

The P-index Framework

A Physics-Based System for Crowd Compression Risk Quantification

Commercial White Paper, Version 1

PneumaTheorem, Inc.

Delaware C-Corporation, incorporated 2026

Kevin Tao, Founder & CEO

Companion academic publication.

This document accompanies the open-access manuscript:

*“A Physical Quantification Framework for Crowd Compression Risk:
The P-index System and Its Contact-Mechanical Foundation,”*

arXiv: [arXiv ID pending -- to be inserted upon release]

The arXiv manuscript contains the complete physical derivation, validation methodology, and bibliography. The present white paper presents the commercial framing and engagement architecture.

PneumaTheorem, Inc.

Legalinc Corporate Services Inc.

131 Continental Dr, Suite 305

Newark, DE 19713, United States

Inquiries: kevin@pneumatheorem.org

Web: pneumatheorem.org

ORCID (Founder): 0009-0005-8136-3886

Confidential — prepared for institutional review.

Executive Summary

The P-index is a single dimensionless quantity that resolves a long-standing failure of crowd safety assessment: density-based indicators cannot distinguish safe high-density scenarios from lethal ones. Tokyo’s Yamanote Line operates routinely at 6–7 persons/m² without recorded compressive-asphyxia incidents; the 2022 Itaewon disaster occurred at 8–10 persons/m² and produced 159 fatalities. The two scenarios are not separated by density. They are separated by physical quantities that density does not measure.

The P-index measures those quantities directly. It is defined as the cumulative seconds, per person-hour, during which individuals in a venue are subject to above-threshold contact pressure. The number is grounded in four mechanical inputs — the time–force lethality envelope from biomechanical literature, the direction–intensity coefficient C (the framework’s central physical primitive), the granular force-chain transmission efficiency β , and the individual force-chain depth L_i — combined through a closed-form three-factor expression. Across the Yamanote Line, the London Underground, and the Itaewon incident, the resulting P-index values differ by more than six orders of magnitude, providing the physical separation that density-based methods cannot.

For the insurance industry, the framework delivers three outputs that map directly into existing catastrophe-modeling workflows: P_{baseline} as the underwriting threshold, P_{max} as the base-rate driver, and a full Exceedance Probability (EP) curve for reinsurance pricing. The output format is identical in structure to the EP curves used by RMS, Verisk/AIR, and KatRisk for earthquake and hurricane risk. The joint uncertainty of the framework is $\sigma_{\log_{10} P} = 0.31$, corresponding to a 95% confidence interval of approximately 0.5× to 2× the point estimate — already within the precision band of mature catastrophe models, with a clear pathway to further compression through subway-commuter calibration data.

For regulators, the framework supplies N_{max} , a venue capacity limit defined as the phase-transition point at which P_{baseline} ceases to be zero. N_{max} is a physics-based ceiling, not an engineering heuristic derived from fire-evacuation time. It can supplement or replace per-person area rules under NFPA 101, ISO 16730, and Eurocode performance-based standards.

For venue operators, the framework supports four operational use cases: iterative evaluation during design, real-time risk monitoring during operation, quantified retrofit prioritization for existing venues, and cumulative-risk evaluation in evacuation planning. The sub-region P-index P_{Ω} provides spatially resolved risk heatmaps suitable for both operational dashboards and post-incident liability documentation.

The framework adopts a *specification-based rather than implementation-based* design, consistent with NFPA 101, ISO 16730, Solvency II, and Eurocode regulatory traditions. It specifies what physical quantities a simulator must output and what consistency tests it must satisfy, not which simulator must be used. Any contact-mechanical engine capable of producing individual-level (F_i, C_i, L_i, τ_i) trajectories qualifies as a compliant back-end. This design ensures technology neutrality, vendor competition, and long-term compatibility.

The P-index framework, the C-value formulation, and the non-repulsive computational archi-

structure on which they are implemented are held by PneumaTheorem, Inc. (Delaware, 2026). Engagement architectures range from standard commercial licensing through to bespoke strategic arrangements, described in Section 6.

Contents

Executive Summary	2
1 The P-index in One Page	6
1.1 The problem: density has failed	6
1.2 The P-index: what it measures	6
1.3 Why this is now possible	7
1.4 What this document presents	7
2 The C-value as Physical Primitive	8
2.1 Definition and role within the P-index	8
2.2 Why the C-value must be allowed to exceed unity	8
2.3 C-value as order parameter: the secondary-indicator argument	9
2.4 Predictability of C-value direction	9
2.5 Semi-binary treatment of psychological factors	10
3 The Non-Repulsive Computational Architecture	11
3.1 Why the social-force lineage cannot deliver the C-value at scale	11
3.2 The FDS+Evac solution path and its envelope	11
3.3 Physical foundation: contact forces only	12
3.4 The two-layer agent geometry	12
3.5 Force network and conduction	12
3.6 Physical status of the C-value within the framework	13
3.7 Computational stability and operating range	13
4 Validation Status	14
4.1 Four independent validation routes	14
4.2 The subway-commuter calibration pathway	15
4.3 Status with respect to fitted parameters	15
5 Applications	17
5.1 Insurance actuarial practice	17
5.2 Regulatory certification	18
5.3 Venue design and operation	19

- 5.4 Methodological extensions 21
- 6 Engagement Architecture 22**
 - 6.1 Engagement structures 22
 - 6.2 Engagement process 23
 - 6.3 Counterparty profile 23
 - 6.4 Pricing posture 24
- 7 Disclosure Boundary, Framework Positioning, and Verification 25**
 - 7.1 Disclosure boundary 25
 - 7.2 Framework positioning 25
 - 7.3 Verification pathway 26
 - 7.4 Closing 26
- References 27**

1. The P-index in One Page

1.1 The problem: density has failed

Crowd safety regulation worldwide rests on density-based indicators — persons per square meter, per-person floor area, occupancy ceilings derived from fire-evacuation time. These indicators were developed in the 1970s, primarily by Fruin, and remain embedded in NFPA 101, Eurocode, and equivalent codes globally.

They have failed at the physical level. Three reference points establish the failure:

- Tokyo Yamanote Line, peak hour: 6–7 persons/m², zero recorded compressive-asphyxia incidents over decades of operation.
- London Underground, Victoria Line peak: 5–6 persons/m², zero recorded incidents.
- Seoul Itaewon, 29 October 2022: 8–10 persons/m², 159 fatalities in a single evening.

A regulatory or actuarial framework that classifies these three scenarios on density alone produces systematic misclassification at both ends of the band. The failure is not statistical (better data fits will not resolve it). It is structural: density measures the instantaneous occupancy state of a venue, whereas the physical basis of compressive asphyxia is the cumulative contact-force exposure sustained by an individual. These two quantities are not equivalent.

1.2 The P-index: what it measures

The P-index resolves the failure by measuring what density does not. It is defined as:

$$P_{\text{venue}} = \frac{\sum_{i=1}^N \max(0, \tau_i^{\text{max}} - \tau_{\text{crit}})}{N \cdot T_{\text{window}}}$$

where τ_i^{max} is the maximum cumulative seconds during which individual i is subject to above-threshold contact pressure within the observation window, $\tau_{\text{crit}} = 30$ s is the lower-bound threshold corresponding to minor injury, N is the total number of individuals in the venue, and T_{window} is the observation window in seconds.

The numerator is the sum of individual-level excess exposure. The denominator is total person-hours. The ratio is a dimensionless quantity interpretable as “the cumulative seconds, per person-hour, during which individuals are subject to above-threshold compression.” P-index equal to zero means no individual has accumulated above-threshold exposure. P-index greater than zero means at least one individual has crossed the risk regime.

Applied to the three reference scenarios:

Scenario	Density	P-index (baseline)
Tokyo Yamanote, peak hour	6–7 persons/m ²	0
London Underground, peak hour	5–6 persons/m ²	$\sim 8.3 \times 10^{-6}$
Itaewon, incident interval	8–10 persons/m ²	$\sim 10^{-3}$ to 10^{-2}

The P-index values differ by at least six orders of magnitude across scenarios whose densities are within a factor of two of one another. The framework supplies the physical separation that density-based methods structurally cannot.

1.3 Why this is now possible

The physical inputs required to compute the P-index — Hertzian contact mechanics, Coulomb friction, granular force-chain physics, and the time–force biomechanical lethality envelope — have been established in their respective literatures for decades. Hertz dates to 1882. Coulomb friction is older. The Cates–Wittmer–Bouchaud–Claudin force-chain framework was published in 1998; Majmudar and Behringer’s contact-force measurements appeared in 2005; Kroll, Still and Griffin’s biomechanical compression-lethality model was published in 2017.

What had not been done, prior to the present framework, was the assembly of these components into a single individual-level cumulative-exposure metric, the formal specification of the four physical quantities a simulator must output to compute that metric, and the alignment of the output format with insurance-industry catastrophe-modeling pipelines. The P-index framework is the assembly.

1.4 What this document presents

This white paper presents the commercial framing of the P-index framework. Section 2 introduces the C-value, the central physical primitive on which the P-index is built, and explains why it is open above unity — a structural requirement that distinguishes the framework from prior crowd-dynamics formulations. Section 3 describes the non-repulsive contact-mechanical architecture on which the framework is implemented, and contrasts it with the social-force lineage including the FDS+Evac solution path.

The complete physical derivation, the four-route validation, and the full bibliography are presented in the companion arXiv manuscript referenced on the title page. Readers conducting due diligence are encouraged to consult both documents together: the present white paper for the commercial and applied frame; the arXiv manuscript for the physical foundations and quantitative validation.

2. The C-value as Physical Primitive

2.1 Definition and role within the P-index

The C-value is the central physical primitive on which the P-index is computed. It is defined as the ratio of an individual’s axial push force — the component of contact force aligned with the local primary force-chain axis — to a physiological baseline:

$$C_i(t) = \frac{f_{\text{axial},i}(t)}{f_{\text{baseline}}}, \quad f_{\text{baseline}} = 200 \text{ N.}$$

The baseline corresponds to the push force exerted by a healthy adult under purposeful, brisk motion without panic — a value derived from human-factors literature on door-push and transfer-walking forces.

The C-value enters the P-index through the closed-form three-factor decomposition of the effective contact load:

$$F_i = f_{\text{baseline}} \cdot \tilde{C} \cdot \frac{1 - \beta^{L_i+1}}{1 - \beta},$$

where \tilde{C} is the upstream-weighted equivalent C-value along the contact chain terminating at individual i , $\beta = 0.978$ is the force-chain transmission efficiency, and L_i is the force-chain depth. F_i is then compared against the time-dependent biomechanical threshold $F_{\text{thr}}(\tau)$ from Kroll, Still and Griffin (2017) to determine whether the individual is in the over-threshold regime, and the exposure timer τ_i accumulates accordingly.

The P-index is the venue-level aggregate of these individual exposures. The C-value is the single physical quantity that, more than any other input, determines whether a venue’s P-index is zero or non-zero.

2.2 Why the C-value must be allowed to exceed unity

A structural property of the C-value, and one that distinguishes the framework from prior crowd-dynamics formulations, is that its denominator is a fixed physiological baseline rather than a maximum. The C-value is therefore not restricted to the interval $[0, 1]$. Under panic conditions, directional convergence, or bidirectional compression, $f_{\text{axial},i}$ can substantially exceed 200 N, and the C-value passes through unity into the regime relevant to compression incidents.

C-value	Physical state	Representative scenario
0.2–0.4	Sparse contact	Subway car under normal flow
0.5–0.8	Routine egress	Scheduled dispersal, commuter flow
≈ 1.0	Standard emergency evacuation	General fire drill
1.5–2.0	Panic with directional convergence	Hillsborough 1989, Love Parade 2010
2.0–2.5	Extreme bidirectional compression	Itaewon 2022, Mecca 2015

Any model that constrains the C-value (or an equivalent push-intensity coefficient) to $[0, 1]$ is

structurally incapable of reproducing the incidents listed in the lower rows. The constraint is not a calibration choice that can be tuned; it is a structural property of the model that excludes the lethal regime from the model's representable state space. Such a model can fit normal-scenario data within $C \leq 1$ and cannot extrapolate to the regime in which insurance and regulatory decisions are made.

The framework's open design of the C-value above unity is therefore a necessary condition — not a preference — for admission into actuarial workflows that price catastrophe risk. Field-type repulsion models, by construction, cannot reach this regime.

2.3 C-value as order parameter: the secondary-indicator argument

A natural methodological objection to any single-parameter risk indicator is that no single quantity can capture multi-dimensional risk. This subsection addresses the objection directly.

The C-value is not a composite metric or a weighted aggregate. It quantifies the directional coherence of inter-agent contact forces — the degree to which forces across the crowd align along a common axis. Once this directional coherence is established, the conventional secondary risk indicators acquire direct causal links to incident occurrence:

- **High C-value combined with high density produces incidents.** The aligned direction converts density into cumulative pressure along a coherent axis; compression follows.
- **High C-value combined with elevated individual exertion produces incidents.** Force contributions act along the established direction and aggregate rather than cancel; the crowd transmits and amplifies localized exertion.
- **High C-value combined with panic response produces incidents.** Panic concentrates motion along the coherent direction, raising both density and exertion simultaneously along the same axis.

Below the critical C-value, the same indicators may take arbitrarily high values without producing incidents. High density without directional coherence remains a dispersed configuration; high individual exertion without alignment dissipates into cancelling vectors; panic without a coherent direction produces disorder but not compression. The causal weight of each secondary indicator is therefore not a fixed property of that indicator; it is governed by the C-value regime in which the indicator is measured.

This is the structural sense in which the C-value functions as an *order parameter* in the statistical-physics sense. It governs the phase in which the system operates. Below a critical value, the system is in a dispersed phase in which secondary indicators do not couple to incident outcomes; above the critical value, the system is in a coherent phase in which they do.

2.4 Predictability of C-value direction

A second operational property is essential. The C-value's direction is predictable from configuration alone, without real-time measurement of individual agents. Crowds respond to instruction, move away from perceived hazards, and converge toward exits. These response patterns are

documented across decades of evacuation research and incident forensics. The direction along which contact forces will align in any given venue can therefore be anticipated from spatial layout, hazard location, and instruction points.

This predictability is the basis for an entire design axis described in later sections: intervention configuration — the placement of guidance lighting, directional signage, evacuation guidance systems, and audio cueing — can be evaluated as a primary safety design variable, with building geometry held fixed. The C-value provides the metric against which intervention configurations are scored.

2.5 Semi-binary treatment of psychological factors

A persistent methodological problem in crowd dynamics is the treatment of psychological variables — panic, agitation, urgency. Traditional approaches model these as continuous coefficients (“panic coefficient 0.7”) multiplied into kinematic quantities. Two structural problems follow:

First, psychological variables are not externally observable. “An individual’s level of panic” cannot be measured at an incident site; it can only be inferred retrospectively from behavior. The coefficient is circularly defined: panic explains behavior, behavior defines panic.

Second, continuous parameters lack calibration anchors. The difference between “panic 0.7” and “panic 0.8” cannot be externally validated. The coefficient becomes a free parameter of model fitting.

The framework circumvents both problems through a semi-binary switching treatment. Two anchor distributions are established: C_{normal} , calibrated from large-sample subway-commuter data (where “not panicking” is operationally well-defined by the absence of incidents over 10^9 -scale ride observations), and C_{panic} , calibrated from contact-mechanical reconstruction of historical incidents. The physical impact of panic is then defined as the difference between two measurable distributions:

$$\Delta C_{\text{panic}} = \langle C_{\text{panic}} \rangle - \langle C_{\text{normal}} \rangle.$$

This difference is itself a measurable quantity. The framework does not need to answer the unmeasurable question “what is this individual’s level of panic”; it answers the measurable question “does the C-value distribution in this scenario belong to the normal regime or the panic regime.”

The natural commercial consequence is the dual-layer output P_{baseline} (using C_{normal}) and P_{max} (using $C_{\text{normal}} + \Delta C_{\text{panic}}$), corresponding to underwriting threshold and base-rate driver respectively.

3. The Non-Repulsive Computational Architecture

3.1 Why the social-force lineage cannot deliver the C-value at scale

The social-force model, introduced by Helbing and Molnár (1995) and developed through extensive subsequent literature, derives crowd dynamics from a repulsive field acting between agents. Within its design envelope — low-to-moderate density, evacuation flow studies, macroscopic statistics — the formulation has been highly productive and remains the dominant tool in academic crowd dynamics.

The formulation encounters three structural costs at the densities and configurations relevant to compression incidents.

Cost 1: numerical instability at compression-relevant densities. Above approximately four agents per square meter, the superposition of repulsive fields produces high-frequency oscillation in agent velocity and orientation. Computation of directional coherence on oscillating velocity fields requires temporal averaging windows that exceed the timescale of incident onset. A C-value computed under these conditions reports a smoothed quantity that lags the configuration it purports to measure.

Cost 2: the freezing-by-heating regime. Helbing, Farkas and Vicsek (2000) documented a regime in which increasing agent agitation — precisely the regime in which compression incidents are most likely — produces locked configurations: agents cease to move while continuing to exert force. Directional coherence in this regime is undefined: forces are present, velocity is zero, and the conventional definition of force direction loses operational meaning.

Cost 3: integration-fidelity dependency. The repulsive field reproduces avoidance through oscillation that must be integrated to recover the intended quantity. Reproducibility of any derived measurement — directional coherence among them — across computing environments with different floating-point behavior, different timesteps, and different rendering frame rates is not preserved at the level required for actuarial or regulatory use.

3.2 The FDS+Evac solution path and its envelope

Within the social-force lineage, the most complete response to these structural costs is FDS+Evac (Korhonen et al., 2010), which couples fire-dynamics simulation with evacuation modeling and provides extended formulations to manage the high-density regimes in which the standard social-force formulation becomes unreliable. FDS+Evac demonstrates that quantities analogous to the C-value can be computed within the social-force lineage when sufficient computational and methodological resources are committed. It represents the most complete solution path available within that lineage.

The engineering envelope of FDS+Evac, however, defines the cost of that completeness. Simulations require coupled CFD computation, expert calibration of agent and environmental parameters, simulation durations that scale with the resolution of the coupled fluid dynamics, and post-processing expertise to extract meaningful indicators from the coupled output. The resulting computational regime is suited to detailed pre-construction fire-safety analysis under

engineering supervision. It is not suited to actuarial pre-binding assessment, real-time monitoring, batch evaluation across many configurations, or the rapid-iteration design workflows that intervention-based safety planning requires.

The methodological question that follows is whether a different formulation can produce the C-value, and hence the P-index, without these structural costs.

3.3 Physical foundation: contact forces only

The framework’s answer is to take the contact force as the only inter-agent force, and to derive crowd dynamics from the network of contact forces directly. Between human bodies in a crowd, the physically present forces are contact forces. Bodies not in contact exert no force on one another; bodies in contact exert force along the contact direction, modulated by the soft-tissue compressibility of the contact region.

This is the same physical foundation invoked by the social-force model in the limit of physical contact (the Hertzian contact term in Helbing–Molnár). The framework differs in that it does *not* introduce a repulsive field to reproduce non-contact avoidance. Non-contact avoidance is instead derived from agent-level intentional motion: agents alter their intended trajectory in response to perceived obstacles. No oscillating auxiliary field is introduced; no integration-fidelity dependency is incurred.

3.4 The two-layer agent geometry

Agents in the formulation are represented by a two-layer geometric model. The outer layer is a visual envelope corresponding to shoulder width and the physiological extent of the body. The inner layer is a rigid incompressible core corresponding to the projected area of the skeletal structure and chest cavity. Visual envelopes may overlap; rigid cores may not.

This geometry corresponds to crowd compression physiology directly. The soft tissue of shoulder and torso compresses under crowd pressure within a bounded range; the chest cavity and skeletal structure do not compress. Compression incidents occur when configurations approach the rigid-core packing limit, at which point further crowd density is converted into force on the chest cavity rather than further packing. The two-layer model permits stable simulation up to the rigid-core packing limit — the empirical range of compression incidents — without the numerical instability of formulations in which agent radius and collision radius are identified as a single quantity.

3.5 Force network and conduction

Contact between rigid cores establishes an edge in the contact network. Force exerted at any node is conducted along edges to adjacent nodes, with conduction efficiency modulated by the physiological properties of the contact region (the parameter β introduced in Section 2, calibrated to $\beta \approx 0.978$ from biomechanical and granular-physics measurements).

The contact network is the substrate on which crowd dynamics are computed; the topology of the network, rather than the spatial distribution of agents, is the primary computational object.

This shift — from spatial distribution to network topology as primary object — is the structural difference between the framework and the social-force lineage.

The freezing-by-heating regime characteristic of social-force formulations at high agitation does not arise: locked configurations are not produced by superposition of fields, and the force network responds to agitation by reconfiguring rather than locking.

3.6 Physical status of the C-value within the framework

The C-value, defined as directional coherence of inter-agent contact forces, is a geometric property of the contact network. Each edge in the network carries a force vector measured directly at each integration step; the C-value measures the degree of alignment among these vectors. The quantity does not depend on integration fidelity of an oscillating field; its reproducibility across computing environments is preserved at the level required for actuarial and regulatory use.

The causal chain from C-value to incident, established in Section 2, is therefore grounded in directly measured quantities throughout. High C-value combined with high density produces directed pressure on rigid cores. High C-value combined with elevated exertion produces aggregating force along the coherent axis. High C-value combined with panic concentrates motion along the same axis. Each link in the chain is computed on directly measured contact forces.

3.7 Computational stability and operating range

The formulation has been demonstrated to remain stable across the density range relevant to compression incidents. Visual envelope density up to approximately five agents per square meter, and rigid-core density up to the physiological packing limit of approximately six to eight agents per square meter, are within the formulation's validated operating range. Deterministic reproducibility under fixed random seed is preserved across this range.

The validated operating range encompasses the regime in which social-force formulations enter numerical instability, and encompasses the regime in which compression incidents are documented to occur. The architecture's design envelope — larger agent counts, longer simulation durations, larger spatial domains — extends beyond the currently validated range, with extended validation documented as part of standard due-diligence procedures.

4. Validation Status

This section summarizes the framework’s validation status. The complete validation methodology, including the four physical-consistency tests, the public-dataset comparison protocol, the historical-incident back-casting, and the Monte-Carlo sensitivity analysis, is documented in the companion arXiv manuscript. The present section presents the conclusions in the form relevant to commercial and regulatory due diligence.

4.1 Four independent validation routes

The framework has been validated along four independent routes, designed to address the four distinct failure modes that a crowd-risk model can exhibit.

Route 1: physical-consistency tests. Four internal tests (A through D) verify the model’s algebraic self-consistency, the individual-level structure of the exposure timer τ_i , the stability of the contact-detection threshold F_{\min} , and time-step convergence. The pass criteria are quantitative (e.g., $\geq 95\%$ of (F_i, C_i, L_i) triples must satisfy the three-factor decomposition with relative error below 5%; the individual τ_i distribution must differ from a uniform redistribution by Kolmogorov–Smirnov $p < 0.01$). Any model claiming to output a framework-compliant P-index must attach pass certificates for all four tests to its report. This requirement corresponds to the standard “internal validation documentation” disclosure clauses of Solvency II for catastrophe models.

Route 2: public-dataset comparison. The framework’s low-to-moderate-density outputs are compared against four publicly accessible crowd-dynamics experimental datasets: BGU–Wuppertal corridor experiments (Boltes & Seyfried, 2013), Sieben exit experiments (Sieben et al., 2017), Wang & Weng marching-column force-transmission experiments (2018), and Feldmann push-transmission experiments (2023). The latter two datasets provide direct external validation of the force-chain transmission efficiency β . Median errors in F_i are below 15%; P99 errors are below 20%; Kolmogorov–Smirnov distances in τ_i distributions are below 0.15. All within engineering-acceptable ranges.

Route 3: historical-incident back-casting. Five historical compressive-asphyxia incidents with sufficient public records are reverse-engineered: Hillsborough 1989, Itaewon 2022, Love Parade 2010, Mecca 2015, and Astroworld 2021. The framework’s back-cast C-values align with scene characteristics: bidirectional compression (Itaewon, Mecca) corresponds to $C \geq 2.0$; single-direction heavy compression (Hillsborough, Love Parade) to $C \approx 1.5$ – 2.0 ; concert front-row (Astroworld) to $C \approx 1.3$ – 1.7 . A log–log regression of fatality count against $P \cdot N$ yields:

$$\log_{10}(\text{deaths}) \approx 0.95 \cdot \log_{10}(P \cdot N) + 0.3, \quad R^2 \approx 0.93.$$

The sample size of five is too small for a statistically robust regression, but is sufficient to establish order-of-magnitude correlation between the P-index and eventual casualty scale.

Route 4: cross-scenario extrapolation error. A Monte-Carlo sensitivity analysis quantifies the framework’s predictive uncertainty across three error sources: scene-dependent variation in

β , conversion of C between psychological regimes, and individual-tolerance variability ($\pm 20\%$). Each parameter is sampled 1,000 times within its calibration uncertainty interval. The joint uncertainty is:

$$\sigma_{\log_{10} P} = 0.31,$$

corresponding to a 95% credible interval of approximately $0.5\times$ to $2\times$ the point estimate — less than one order of magnitude. This precision meets the industry standard for catastrophe models (RMS, Verisk/AIR, KatRisk all operate within ± 1 order of magnitude on annual loss estimates).

4.2 The subway-commuter calibration pathway

The four routes above establish the framework’s current precision. The framework’s long-term precision trajectory is determined by a separate property: the contact-mechanical foundation that unifies routine commuter data with historical incident scenarios.

Subway commuter data ($\approx 10^9$ rides per year per major metropolitan system; Tokyo Metro $\approx 3.5 \times 10^9$, London Underground $\approx 1.4 \times 10^9$, New York Subway $\approx 1.7 \times 10^9$) and historical incident data share the same underlying contact-mechanical basis (Hertzian contact, Coulomb friction, Newton’s third law, β -attenuation). They differ only in the regime of the C -value distribution: commuter data populates C_{normal} ; incidents populate C_{panic} . The semi-binary treatment of psychological factors (Section 2) provides the formal mapping between the two regimes.

The strategic consequence is direct. The framework’s calibration precision can be improved monotonically through commuter-data accumulation, at a rate determined by commuter-data acquisition rather than by incident occurrence. Joint $\sigma_{\log_{10} P}$ is expected to compress from 0.31 to 0.10–0.15 as cross-city commuter databases are built. The 95% credible interval would correspondingly narrow from ± 1 order of magnitude to approximately ± 0.3 – 0.4 orders.

This positions the framework’s long-term precision ceiling structurally above that of any traditional risk model relying on incident statistics. In insurance terminology, this corresponds to the fundamental advantage of a calibrated physical model over a purely statistical fit. The validation precision presented in this section is therefore a floor, not a ceiling.

4.3 Status with respect to fitted parameters

The framework does not employ fitted parameters or fitted functional forms in the technical sense established in computational physics. All numerical values are anchored in published physiological, biomechanical, or physical sources, with linear extrapolation between anchoring points where direct measurement is not available. All functional forms are derived from mechanical first principles. The same parameter set is used across all validation scenarios.

This status is structural — it follows from the framework’s development methodology — and is verifiable through standard due-diligence examination of parameter sources. All physiological, biomechanical, and environmental parameters are exposed as user-adjustable inputs, with

adjustment scope bounded by the literature-anchored physical range; certification of outputs against this range is part of the engagement structure described in Section 6.

5. Applications

The P-index framework is designed to integrate directly with three classes of existing commercial and regulatory workflow: insurance actuarial practice, regulatory certification, and venue design and operation. This section presents each in turn. A unifying feature of all three is that the P-index output integrates as a drop-in component within established workflows; no restructuring of internal decision architectures is required.

5.1 Insurance actuarial practice

Industry context

Liability insurance for crowd-dense venues currently rests on two foundations: historical claims statistics (actuarial loss triangles) and qualitative engineering-consultancy assessment. The former is constrained by the rarity of compressive-asphyxia incidents (per-venue annual rate $< 10^{-3}$), yielding insufficient statistical power. The latter depends on expert judgment and lacks a reproducible quantitative basis.

A structural problem common to both: compressive-asphyxia fatalities are routinely classified into broader incident categories (fire, stampede, civil unrest), making liability attribution and rate-setting imprecise. Mature catastrophe models exist for earthquake, hurricane, and flood; no equivalent physics-based model for crowd risk has existed prior to the present framework.

Three-tier output aligned with insurance pricing

The framework delivers the three-tier output required by the insurance industry, corresponding to Solvency II’s standard disclosure requirements for catastrophe models.

Tier 1: P_{baseline} as the underwriting threshold. Any venue with $P_{\text{baseline}} > 0$ presents asphyxia-event risk under normal operating conditions and should be refused coverage or required to undergo structural modification. This criterion is binary (insure / do not insure) and corresponds to the traditional industry concept of “uninsurable risk.” The rigidity of $P_{\text{baseline}} = 0$ as a physical criterion (Section 2) means that the framework cannot be calibrated to produce $P_{\text{baseline}} = 0$ for a venue at which an incident later occurs — any such fit is exposed in post-incident back-casting, providing the auditor a clear record of pre-binding assessment integrity.

Tier 2: P_{max} as the base rate. For venues with $P_{\text{baseline}} = 0$, the base rate is computed from P_{max} (the worst-case P-index, sampled under $C_{\text{normal}} + \Delta C_{\text{panic}}$). Conversion to Expected Annual Loss follows the standard Monte-Carlo simulation:

$$\text{EAL} = \lambda_{\text{event}} \cdot N \cdot \int P(L) L dL,$$

where λ_{event} is the annual event rate, N is mean venue occupancy, and $P(L)$ is the probability density of loss per event. Industry-standard uncertainty loadings, operating costs, and profit margins are added to the EAL to yield the base premium.

Tier 3: full EP curve for reinsurance pricing. The framework’s Monte-Carlo output yields a complete Exceedance Probability curve, directly compatible with excess-of-loss reinsurance pricing. Reinsurers compute layer premiums from EP-curve loss quantiles at various return periods, consistent with existing earthquake and hurricane reinsurance pricing workflows. The EP curve is the standard interchange format between primary carriers and reinsurers; the framework outputs it natively.

Differentiated pricing through P_Ω

The sub-region P-index P_Ω provides a basis for differentiated pricing of composite venues (integrated stadiums, exhibition centres, shopping complexes). Insurers can price each functional zone by its own P_Ω and person-hour exposure, rather than applying a uniform venue-wide rate. This delivers two commercial benefits: risk-sensitivity matching (satisfying Solvency II’s premium–risk–correlation requirements) and economic incentive (encouraging operators to prioritize modification of high- P_Ω zones, thereby lowering overall P_{\max} and premium). The mechanism is logically identical to existing property-insurance zonal rating (e.g., the differential between coastal high-risk and inland low-risk zones).

5.2 Regulatory certification

Augmenting existing certification systems

The regulatory certification system for crowd-dense venues includes building fire codes (NFPA 101, Eurocode), event permitting, and structural safety certification. Existing systems use static occupancy limits (persons per square metre) and evacuation time as primary criteria, but lack a physical quantification tool for dynamic compression events.

The framework does not replace existing certification; it augments it. Alongside existing static criteria, $P_{\text{baseline}} = 0$ is added as a dynamic-risk criterion; N_{\max} is added as a physics-based occupancy ceiling (replacing or supplementing per-person area rules); and sub-region P_Ω becomes the basis for mandatory-modification requirements in high-risk zones.

N_{\max} as legal capacity

N_{\max} is defined as the phase-transition point at which P_{baseline} ceases to be zero. It carries three regulatory implications.

Physics-based capacity ceiling. Traditional per-person area rules (e.g., 0.46 m²/person, corresponding to about 2.2 persons/m²) are engineering estimates from fire-evacuation time, with no direct link to compressive-asphyxia risk. N_{\max} is instead a contact-mechanics-based capacity ceiling, directly corresponding to the physical guarantee “no asphyxia event under this occupancy.”

Legal-liability boundary. If actual occupancy exceeds N_{\max} and an incident follows, operator liability is unambiguous: exceeding the physics-based ceiling is itself a violation. This boundary is easier to substantiate in court than vague criteria such as “overcrowding” or “management

negligence.”

Alignment of insurance and regulation. Section 5.1 identified $P_{\text{baseline}} = 0$ as the underwriting threshold. If regulators adopt N_{max} as legal capacity, the venue’s legal capacity and insurable capacity automatically align, avoiding operational impasses from conflicts between the two.

Dynamic permitting

For concerts, festivals, and sporting events, traditional permitting focuses on “expected attendance” and lacks physical quantification of density evolution, flow design, and contingency planning during the event. The framework supports dynamic permitting: applicants submit full-event scenario simulations whose model outputs include P_{max} , the temporal evolution of $P_{\Omega}(t)$, the EP curve, and a comparison against N_{max} . The regulator decides whether to grant the permit, whether to require flow modifications, and whether to adjust attendance caps based on these outputs.

The benefits: quantitative review replaces subjective “too many people may be dangerous” judgment; liability clarification, since after permit issuance, if an incident occurs, the cause can be cleanly traced to simulation inaccuracy, execution deviation, or unforeseen factors; and improvement feedback, since post-event data can be compared against pre-event simulation to continuously calibrate the model.

Cross-jurisdictional portability

The framework rests on universal physical foundations — Hertzian contact theory, Coulomb friction, Newton’s third law — and does not depend on cultural or behavioral assumptions. The β and C parameters calibrated in one jurisdiction apply directly to venues in any other, without recalibration. This property benefits multinational venue operators (international stadium chains, multinational tour promoters) by allowing a single certification process to serve multiple jurisdictions, lowering compliance cost. For regulators, portability also means mature certification practice from other jurisdictions can be adopted, accelerating domestic system development.

5.3 Venue design and operation

Design stage: physical basis for structural decisions

In the new-venue design stage, architects and engineering consultants choose among multiple spatial layout options. Traditional decisions rely on heuristics and fire-evacuation calculations, lacking quantitative evaluation of dynamic compression risk.

The framework provides an iterative evaluation tool during design: each candidate layout is simulated, with P_{max} , the P_{Ω} heat map, and N_{max} as outputs. The design team adjusts the layout (exit positions, corridor widths, corner geometry, staircase design) to minimize risk.

Design variable	Principal effect on P_{\max}
Exit width	Each additional 1 m reduces pre-exit P_{Ω} by $\approx 30\text{--}50\%$
Corner radius	Smoothed corners lower local L_i by 2–4 layers
Stair–landing ratio	Landing length \geq stair horizontal span avoids pre-stair pile-up
Barrier placement	Improper barriers are a leading cause of high L_i (Hillsborough)

Quantitative relationships depend on venue type, flow direction, initial density, and other factors; the figures above are typical ranges expected within this framework. Detailed design-optimization analyses are delivered as commercial consulting outputs.

Operation stage: real-time risk monitoring

For operating venues, the framework integrates with real-time crowd-monitoring systems. Density and dynamic data from sensors (video analytics, pressure mats, inertial sensing) feed into the framework, which outputs the live $P_{\Omega}(x, y, t)$ heat map and short-horizon forecasts.

Two real-time applications follow. *Warning triggers*: when any sub-region’s P_{Ω} approaches a warning threshold (e.g., 50% of the venue’s P_{\max}), the system automatically triggers operator response — zonal evacuation, entry restriction, or flow re-routing. Traditional density-monitoring systems can only trigger *after* density exceeds limits; this framework triggers as the physical risk quantities (C , L_i) rise, advancing warning by tens of seconds to several minutes. *Post-incident analysis*: if an incident occurs, the recorded $P_{\Omega}(x, y, t)$ provides a precise spatiotemporal evidence chain for liability attribution, insurance claims, and design improvement.

Retrofit stage: quantified prioritization

For safety-retrofit budget allocation of existing venues, traditional practice relies on expert assessment and experiential ranking. The framework provides quantitative retrofit prioritization. A P_{Ω} heat-map analysis identifies the sub-regions contributing most to the venue’s total P-index. For each, feasible retrofit options (widening exits, adding flow separation, adjusting routes) are estimated together with their expected P-index reduction. Dividing retrofit cost by P-index reduction yields a “cost per unit of risk reduction,” by which retrofit priorities are ranked. This replaces subjective judgment with quantitative decision-making, maximizes risk reduction under fixed budget, and provides liability clarity if an incident occurs in a deprioritized zone.

Evacuation plan design

Traditional evacuation plans use “total evacuation time” as the design criterion, ignoring compression risk during evacuation itself. In emergency scenarios such as fire, evacuation can cause compressive-asphyxia incidents (the Station Nightclub fire of 2003 and the Kiss Nightclub fire of 2013 both involved fatalities from evacuation crush rather than from the fire itself).

The framework provides dynamic risk assessment during evacuation: simulating the full evacuation, outputting $P_{\Omega}(x, y, t)$ and the temporal evolution of N_{\max} . If at any moment P_{Ω} exceeds

threshold, the plan must be adjusted (additional exits, staggered evacuation, guided flow separation). This shifts evacuation-plan design from “minimize evacuation time” to “minimize cumulative risk during evacuation,” better aligning with the actual safety objective.

5.4 Methodological extensions

The non-repulsive computational architecture enables several methodological extensions that operate as natural development pathways from the validated core.

AI-driven scenario generation. The framework’s outputs are interpretable across diverse spatial layouts, agent populations, and environmental configurations without per-scenario re-validation. This cross-scenario transferability supports generative-model batch evaluation: a generative model produces simulation configurations at scale, and the physical substrate evaluates each on its own terms.

Active intervention design. Building geometry held fixed, the configuration of intervention devices — guidance lighting, directional signage, evacuation guidance systems, audio cueing — can be designed as a primary variable. The C-value’s predictability of direction (Section 2) provides the methodological basis: intervention devices act as instruction points, adjusting the predictable direction, adjusting the regime in which secondary indicators operate.

Vehicular traffic extension. The architecture’s separation between agent representation, force network, and environmental configuration permits module substitution. Replacing pedestrian agents with vehicular agents, and pedestrian environments with road networks, extends the formulation to vehicular traffic dynamics. The contact network becomes a traffic interaction network; the C-value becomes a traffic flow coherence indicator.

Parameter certification framework. The literature-anchored parameter architecture supports the development of a certification mechanism for simulation outputs. Outputs generated using parameters within the physically defensible range receive methodological certification; outputs using parameters outside this range require user-provided physical justification before certification. The certification mechanism is intended as the methodological reference point for downstream applications in actuarial assessment, regulatory compliance, and design verification.

Detailed treatment of these extensions, and their integration with the engagement structures of Section 6, is delivered as part of the corresponding commercial deliverables.

6. Engagement Architecture

The P-index framework, the C-value formulation, the non-repulsive computational architecture on which they are implemented, the full parameter specifications, the numerical thresholds governing physiological state transitions, the validated and adjustable parameter sets, the case-study materials developed through historical-incident reconstruction, and the certification framework governing methodological validation of simulation outputs are held by PneumaTheorem, Inc. (Delaware C-Corporation, incorporated 2026).

The framework and the methodology are methodologically inseparable. P-index measurement requires a computational environment free from auxiliary-field artifacts; this constrains valid implementations to a specific class of non-repulsive contact-mechanical architectures. The companion arXiv manuscript specifies what physical quantities a compliant implementation must output and what consistency tests it must pass; the implementation itself is held proprietarily.

6.1 Engagement structures

Two engagement structures are available. Both are open; counterparty selection between them depends on the nature of the intended deployment and the degree of independence sought.

Structure A: Full acquisition

Complete intellectual-property transfer, including the methodological framework, the computational implementation, the parameter certification authority, and all derivative implementation pathways constructible from the methodological foundation.

Suitable for counterparties seeking immediate strategic independence and complete control over downstream deployment. Acquisition includes transfer of:

- All rights to the P-index framework, the C-value formulation, and the non-repulsive computational architecture.
- The full algorithmic specification, source code, parameter datasets, and case-study reconstructions.
- The certification authority for downstream methodological validation.
- Founder transition support over an agreed period to ensure operational continuity.

Acquisition terms are bespoke and reflect the framework's pre-commercialization status. Inquiries through the channel below.

Structure B: Long-term collaboration / licensing

Structured engagement supporting partial licensing, time-limited exclusivity, field-of-use exclusivity, certification-rights arrangements, or combinations thereof.

Suitable for counterparties seeking lower initial commitment and progressive evaluation. The

collaboration framework is designed to remain flexible across the early-stage development period. Typical engagement components include:

- Non-exclusive or field-of-use exclusive licensing of the P-index methodology for actuarial, regulatory, or operational deployment.
- Joint validation projects on counterparty-specific scenarios (venue assessments, portfolio risk modeling, regulatory pilot programs).
- Methodological certification of counterparty-internal simulation outputs against the framework’s parameter and consistency standards.
- Co-development of extensions (vehicular traffic, robotic systems, sport-specific applications) within agreed scope.
- Multi-year strategic partnership with renewable terms.

Pricing and exclusivity terms reflect the framework’s current pre-commercialization status. Terms accessible during the present phase may not be reproducible at later stages of framework development and market exposure.

6.2 Engagement process

The engagement process is structured to support institutional due diligence without unnecessary procedural overhead.

Stage 1: Initial inquiry. Written technical or commercial inquiry through kevin@pneumatheorem.org. Inquiries are responded to within five business days. There is no application form, no scoring process, and no required pre-qualification. The companion arXiv manuscript, the present white paper, and the framework’s deterministic reproducibility under fixed random seed together provide a self-contained basis for initial counterparty evaluation.

Stage 2: Technical and commercial discussion. Following initial inquiry, an introductory call (typically 30–60 minutes) clarifies the counterparty’s intended deployment, the engagement structure preference, and the technical scope. Mutual non-disclosure arrangements can be put in place at this stage if proprietary discussion is required.

Stage 3: Due diligence. Counterparties conducting due diligence receive access to extended technical material under appropriate confidentiality terms. Standard due diligence covers framework independence (Section 7), parameter calibration sources, validation reproducibility on independent hardware, and the extended operating range beyond what is documented in the present material.

Stage 4: Engagement formation. Terms are formalized through bilateral agreement appropriate to the chosen engagement structure. PneumaTheorem, Inc. does not operate an open application process; engagements are established through direct discussion.

6.3 Counterparty profile

The framework is targeted at three classes of institutional counterparty:

Insurance and reinsurance carriers with crowd-density-related underwriting exposure — including event-cancellation insurance, public-liability insurance for venues, and reinsurance treaties incorporating crowd-risk components. The three-tier output (P_{baseline} , P_{max} , EP curve) integrates with existing catastrophe-model pipelines (RMS, Verisk/AIR, KatRisk) without restructuring of internal actuarial workflows.

National and supranational regulatory bodies responsible for venue safety certification, event permitting, and building-code development. The framework’s specification-based design (consistent with NFPA 101, ISO 16730, Solvency II, and Eurocode traditions) allows direct incorporation into existing performance-based regulatory standards.

Venue operators, design consultancies, and event organizers with portfolios of dense-crowd venues. Iterative design evaluation, real-time risk monitoring, quantified retrofit prioritization, and dynamic evacuation risk assessment together cover the full lifecycle of operational safety management.

Adjacent counterparties — transit operators (for the subway-commuter calibration pathway), autonomous-systems developers (for mixed pedestrian–vehicular environments), and academic institutions (for joint research engagements) — are also welcomed within the long-term collaboration structure.

6.4 Pricing posture

Specific pricing is not published in this document. Engagement terms vary substantially across the structures described above and across the scope of intended deployment. Inquiries are responded to with indicative terms appropriate to the counterparty’s profile and intended use; bilateral negotiation establishes the final terms.

The framework is in a pre-commercialization phase. Terms accessible to early counterparties — including pricing, exclusivity scope, and certification-authority arrangements — reflect this phase and are expected to evolve as the framework matures, additional validation accumulates, and the calibration pathways described in Section 4 are realized.

7. Disclosure Boundary, Framework Positioning, and Verification

This section establishes the boundary of methodological disclosure in this document, the positioning of the framework with respect to existing crowd simulation tools, and the verification pathway by which the framework’s independence may be confirmed.

7.1 Disclosure boundary

This document, together with the companion arXiv manuscript, presents the methodological foundation of the P-index framework, the C-value as its central physical primitive, and the non-repulsive computational architecture on which the framework is implemented. The level of disclosure is calibrated to establish the physical basis for the framework’s outputs and to support due-diligence evaluation of the methodology.

The complete algorithmic specification of the framework, the implementation details of its mechanisms, the specific numerical thresholds governing physiological state transitions, the full set of validated and adjustable parameters, the case-study materials developed through application of the framework to historical incidents, and the certification framework governing methodological validation of simulation outputs and downstream adoption pathways are retained as deliverables within the engagement structures described in Section 6.

The boundary between this document and those deliverables reflects the framework’s status as a pre-commercialization technical asset; the boundary is not a constraint on future engagement but a property of the present document. Counterparties under appropriate confidentiality arrangements have access to extended material as part of standard due-diligence procedures.

7.2 Framework positioning

The framework is positioned as a crowd simulation methodology with associated computational architecture. Its outputs — agent trajectories, force distributions, contact networks, the C-value as a derived order parameter, and the P-index as the venue-level risk aggregate — are computational results of the simulation methodology and do not constitute claims on prior art in crowd dynamics research, social-force modeling, fire safety engineering, granular physics, or computational fluid dynamics.

The framework was developed independently. No proprietary code, datasets, or methodology from existing crowd simulation tools — including but not limited to Pathfinder (Thunderhead Engineering), MassMotion (Oasys), FDS+Evac (NIST), and Viswalk (PTV Group) — was referenced, incorporated, or relied upon during the development of the framework’s algorithms, data structures, or methodological foundations.

The repulsive-force comparison module implemented for validation purposes follows Helbing–Farkas–Vicsek (2000) with published parameter values, used solely for comparative demonstration as documented in the companion arXiv manuscript.

The biomechanical lethality envelope follows Kroll, Still and Griffin (2017), cited as published

literature in the standard academic sense. The force-chain transmission framework draws on Cates, Wittmer, Bouchaud and Claudin (1998) and Majmudar and Behringer (2005), again as standard literature citation. None of these citations constitutes derivative-work relationship; the integration of these published results into a single individual-level cumulative-exposure metric, and the specification-based framework constructed around that integration, are original to PneumaTheorem, Inc.

7.3 Verification pathway

Independence of development is verifiable through standard due-diligence procedures, including examination of development records, version control history, software dependencies, and the absence of derivative relationships to the source code and data structures of the tools enumerated above.

The framework's deterministic reproducibility under fixed random seed provides an additional verification pathway: outputs can be reproduced on independent computing environments by parties conducting due diligence, confirming the framework's operation without reliance on the specific computing environment in which it was developed.

The extended operating range of the framework — agent counts, simulation durations, and spatial domains beyond those documented in the present material — can be evaluated through direct demonstration on independent hardware. Such demonstration is part of standard due-diligence procedures available to counterparties under appropriate confidentiality arrangements.

The companion arXiv manuscript is publicly accessible and provides the full physical derivation and the four-route validation methodology in academic-publication form. Independent reviewers can replicate the validation analyses described therein using the public datasets cited (BGU–Wuppertal, Sieben, Wang & Weng, Feldmann; database: ped.fz-juelich.de/db) and the historical-incident records cited.

7.4 Closing

This document presents the P-index as a physics-based quantitative indicator of crowd compression risk, the C-value as the order parameter on which the P-index is computed, the non-repulsive computational framework on which the C-value is defined, the validation status of the framework, the three principal application domains in which the framework is deployed, and the engagement structures under which the framework may be acquired or licensed.

Subsequent technical or commercial inquiries may be addressed to:

Kevin Tao, Founder & CEO

PneumaTheorem, Inc.

Delaware C-Corporation, incorporated 2026

Email: kevin@pneumatheorem.org

Web: pneumatheorem.org

ORCID: 0009-0005-8136-3886

References

The following references are cited in this white paper. The full bibliography (18 entries) supporting the framework’s physical derivation and validation is provided in the companion arXiv manuscript.

Boltes, M., & Seyfried, A. (2013). Collecting pedestrian trajectories. *Neurocomputing*, 100, 127–133.

Cates, M. E., Wittmer, J. P., Bouchaud, J.-P., & Claudin, P. (1998). Jamming, force chains, and fragile matter. *Physical Review Letters*, 81(9), 1841–1844.

Feldmann, S., & Adrian, J. (2023). Forward propagation of a push through a row of people. *Safety Science*, 164, 106173.

Fruin, J. J. (1971). *Pedestrian Planning and Design*. Metropolitan Association of Urban Designers and Environmental Planners, New York.

Helbing, D., & Molnár, P. (1995). Social force model for pedestrian dynamics. *Physical Review E*, 51(5), 4282–4286.

Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating dynamical features of escape panic. *Nature*, 407(6803), 487–490.

Korhonen, T., Hostikka, S., Heliövaara, S., & Ehtamo, H. (2010). FDS+Evac: An agent based fire evacuation model. In *Pedestrian and Evacuation Dynamics 2008* (pp. 109–120). Springer.

Kroll, M. W., Still, G. K., Neuman, T. S., Graham, M. A., & Griffin, L. V. (2017). Acute forces required for fatal compression asphyxia: A biomechanical model and historical comparisons. *Medicine, Science and the Law*, 57(2), 61–70.

Majmudar, T. S., & Behringer, R. P. (2005). Contact force measurements and stress-induced anisotropy in granular materials. *Nature*, 435(7045), 1079–1082.

Nolan, J. P., Soar, J., Cary, N., et al. (2021). Compression asphyxia and other clinicopathological findings from the Hillsborough stadium disaster. *Emergency Medicine Journal*, 38(10), 798–802.

Sieben, A., Schumann, J., & Seyfried, A. (2017). Collective phenomena in crowds—where pedestrian dynamics need social psychology. *PLoS ONE*, 12(6), e0177328.

Tao, K. (2026). A physical quantification framework for crowd compression risk: The P-index system and its contact-mechanical foundation. *arXiv preprint*, [arXiv ID pending].

End of document.

PneumaTheorem, Inc. — Confidential — Prepared for institutional review.