergencies must satisfy three conditions, each corresponding to a documented neurobiological constraint identified in the preceding sections.

The first is that the system performs the environmental and spatial computation on behalf of the user, rather than supplying data for the user to process. The prefrontal cortex is disconnected at precisely the moment it would be needed: any architecture that presupposes processing by the victim inherits the failure it is meant to resolve. The second is that the output be a single directive, not a set of options. Presenting alternatives reactivates deliberation — the function that has been suppressed — and re-introduces milling latency by that route. The third is that the action required from the user be exclusively locomotor — following a line — because movement toward an already-determined destination is compatible with residual motor control, while choosing the destination is not.

The infrastructure developed within the SafeWay project is built precisely to this specification. The system absorbs the variables — the ongoing threat, topography, hazard physics — processes them, and returns a single safe point and a single route, computed from the real physics of the terrain around the user. The product communication's distinction between a determination and a suggestion is not a rhetorical choice: it is the operational translation of the neuro-cognitive constraint. A suggestion requires evaluation, and evaluation is not available; a determination requires only execution.

In functional terms, the device acts as an external surrogate of the user's prefrontal cortex: it executes the contention scheduling the human can no longer perform and returns a singular, unambiguous output that requires no logical analysis to be followed. By substituting spatial decoding with a single escape line, the system intervenes at the exact point where delay is generated — the milling-and-confirmation-seeking phase — and reduces the probability of triggering the neuroceptive signals that precipitate tonic immobility. The conversational agent SURV.AI extends the same logic to the subsequent phase, maintaining interaction within the single sequential instruction format that the dispatch literature identifies as the only one viable under shock.

The expected effect, re-read through Table 4, is an intervention on the subtrahend. The system does not lengthen technical lead time, which depends on detection networks; it compresses behavioural latency, which is the dominant term in the balance and the one on which training and information have been shown unable to act. It is in that compression that the residual survival margin is generated.

10. Conclusions

Decisional paralysis in the face of sudden catastrophe — whether manifesting as preparatory parasympathetic freezing mediated by the amygdala-PAG circuit, or as tonic immobility induced by dorsal vagal collapse — is not a character weakness or an individual preparedness deficit. It is a phylogenetically rooted, subconscious mechanism, inscribed in the neural architecture of the species, that is for the majority of people unavoidable.

Institutional confidence in the capacity of training and information programmes to eradicate this vulnerability in the civilian population is in contradiction with the neurochemical and behavioural evidence. Overwhelming stress acutely inhibits prefrontal synaptic connectivity, degrades working memory, attention, and logical capacity, and renders learned instructions effectively irrecoverable in the field. In the absence of the possibility of overlearning-conditioning at population scale, the civilian mass will remain exposed to Dysexecutive Survivor Syndrome and to the pre-evacuation latencies it generates.

The specific contribution of this report is to have placed that clinical conclusion alongside the numbers of operational alerting. Subtracting behavioural latency from technical lead time returns, for near-field tsunami, fast-catchment flash flood, and first-event outbreak tornado, a negative residual margin. The warning system works; the recipient, during the minutes for which the warning is valid, is unavailable to receive it in the way the system presupposes.

A reorientation of investment priority follows. Once detection times have been reduced to near the physical limit, the further marginal gain is modest; the margin still available resides entirely in the compression of human delay. Because that compression is not achievable through education outside of overlearning regimes inapplicable at scale, the only practicable path is externalising the decisional process to a computational architecture immune to neurochemical shock. By absorbing the complexity of environmental computation and converting a chaotic alarm into a single vector directive, such a system does not improve human reaction: it bypasses it. That is, given the present state of knowledge, the only known means of restoring to individuals caught in a natural disaster a concrete mechanical probability of reaching safety before the environment forecloses every escape.

Bibliography

A. Neuroscience, Survival Psychology, Emergency Behaviour

[A1] Reducing cognitive load in dispatcher-assisted CPR: a simulation-based educational study among non-medical university students. PMC12955211.

[A2] Leach J., Survival Psychology: The Won't to Live.

[A3] The Psychology of Survival — Part 1, Backcountry Survival.

[A4] History of survival psychology. Nomad Bushcraft.

[A5] Human Responses to Disasters: A Pilot Study on Peritraumatic Emotional and Cognitive Processing.

[A6] Leach J., Why people 'freeze' in an emergency: temporal and cognitive constraints on survival responses. PubMed 15198281.

[A7] Robinson S. J., Higgins T. D., Human behaviour in emergency situations, School of Psychology.

[A8] Leach J., Maladaptive behavior in survivors: dysexecutive survivor syndrome. PubMed 23316543.

[A9] Maladaptive Behavior in Survivors: Dysexecutive Survivor Syndrome (full text).

[A10] Analysis of influencing factors in pre-evacuation time using Interpretive Structural Modeling.

[A11] Evacuation behaviors and emergency communications: an analysis of real-world incident videos. PMC10620751.

[A12] Study of critical fire safety parameters and development of an evacuation strategy for high rise buildings.

[A13] Movement on Stairs During Building Evacuations. NIST TN 1839.

[A14] Modeling and Simulation of Crowd Pre-Evacuation Decision-Making in Complex Traffic Environments. PMC9779157.

[A15] 1985 Manchester Airport disaster — event summary.

[A16] Air Accidents Investigation Branch, Aircraft Accident Report 8/88, G-BGJL.

[A17] 40 Years On: The Air Tragedy That Changed The Way We Fly Forever.

[A18] Galea E. et al., Numerical investigation of the fatal 1985 Manchester Airport B737 fire. Greenwich Academic Literature Archive.

[A19] Cloudberg A., Fire on the Runway: the Manchester Airport Disaster and British Airtours flight 28M.

[A20] How tragedy at Manchester airport saved passengers in Tokyo plane fire. The Independent.

[A21] Prefrontal cortex executive processes affected by stress in health and disease. PMC5756532.

[A22] Arnsten A. F. T., Stress signalling pathways that impair prefrontal cortex structure and function. PubMed 19455173.

[A23] Stress and Exploitative Decision-Making. PMC5647765.

[A24] Arnsten A. F. T., Stress weakens prefrontal networks: molecular insults to higher cognition. PubMed 26404712.

[A25] Decision-making under stress: a psychological and neurobiological integrative model. PMC11061251.

[A26] Locus coeruleus–amygdala circuit disrupts prefrontal control to impair fear extinction. PMC12993942.

[A27] Chronic stress weakens connectivity in the prefrontal cortex. PMC8408896.

[A28] Molecular mechanisms of stress-induced prefrontal cortical impairment. PubMed 18685145.

[A29] The central amygdala to periaqueductal gray pathway. PMC6292805.

[A30] Synaptic correlates of fear extinction in the amygdala. PMC2847017.

[A31] Dorsal periaqueductal gray–amygdala pathway conveys innate and learned fear responses. PMC3767534.

[A32] Neural Correlates of Fear in the Periaqueductal Gray. PMC5157112.

[A33] Role of periaqueductal gray neurons in parasympathetically mediated fear bradycardia. PubMed 27335434.

[A34] The cortical connectivity of the periaqueductal gray. PMC5332157.

[A35] Fear bradycardia and activation of the human periaqueductal grey. PubMed 23110885.

[A36] Freeze for action: neurobiological mechanisms in animal and human freezing. PMC5332864.

[A37] Role of periaqueductal gray neurons in fear bradycardia (full text). PMC4923232.

[A38] Kozlowska K. et al., Fear and the Defense Cascade: clinical implications and management. PMC4495877.

[A39] Translational Challenges With Tonic Immobility. PMC3237674.

[A40] Exploring Human Freeze Responses to a Threat Stressor. PMC2489204.

[A41] Shame, tonic immobility and reactions to stressful events as phylogenetically conserved submissive defence mechanisms. PubMed 41021541.

[A42] Tonic immobility differentiates stress responses in PTSD. PMC5102644.

[A43] Tonic immobility during sexual and physical assaults produces stronger memory effects. PubMed 36933228.

[A44] Polyvagal Theory: current status, clinical applications, and future directions. PMC12302812.

[A45] Polyvagal Theory Explained Simply: Understanding Your Nervous System.

[A46] Why Your Body Freezes Under Stress: A Polyvagal Guide.

[A47] Porges S. W., Polyvagal theory: a journey from physiological observation to neural innervation and clinical insight.

[A48] Porges S. W., Polyvagal Theory: A Science of Safety. PMC9131189.

[A49] Polyvagal theory: from physiological observation to clinical insight (full text). PMC12479538.

[A50] Porges S. W., The vagal paradox: a polyvagal solution. PMC10724739.

[A51] Leach J., Survival Psychology (compressed edition).

[A52] Leach J., Cognitive paralysis in an emergency: the role of the supervisory attentional system.

[A53] Stress Reactivity and Cognitive Performance in a Simulated Firefighting Emergency.

[A54] The Effects of Stress Exposure on Prefrontal Cortex: translating basic research into treatments for PTSD.

[A55] Quantifying Cognitive Load of Emergency Dispatchers. AEDR Journal.

[A56] Application of human factors to improve usability of clinical decision support for diagnostic decision-making. PMC7490974.

[A57] Decision-Making During High-Risk Events: a systematic literature review. PMC10564111.

[A58] Cognitive-Aware AI in Emergency Response: enhancing human decision-making through adaptive real-time alerts.

B. Operational Sources: Lead Times and Evacuation Behaviour

[B1] National Weather Service / NOAA, NTHMP Tsunami Information Guide.

[B2] PTWC / NOAA, documentation on seismic detection times and tsunami fatality distribution.

[B3] NOAA NCEI, Tsunami Travel Time Maps.

[B4] NOAA / Department of Commerce OIG, Review of the National Weather Service's Actions — GPRA targets and Flash Flood Warning lead-time series.

[B5] A review of the entire warning system, from prediction and detection to response. Bulletin of the American Meteorological Society, 94(11).

[B6] AMS 23rd Conference on Severe Local Storms, analysis of tornado warning lead time and probability of detection.

[B7] NOAA, tornado warning lead-time and POD historical series 1986–2004.

[B8] NOAA Performance Dashboard, Severe weather warnings for flash floods: lead time.

[B9] National Weather Service, Calculation of Lead Time for the Flash Flood 4-Panel technique (TA 2502).

[B10] Department of Commerce OIG, Independent Evaluation of NOAA/NWS Tornado Forecasting and Warning Services.

[B11] NOAA, Examining Subdaily Tornado Warning Performance and Associated Factors.

[B12] Weathernews, Tsunami Survey — evacuation onset times among survivors and fatalities.

[B13] Tsunami evacuation: lessons from the Great East Japan earthquake. GNS SR 2012-017.

[B14] Population vulnerability to tsunami hazards informed by post-event surveys (American Samoa, 2009).

[B15] Natural Hazards and Earth System Sciences, 9, 1075–1086 (2009) — Individual Decision Time.

[B16] Least-cost tsunami evacuation modelling. NHESS Discussions.

[B17] Proceedings of the Eastern Asia Society for Transportation Studies, vol. 14 (2023) — Evacuation Preparation Time.

[B18] Sene K., Flash Floods: Forecasting and Warning.

[B19] Large-scale geolocation data reveal evacuation behaviour during the 2024 Noto Peninsula earthquake and tsunami. Communications Earth & Environment.

[B20] Are current tsunami evacuation approaches safe enough? Stochastic Environmental Research and Risk Assessment.

This report was prepared in support of the SafeWay project (PS 1978 Limited). Lead-time and evacuation latency data are drawn from documentation and literature of official bodies (NOAA/NWS, USGS, PTWC, FEMA); the neurobiological and psychological apparatus is drawn from the peer-reviewed literature cited in the bibliography. Contact:
baccaro.lab@thesafeway.it