A Physical Quantification Framework for Crowd Compression Risk

Yixuan Tao*

July 30, 2026

Pneuma Theorem, Inc.
Delaware, United States

Abstract. Criteria for crowd safety are conventionally expressed as densities. What injures a person is a force, and the relation between the two is not a function of density alone: it depends on how many bodies lie along a connected path, on whether exertion is oriented so as to accumulate, and on whether the configuration permits accumulation at all.

The difficulty is sharper than a matter of precision. Sustained muscular exertion decays, and a body that ceases to hold itself upright does not stop delivering force but leans — transferring weight rather than effort, which is a larger contribution and one that does not fatigue. At a fixed density the borne load therefore rises with time. A criterion whose quantity is stationary while the quantity it stands for is not is not an imprecise relation but no relation.

This paper separates what follows from geometry and conservation from what must be supplied from outside. A crowd is treated as having two phases, distinguished by whether floor area for a recovery step remains available; the framework addresses the second, in which contact is continuous and a disturbance is transmitted rather than absorbed. There the borne load accumulates in closed form along a connected path, and accumulates at all only where a contributor delivers more than a receiving body absorbs — where that fails, each individual bears one neighbour's exertion at any depth and for any duration. Whether it holds is settled by the boundaries: an open region divides the alignment among the directions of travel it admits and a bounded run does not, and the difference is large enough to place the same people, at the same occupancy, on opposite sides of the condition.

What the relations bound is the length of a connected run rather than a number of people. The deliverable before an event is accordingly a partition spacing together with the load it guarantees — a statement of geometric impossibility rather than of statistical rarity — and after an event a trajectory of borne load, to which a reader applies whatever dose–response function they accept. No physiological value enters any derivation.

The account makes a prediction that criteria expressed in density do not: a dense crowd with no common target does not accumulate load at all. Everyday commuter crowding is that configuration, at an exposure large enough that a non-zero per-exposure risk would long since have been recorded as a statistical fact.

Keywords: crowd compression risk, force chain, fluid–solid transition, jamming, P-index, C-value, compressive asphyxia, venue capacity

*Pneuma Theorem, Inc.Correspondence: kevin@pneumatheorem.org. ORCID: 0009-0005-8136-3886.

1. Introduction

1.1 A criterion stated in a proxy

Criteria for crowd safety are conventionally expressed as densities: so many persons per square metre, admissible or not. What injures a person in a crowd is not a density but a force, and the relation between the two is not a function of density alone.

A note on where that convention comes from. The practical vocabulary of density grades descends from Fruin (1971, 1993), and it is worth being accurate about what is being objected to here. Fruin did not overlook the force. His later treatment discusses the load a compressed crowd delivers, expressed as a fraction of body weight, and identifies compressive loading rather than density as what does the harm.

What followed retained one half of that. Density admits direct measurement, survives being written into a rule, and can be checked on site by anyone; force does none of these. A field settles on the vocabulary it can operationalise, and the grades were operationalisable while the force discussion was not. The objection in this paper is therefore to what the convention kept, not to the author from whom it descends — and the question it raises is the one he raised first.

The force borne by an individual is the accumulated exertion of those upstream of them, and how much of it arrives depends on three things the density does not record. It depends on how many bodies lie along a connected path — a matter of the geometry of the space rather than of how many people are in it. It depends on whether the exertion is oriented so as to accumulate, which is not always so and is sometimes not so at all. And it depends on whether the configuration permits accumulation in the first place, which is a threshold phenomenon rather than a gradual one. Two configurations at equal density can differ in the load they deliver by a large factor, and a criterion expressed in the density will not distinguish them.

The third of these is settled by the boundaries, and it is worth saying so at the outset. The orientation of exertion is not a behavioural variable in the way it first appears. A region open on its flanks distributes its occupants among several directions of travel, and what any one of them delivers along a given axis is divided among those directions. A run bounded on both flanks admits one direction, and no such division occurs. The difference between the two is not a refinement: it is large enough that the same people, at the same occupancy, exerting the same amount, stand on opposite sides of the condition that decides whether load accumulates at all.

Two consequences follow, and both run against the intuition the density vocabulary encourages.

The first is that walls are not an aggravating factor. On the reading in which the hazard is a matter of density, a bounded approach is a place where a crowd happens to get denser and its boundaries make a bad situation worse. On the reading developed here they are not aggravating: they are the condition. Nothing about the occupants distinguishes the two cases.

The second is that a large space is not dangerous because it is large. In an open region the exertion is not aligned and the accumulation does not begin, whatever the extent. What a large space permits is that enough people be admitted to fill it, and whether they are is a question about admission rather than about the space. A narrow bounded run is unfavourable directly, and without any question of numbers.

§4 derives the factor by which the division occurs and §5 the condition it is compared against; both are properties of the built form, and neither requires an account of what the crowd intends.

The obvious remedy — to state the criterion in force instead — runs into a second difficulty. It requires knowing what force injures, and the quantitative human data in the range that matters for sustained crowd loading are sparse, of mixed evidential character, and in one instance traceable only to a source that has not been located. A framework that anchored its derivations on a value drawn from that literature would inherit its provenance, and any finding it produced could be contested by disputing the anchor rather than the mechanics.

This paper takes neither course. It separates what can be derived from geometry and conservation from what must be supplied from outside, derives the first as far as it will go, and concedes the second entirely.

1.2 Two phases

The separation is possible because a crowd is not one system across the range of densities at which it is observed. Below a geometric crossing an individual displaced from equilibrium can step, recover, and dissipate the disturbance; above it that step has nowhere to land, contact with neighbours is the only support available, and the disturbance is passed on instead of absorbed. The quantity that separates the two regimes is not the number of people per unit area but the floor area available for a step.

This paper treats the second regime and makes no claim about the first (§2). The restriction is less costly than it appears and it carries an explanation with it: models that represent contact by a bounded repulsive potential never jam and therefore cannot exhibit the mechanism at all, and controlled experiments are conducted below the crossing because the regime above it cannot ethically be produced in volunteers. That the mechanism described here is absent from both bodies of work is a property of their domains rather than evidence about this one.

1.3 What is claimed

The framework establishes the load a configuration can produce, how that load accumulates along a connected path, and the conditions under which the accumulation is irreversible. It establishes these from the geometry of the occupied space, the conservation of the population within it, and the monotonicity of the accumulating map.

It does not establish which loads injure. Every physiological value appears in one chapter (§7), is catalogued rather than adopted, and enters no derivation. Nor does the framework treat rescue capability as a premise, its own relations bounding the lateral access on which rescue depends; nor does it allow egress rate to carry weight in a capacity finding, that quantity being an unverified behavioural input whose failure transfers a shortfall upstream rather than merely withholding a benefit.

These three refusals are one requirement applied three times:

> An indicator must not depend on a quantity the framework does not possess.

The architecture that follows is what satisfying it costs. Its principal consequence is that the framework's deliverable before an event is not a risk figure but a demonstration that no accumulation can occur — a statement of geometric impossibility rather than of statistical rarity — and after an event is a trajectory of borne load, to which a reader applies whatever dose function they accept.

What a positive finding does not predict. A finding that a configuration permits accumulation is not a forecast that an incident will occur there. The two are easily conflated, and the distinction is worth stating before the derivations rather than after them.

Recorded events of this kind mostly proceed in one way. A configuration that already satisfies the standing condition is carried past saturation by a single strong external factor — something that arrives at once and gives a large number of people the same reason to move, or removes the space ahead of them faster than they can respond. The configuration is what makes the outcome available; the factor is what realises it. This framework treats the first and says nothing about the second.

Reaching saturation by accumulation alone, with no such factor, requires a conjunction that is severe and must hold together: a region occupied to contact over sufficient length, occupants who cannot recover room, and a state that is sustained rather than transient. That conjunction is uncommon, and its rarity is the reason configurations of this kind are ordinary while events of this kind are not.

The rarity does not weaken the finding; it separates two questions that are not the same question. Whether a configuration can carry the load is a property it holds whether or not the factor ever arrives, and it is answerable from the drawing. Whether the factor arrives is not a property of the geometry and is not answered here. And where the conjunction does occur, the outcome is not marginal: the length that carries the load carries it for everyone along it, which is a structural statement rather than an empirical one.

1.4 How this paper is arranged

Three quantities carry the analysis and answer different questions. The C-value (§4) is the input: what the scene supplies. The state and its rate (§6) are the process: what accumulates, and whether it is still accumulating. The exposure integral (§3) is the conclusion: what was borne, and for how long. They are presented in the order conclusion, input, process, because the form the conclusion must take is what dictates the treatment of the other two.

§2 establishes the domain. §3 constructs the exposure architecture and the certificate. §4 treats the exertion, §5 its accumulation along a chain, and §6 the operational readout. §7 catalogues the external values. §8 sets out the uses, §9 the relation to existing models and the limitations.

Two layers, and a test that separates them. All material outside §7 is derivation and carries no physiological value. The claim is testable:

> Delete §7. No equation in this paper changes, no proposition loses a premise, and no result requires restatement.

The separation is not a presentational convenience. It is what allows the framework to be used by parties who disagree about the values in that chapter, and it is the property a reader should check first.

What is declared, and where each declaration bites. The derivations rest on a reference body mass (International Commission on Radiological Protection, 2002) and a declared exertion scale. Those two suffice for every result this paper expresses as a force: the accumulation, the growth condition, the saturation, and the bounds that follow from them.

Expressing a result as a *length* requires one further declaration — the density at which bodies can be packed no closer — because a chain depth is a count of people and converting it to metres needs the spacing between them. That quantity is fixed by the plan area an upright body occupies, and it cannot be obtained from the first two: a mass does not determine an outline. It is therefore a third declaration rather than a consequence, and §5.1 states what it rests on.

The division is worth keeping in view, because the two kinds of result are put to different uses. A load in newtons, to be compared against a level, is a forensic output and needs two declarations. A partition spacing in metres, to be built to, is a design output and needs three.

The scale is obtained as the product of the reference weight and a friction coefficient, taken at a stated reference condition: dry athletic footwear on a hard smooth indoor floor, at the static value — the coefficient at which sliding begins, rather than the lower one that sustains it, since what bounds a person's exertion is whether their footing gives way and not how it behaves afterwards. The value declared is at the upper end of what that condition yields.

The scale is thereafter fixed. It is not a measurement of any particular venue and does not vary with one: a reader preferring a different reference condition rescales throughout, and the structure of every result is unchanged. What the choice does affect is the comparison of a result against an externally supplied level, since those are quoted in absolute units.

A reader in possession of these two declarations can check the arithmetic without accepting anything on the author's authority, which is the standard this paper sets itself and the reason the declarations are kept to two.

2. Domain of Application

2.1 Two phases

A sparse crowd and a dense crowd are not the same system observed at two values of a parameter. They differ in which mechanical interactions are available at all. In the sparse case an individual who is jostled can take a step; contact with neighbours is intermittent, and a disturbance is absorbed locally by the person who receives it. In the dense case neither is true: contact is continuous, and the step that would absorb the disturbance has nowhere to land. The disturbance is then not absorbed but passed on.

The quantity that separates these regimes is not the number of people per unit area. It is the floor area available for a step.

Definition 1 (Lateral free volume)**.** *For an individual whose feet occupy a support base at position x, let $S(x)$ denote the set of ground positions reachable by a single step from x, and let $O(t)$ denote the region occupied at time t by the support bases of other individuals. The* lateral free volume *available to that individual is*

$$\varphi(x,t) \;=\; \frac{\big|\, S(x) \setminus O(t) \,\big|}{a_{step}}, \tag{1}$$

where a_{step} is the ground area required by a single foot placement.

$\varphi \geq 1$ means at least one step is geometrically available; $\varphi < 1$ means none is. We take this as the phase criterion:

$\varphi \geq 1$	**fluid phase**
$\varphi < 1$	**solid phase**

Two conventions, and one thing the criterion does not claim. The occupied region $O(t)$ is taken to include the support base of the individual in question, so that ground already beneath their own feet is not counted as available for a new placement. Without that convention φ is systematically overstated, and the overstatement is largest exactly where the criterion is being applied.

The criterion is also a *necessary* condition for recovery rather than a sufficient one, and §2.1 should be read that way throughout. A recovery step requires more than unoccupied ground: it requires a landing area for the swinging foot, an interval in which to place it, and room for the lateral displacement of the trunk that accompanies it. Definition 1 counts only the first. It follows that $\varphi < 1$ implies recovery is unavailable, while $\varphi \geq 1$ does not imply that it is available. The framework uses the implication that holds, and §5.8 records the open question of whether the two conditions coincide.

The terminology is borrowed, and from where. A system that ceases to admit further approach, and that thereafter transmits load along connected paths rather than distributing it uniformly, is a jammed one in the sense established for granular materials by Liu and Nagel (1998); O'Hern et al. (2003). The phase language used here is taken from that literature rather than coined for this purpose, and the load-bearing structures of §5 are the force chains of Cates et al. (1998), whose existence in granular systems has been measured directly (Majmudar and Behringer, 2005).

The borrowing is of vocabulary and of the qualitative structure it names. It is not a claim that a crowd is a granular material in any stronger sense: the constituents here are self-propelled, the contacts are compliant, and the exertion has a direction that no grain possesses. What transfers is the observation that a system can cease to be compressible and that what follows is a different mechanics, which is the whole of what this chapter asserts.

Three things change together at the crossing, and it is worth stating them separately because each carries a different part of the argument that follows.

1. **Contact becomes continuous.** Below the crossing, contacts form and break; above it, an individual is in simultaneous contact with neighbours on all sides for as long as the condition persists.

2. **Transmission becomes possible.** A force chain requires a connected path of simultaneous contacts. Intermittent contact does not support one; continuous contact does.

3. **Recovery becomes unavailable.** The step that restores balance after a perturbation is the same step that Definition 1 counts. When $\varphi < 1$, a perturbed individual cannot recover by stepping and must instead be supported by contact with neighbours — which is to say, must transmit.

Items 2 and 3 are the two halves of the mechanism this paper quantifies. They are not independent conditions that happen to coincide: they are consequences of the same geometric fact, and they cross together because they are governed by the same quantity. Their separate consequences are developed later — the bearing of the crossing on the direction of exerted force in §4, and its bearing on accumulation in §5. The present chapter establishes only the geometric boundary itself.

Why density is only a proxy. φ depends on the spatial arrangement of bodies, not only on their number per unit area. Two crowds of equal mean density differ in φ if one is uniformly arranged and the other contains voids; conversely, φ can fall below unity locally in a crowd whose mean density is well below any nominal limit. Density is monotonically related to φ under uniform arrangement, and is therefore a serviceable proxy where uniformity can be assumed — but it is a proxy, and the assumption has to be stated. Where geometry produces convergence, taper, or a bounded run, the assumption fails precisely where it matters most.

Two stages within the solid phase. The crossing at $\varphi = 1$ is not the same as the exhaustion of all compressible gap. Between them lies a regime in which stepping is already unavailable — so contact is continuous and transmission is possible — but the bodies are not yet in maximal contact, and further compression can still occur by closing the remaining interstitial space. Only when that space is exhausted does the crowd reach the packing limit $\rho_{\max}$, at which no further compaction is geometrically possible.

We therefore distinguish, within the solid phase:

$\varphi < 1$, gap remaining	**solid, compressible**
gap exhausted ($\rho = \rho_{\max}$)	**solid, compacted**

The distinction matters for what quantity accumulates rather than for whether accumulation is possible. In the compressible stage the accumulating quantity is compaction: advance occurs by closing gap. In the compacted stage there is no gap left to close, and further advance appears as force. Both stages belong to the solid phase in the sense that matters mechanically — contact is continuous and transmission is possible throughout — and both are within the domain of this framework.

That the accumulating quantity is not the same in the two stages is a limitation of the treatment given here rather than a feature of it, and it is stated as one in §9.7. The relations of §5 are unaffected: they hold wherever contact is continuous, and do not require the two stages to be integrated over as a single interval.

2.2 Scope

This framework treats the solid phase. The fluid phase is outside its domain, and results derived here are not claimed to hold there.

That restriction is narrower than it appears, because the fluid phase is also outside the domain of the mechanism the framework is concerned with. Compressive loading that accumulates along a chain requires continuous contact, and by Definition 1 continuous contact is what the crossing supplies. A framework that claimed to cover both phases with one set of relations would be claiming that the same mechanism operates in a regime where its preconditions fail.

Why existing models do not exhibit this mechanism. The restriction also explains a feature of the literature that would otherwise count against the present account. Microscopic models of dense pedestrian crowds generally represent inter-personal interaction by a bounded soft-core repulsive potential — for example an exponential form, whose gradient is bounded above by a constant amplitude. A bounded repulsion admits mutual interpenetration: two agents can always be brought closer at finite cost. A system of such agents remains compressible at all densities and does not jam, and a system that does not jam does not develop force chains.

The consequence is structural rather than a matter of parameter choice. No setting of a bounded soft-core potential produces the solid phase of §2.1, because the phase is defined by the unavailability of further approach and such a potential never makes approach unavailable. Models of this class therefore occupy the fluid phase by construction, and their not exhibiting chain-borne loading is a property of their domain rather than evidence about ours.

The same reasoning applies to the empirical record. Controlled experiments on dense crowds are conducted in the fluid phase, for reasons that are sound: the solid phase cannot ethically be produced in volunteers. The absence of chain-borne loading from the experimental literature is therefore a sampling property of what can be measured, and carries no evidential weight against a mechanism that operates where measurement is unavailable.

The fuller comparison with existing accounts of dense-crowd mechanics, including the one empirical indicator that addresses the same practical question as §6, is given in §9.1 rather than here, since it requires the relations of §5 to be in hand.

What is not claimed. Nothing in this chapter asserts that the solid phase is injurious, or that the fluid phase is safe. The chapter establishes only that the two regimes support different mechanics, and that the framework's relations belong to one of them. Whether a given load is injurious is not a question this framework answers; the values on which such a judgement would rest are catalogued in §7 and are supplied from outside.

3. The P Architecture

3.1 Three roles

The framework carries three quantities and they answer different questions. Keeping them apart is what allows the third to be defined without reference to any physiological value.

	Role	Where
C	the input — what the scene supplies	§4
R	the process — what accumulates, and how fast	§6
P	the conclusion — what was borne, and for how long	this section

P is a conclusion rather than a design target, and §3.6 states why. The consequence for this section is that P must be constructible without knowing what load is injurious, since that is precisely the question a conclusion is not entitled to assume.

P is a form, and its dimension is not the framework's to fix. This should be said at the outset rather than discovered later. The quantity defined in §3.3 is an integral of a dose–response function that the framework does not supply, normalised by a capacity that the framework does not supply either. Its dimension, its scale, and the numerical value at which it becomes significant are all determined by those two choices, and different readers making different choices will obtain figures that are not comparable with one another.

What the framework fixes is the structure: which interval is integrated over, what the integrand takes as its argument, and the truncation that makes the result monotone in exposure. What it does not fix is the number. The name should not be read as implying a standardised index with a conventional threshold; there is no such thing here, and §7 explains why the framework declines to supply one.

3.2 Onset is two conditions, not one

An exposure integral needs a lower limit. The natural candidate is the moment at which loading begins to accumulate, and the difficulty is that this phrase has two readings which earlier formulations did not separate: the moment at which the *mechanism* engages, and the level at which the load becomes *injurious*. The second is a physiological quantity and is supplied from outside (§7). The first is mechanical, and this section shows that the framework owns it.

Two conditions must hold together.

The geometric condition. Accumulation requires continuous contact, which by Definition 1 is the solid phase: $\varphi < 1$. Below that crossing a disturbance is absorbed by the individual who receives it, and nothing accumulates. This condition is established in §2.1 and needs nothing further here.

The mechanical condition. Continuous contact makes transmission possible but does not make a chain *grow*. Whether it grows is settled by Proposition 4: writing f for the exertion a single contributor delivers and δ for the amount the receiving body absorbs irrespective of load, the borne force exceeds one contributor's worth at some depth if and only if

$$\beta f > \delta. \tag{2}$$

Where the inequality fails, every individual bears exactly f however deep the chain, and no accumulation occurs to be integrated.

Both quantities belong to the framework — the first to the exertion side, the second to the receiving body — and neither is a physiological threshold. The condition is therefore evaluable without one, which is what this section needs of it. The derivation is given where the recursion it rests on is established (§5.2.1).

The start of the clock.

$$t_0 \;=\; \inf\Big\{\, t \;:\; \varphi(t) < 1 \;\wedge\; \beta f > \delta \,\Big\} \tag{3}$$

Both conditions are necessary. The first without the second gives a dense crowd in which each individual bears one neighbour's exertion and nothing more, at any depth. The second without the first gives a growth condition that has no chain to act on.

3.3 Two integrals

What follows t_0 divides into two accumulations, and the division is one of ownership rather than convenience.

Mechanical loading. The trajectory of borne load,

$$F(t) \;=\; A_{L(t)}(t), \qquad t \geq t_0, \tag{4}$$

follows from the geometry, the chain depth and the exertion. Every term is derived within the framework or measured (§6.3).

Physiological dose. What that loading does to a body,

$$D(t) \;=\; \int_{t_0}^{t} g\big(F(s)\big)\,\mathrm{d}s, \tag{5}$$

depends on a dose–response function g which the framework does not possess and does not propose.

The framework therefore delivers $F(t)$ and t_0; the user supplies g and performs the integration. Externalising the threshold is on this account not a policy choice but a structural consequence of the split: what is delivered is the argument of the integrand, not the value of the integral.

Truncation belongs in the definition. Where a normalised scalar is wanted,

$$P \;=\; \frac{1}{D_c}\int_{t_0}^{\infty} \max(g(F(s)) - g(F_{\mathrm{ref}}),\, 0)\,\mathrm{d}s, \tag{6}$$

with D_c and F_{ref} both supplied by the user. The truncation is not a numerical convenience. Without it an interval spent below the reference level would offset an interval spent above it, which would assert that safe time discharges accumulated damage. Loading is reversible; damage is not. The asymmetry is a property of the quantity and is written into its definition rather than attached as a caveat.

3.4 The design side

A certificate issued before an event cannot rest on an integral whose integrand is not the framework's to supply. It rests instead on the derivative.

Definition 2 (Structural zero)**.** *A configuration satisfies the structural zero if* $\dot{F} \equiv 0$ *over the whole of the occupied domain and over every admissible arrangement of its occupants.*

A note on the word "certificate". It is used throughout this paper in the sense it carries in mathematics and in verification — an object that allows a claim to be checked — and not in the sense of an instrument that confers a status. What is offered is a demonstration together with the premises it rests on, so that a reader may determine for themselves whether the demonstration holds and whether the premises do.

The distinction matters in this subject more than in most, because the word is also the name of a regulatory instrument. Nothing in this paper is issued by, or on behalf of, any body empowered to approve a configuration for use, and no such standing is claimed anywhere in it. Whether a configuration may lawfully be occupied is a question for whoever is empowered to answer it; what follows is an input to that question and not an answer to it.

Two independent routes establish it, and they bound different things.

Route 1 — bounding the chain. Where partitions maintain a clear spacing s, no coherent chain exceeds that length, so $F \le f\,\Phi(s/p)$ with Φ as in §5. The bound is geometric: it holds because the space required for a longer chain is not present.

Route 2 — bounding the growth. Where $\beta f \le \delta$, Proposition 4 gives $F \equiv f$ irrespective of depth. This route is stronger in an important respect: it does not constrain the geometry at all. A configuration certified by Route 2 remains certified however long its runs, because no run accumulates.

Neither route requires the true value of β. Route 1 is evaluated at the upper end of the admissible interval, and an upper bound on transmission gives an upper bound on load. Route 2 is a comparison of two framework quantities. In both cases the certificate is valid whatever the true transmission proves to be, which is the property a certificate needs and an estimate cannot have.

Nor does either require the origin. The criterion of Definition 2 is a derivative, and a derivative is invariant under translation of the origin. Whatever reference level a user adopts, the structural zero is unaffected. The framework's weakest quantity — the scene C-value, for which no direct measurement exists (§4) — enters the design side only through f, and only as an upper bound.

What the certificate says. That $P \equiv 0$, for every choice of g and every choice of D_c, on grounds of geometric impossibility rather than statistical rarity. The distinction is the whole of the claim: a configuration that is rarely dangerous is one perturbation away from being dangerous, whereas a configuration in which the accumulating mechanism is absent does not become dangerous by being perturbed.

3.5 The forensic side

After an event the question reverses: not whether accumulation was possible, but whether it occurred, and for how long.

First filter — the load condition. That the load borne along the realised direction of convergence exceeded whatever level the reader adopts. Two requirements attach. The level is the reader's, not the framework's. And the chain must be the one that formed: a claim of exposure is not established by the longest chord the geometry admits in some direction, but by a chain aligned with the convergence that actually occurred. The burden here falls on the party asserting exposure, which is the correct allocation and the reason the requirement is stated rather than assumed.

Second filter — the exposure condition. That the load persisted. Relief may arrive by dispersal, by removal, or by the collapse of the configuration; whether it arrived in time is a question about the scene and not about the mechanics.

The filters differ in reversibility. If a crowd disperses, F falls and the first accumulation reverses. D does not: relief stops further accumulation and discharges none of what preceded it. This is the formal content of the framework's refusal to treat rescue capability as a load-bearing premise. A rescue that arrives halts the mechanism; it does not undo it.

3.6 Why P is not a design metric

Earlier formulations placed P at the centre of design assessment. The difficulty was not that the quantity was wrongly constructed but that it depended on a value the framework does not own, so that disagreement about that value propagated into every figure computed from it.

The generalisation is worth stating as a principle, because the framework applies it in three places:

> An indicator must not depend on a quantity the framework does not possess.

It is why thresholds are externalised (§7); why rescue capability is not a premise (§3.5); and why egress rate carries no weight in a capacity finding (§8.3). The three are one principle applied three times.

A permanent zero is the correct form. Under this architecture P vanishes for any configuration that passes §3.4, and takes a positive value only in reconstruction. That is not a defect of resolution. Discrimination on the design side is carried by the margin in the bound and by the robustness of the premises the certificate rests on — which is also what an underwriter is pricing. A certificate is a statement about premises; its value is the difficulty of violating them.

3.7 A framework-side scalar

Delivering only a trajectory is defensible but hard to compare across scenes. Where a single number is wanted on the framework's own side of the split,

$$J = \int_{t_0}^{\infty} F(s)\,\mathrm{d}s, \qquad [\,J\,] = \mathrm{N\,s}, \tag{7}$$

is available and contains no physiological input. Its limitations are specific and must travel with it.

It is ordinal, not absolute. J supports the statement that one configuration accumulates several times what another does. It supports no statement about injury, and none should be inferred from its magnitude.

It is linear in time. A low load borne for a long period accumulates the same J as a high load borne briefly, and the two are not equivalent in their effect. A dose function with the superlinearity that injury data suggest would suppress the low-load contribution — but such a function is exactly what the framework has externalised, so J cannot use it and retains the weakness. J is therefore reported alongside the trajectory and never in place of it.

And the residual it depends on requires a value. The growth condition of Proposition 4 compares two quantities, and earlier formulations carried the second of them at zero on the conservative side. That is sound where the quantity is used as a bound and unsound where it is used as a criterion: at zero the comparison is satisfied trivially, and the mechanical half of Eq. (3) degenerates.

A value is therefore needed, and the quantity is mechanical rather than physiological — it admits direct measurement under controlled loading, with no ethical obstacle. Two routes exist and they are not equivalent. A physical experiment yields a method others can repeat and a figure a reader can check. A simulation yields a figure sooner, but where the simulation is not public it yields an input the reader cannot check, and on this framework's own rules such a figure is unverified and cannot support a design-side certificate (§3.4).

4. The C-value

4.1 When alignment applies

Before deriving a bound on the horizontal factor it is necessary to say under what condition a bound of this form applies at all. The condition is the one established in §2.1, and the dependence is not incidental: the direction in which people exert force is governed by the same lateral free volume that separates the two phases.

Consider an individual who intends to advance and finds the way blocked by another person. Displacing the obstruction requires work, and the direction in which that work is least is not, in general, the direction of intended travel. If the blocking person has somewhere to move to, the cheapest displacement is lateral: a small sideways push, which the blocker accommodates by stepping, opens a path. If the blocking person has nowhere to move to, no lateral push accomplishes anything, and the only remaining option is to push in the direction of travel — against not merely the blocker but everything downstream of them.

The mechanical consequence is immediate. Force applied laterally is orthogonal to the axis along which a chain would transmit and does not contribute to it; force applied along the direction of travel does. Alignment is therefore governed by the availability of the lateral option, which is exactly the quantity of Definition 1.

Proposition 3 (Alignment is a phase property)**.** *Consider effort directed at removing an obstruction, resolved in the horizontal plane. In the fluid phase a lateral component is available to it, and the least-effort direction is not aligned with the chain axis. In the solid phase the lateral option is geometrically unavailable, and the effort is necessarily aligned with that axis.*

Why the vertical is not a third option. The resolution is taken in the horizontal plane because the vertical direction does not remove an obstruction. Force applied upward lifts or climbs; it leaves the person ahead exactly where they were, and the path is no more open than before. Vertical force is not thereby irrelevant to crowd loading — where bodies come to rest upon one another it is the whole of the load (§9.5) — but it belongs to a different configuration and is treated there rather than as an alternative within this one.

The practical reading is that the horizontal factor is small in the fluid phase not because individuals exert little force, but because the force they exert is not oriented to accumulate. This answers an objection that would otherwise be fatal: ordinary crowds contain a great deal of visible pushing, jostling and shoving, and no force chains of consequence form in them. The answer is not that the per-contributor force is small in such crowds. It is that the exertion is not aligned, and unaligned exertion does not accumulate along a chain however much of it there is.

A transition, not a gradient. Because alignment is governed by the availability of an option rather than by a continuously varying magnitude, it does not interpolate smoothly between its fluid-phase and solid-phase values. The lateral option is available or it is not. Treatments that model alignment as a smooth function of density therefore misrepresent the structure of the problem in the one regime where the answer matters.

Water and ice. The distinction is the one between a mass of water and a block of ice of the same weight, pressed against the same surface. Both deliver the same total force, and they deliver it quite differently: the water conforms, spreads over whatever it meets, and transmits nothing beyond the local pressure; the ice is correlated throughout, so what bears on one face is carried by the structure to whatever the opposite face rests against. The difference is not the amount of material. It is whether the constituent parts act independently or in concert.

A crowd with no common direction is the water. Its members exert as much as ever, and the exertions arrive at any given body from all sides and largely cancel; nothing is transmitted onward because nothing is aligned to transmit. A crowd with a common direction is the ice. The same people, the same exertion, and now the parts are correlated: what is applied at the back is carried through the intervening bodies to whatever the front is pressed against.

Consequence for what follows. The bound derived below is obtained under the constraint that exerted force lies within a forward sector. That constraint is not a modelling convenience; by Proposition 3 it is a statement that lateral escape is unavailable. The bound is accordingly a solid-phase result. It may be applied where the premise of §2.2 holds, and it may not be applied where it does not.

The reference angle is a convention, and is meant to be one. The routine value of the horizontal factor below corresponds to effort distributed within a sector of about sixty degrees either side of the direction of travel — roughly the span within which a person attends to what is ahead and can bring force to bear on it. No measurement is offered for that figure and none is claimed. It is a reference: arithmetically convenient, physically unsurprising, and stated so that a reader who prefers a different span can substitute it and see immediately what changes.

Its status is the same as that of the scale itself (§1.4). It touches neither of the limits that bound the framework's claims — not the ceiling, which is set by what a body can generate, and not the floor, which is set by the growth condition of §5.2.1. It fixes where the *routine* case sits between them, and a report adopting a different value is making a declaration rather than correcting an error.

The sector fixes the spread; the boundaries fix how much of the crowd is in it. Two quantities enter the horizontal factor and they are independent. The first is the angular spread of a group heading one way, which the reference angle fixes. The second is how many such groups the region admits, which the boundaries fix, and which the reference value does not address at all.

Within one group, the mean exceeds the floor. The reference value is the projection of the *worst-placed* contributor: at the edge of a $\pm 60^\circ$ sector, $\cos 60^\circ = 0.5$. Every contributor in the sector delivers at least that, which is what makes it a floor. The *mean* over the same sector is larger,

$$\frac{1}{2\pi/3}\int_{-\pi/3}^{\pi/3}\cos\theta\,\mathrm{d}\theta \;=\; \frac{\sqrt{3}}{2\pi/3} \;\approx\; 0.83, \tag{8}$$

and it is the mean, not the floor, that governs what a group of many contributors delivers.

Between groups, the factor divides. A region permitting movement in several independent directions distributes its occupants among that many groups. Those heading along a given axis contribute to it; those heading elsewhere do not subtract, but they do not add either. On any one axis the factor is reduced accordingly. Where a planar region admits movement in three broad directions, the factor on one of them is about a third of Eq. (8), or roughly 0.28.

A laterally bounded run admits one. Walls on both flanks do not merely remove an escape. They remove the alternative directions, so the division above does not apply and the factor is the mean itself:

Region	Directions available	C_h
Open ground, no common target	three	≈ 0.28
Laterally bounded run	**one**	≈ 0.83
Bounded run with a common target	one, and narrower	up to 1

The multiple between the first two rows is about three, and it follows from the boundaries alone. No common target, no instruction, and no intention on anyone's part is required to produce it.

Which is where the reference value must not be used. The 0.5 above is a floor on a single contribution within one group. It is not the factor for a bounded run, where the mean applies and no division takes place. A report using the reference value for a corridor has taken a figure roughly a third of the applicable one, and §5.2.1 sets out why an error of that size is not proportionate in its consequences.

Status. Eq. (8) is arithmetic on the declared sector. The decomposition into groups is a declared modelling structure rather than a derived one: it asserts that the directions a planar crowd takes fall into a small number of families set by the boundaries, and it does not obtain that number from anything more basic. What survives without the decomposition is the ordering — fewer available directions give a larger factor — and the results below rely on the ordering rather than on any particular count.

The sustainable fraction is a declared reference too, and a contested one. The routine value of the vertical factor below rests on the proposition that an exertion below some fraction of maximum voluntary contraction can be held indefinitely, and on a figure of about a seventh for that fraction (Rohmert, 1960, 1973). That proposition has not held up, and the framework should not be read as relying on it.

The difficulty is in what "indefinitely" meant. In the study from which the figure comes, a hold exceeding a quarter of an hour was treated as unlimited (Hagberg, 1981). Subsequent work finds fatigue developing at intensities well below that fraction — in the elbow flexors within an hour at under a tenth of maximum (Björkstén and Jonsson, 1977), and in the shoulder flexors sooner still (Mathiassen and Åhsberg, 1999) — and reports asymptotes, where any are found at all, at roughly half the figure carried here.

What follows for this framework, and in which direction. The value is retained as a declared reference, on the same footing as the exertion scale (§1.4) and the reference angle above. It is not retained as a measurement, and a reader adopting a lower fraction is correcting an over-estimate rather than departing from an established result.

The direction is the familiar one. A sustainable fraction that is too large makes the per-contributor exertion too large, hence the chain load too large — conservative where the framework is issuing a guarantee, and anti-conservative where it is asserting exposure (§5.6). It also bears on the growth condition: since the exertion at routine values and the residual it is compared against are of the same order (§5.2.1), a correction of this size moves an ordinary crowd further onto the non-accumulating side rather than across the boundary.

And it implies that the factor is not a constant. If no intensity is genuinely sustainable without limit, then the fraction a body can hold depends on how long it is asked to hold it, and the vertical factor is properly a function of duration rather than a number. Events last from minutes to hours, which is exactly the interval over which the reported curves are still falling. The framework does not carry that dependence at present; §3 is formulated as an integral over time and would accommodate it without structural change.

A note on the ceilings, before they are derived. Two limits on C appear below and one of them exceeds the other, which reads as a contradiction until the two are seen to bound different things.

The first is the point at which a person pushing horizontally can push no harder, because their feet would slide. That is a limit on force a person *generates*, and it is set by the friction available at the ground; on the scale of §4.3 it is unity by construction, since the scale is that same friction limit.

The second concerns a person who is not pushing at all. A body resting its weight on what is in front of it delivers a force whose source is gravity rather than muscular effort, and gravity requires no purchase against the ground to exist. Nothing about the friction available at that person's feet bounds it. Expressed on the same scale, body weight is the reciprocal of the friction coefficient, which is greater than unity for any surface a person can stand on — and it is greater for precisely the reason that makes the first limit what it is.

So the ordering is not a paradox but a consequence: a passive body can deliver more than an active one because it is not delivering the same kind of force. §9.5 carries this through to the configuration in which the whole load is of the second kind, where the two coefficients cancel and the framework returns the weight above.

And the two meet at a definite angle. The lean route of Eq. (13) expresses the vertical factor as the tangent of the lean angle divided by the friction coefficient, a quantity that grows without bound as the angle does. It does not apply without bound, and what limits it is the same friction that fixes the scale. A person leaning at angle θ requires a horizontal reaction at the feet of $W_0 \tan\theta$, and the feet supply at most μW_0. The route therefore holds for

$$\theta \;\leq\; \arctan\mu, \tag{9}$$

and at that angle it returns exactly unity — the active ceiling.

That the two coincide there is not a coincidence. The active ceiling is the friction limit expressed on a scale which is itself the friction limit, and the lean reaches it precisely when it has consumed the whole friction budget. Beyond that angle a person is no longer self-supporting: they are held by whatever is in front of them, and the case is the passive one, whose ceiling is stated below. The lean route and the passive ceiling are therefore not two estimates of one quantity between which a reader must choose. They describe adjacent intervals of a single angle, and they agree at the boundary.

More generally, the ceiling belongs to the source. A C-value expresses a force on the scale. It does not say what produced that force, and the limits below bound one producer only: exertion a person generates horizontally by pressing against the ground. Other producers are bounded by other things.

Source of the force	What bounds it
Self-generated, horizontal, through the feet	purchase against the ground
Self-generated, at chest height	postural equilibrium, which may bind first
A body resting its weight forward	its weight
A body that has fallen and is supported	its weight, and what the fall left
An agent that is not a person	the agent, and nothing carried here

The last row is the important one. Where a vehicle, a machine, or a structure giving way is delivering the load, nothing about human physiology or footwear limits it. The framework can express such a force on the scale once its magnitude is known — and the relations of §5 then describe how it propagates — but it neither computes that magnitude nor bounds it. Which sources a given configuration admits is a fact about the scene, and a report states it (§8.8) rather than inheriting a default from this chapter.

One restriction on that generalisation. The propagation relations are quasi-static (§9.7). A sustained external load — a fallen body that remains, a barrier bearing on the crowd, a vehicle that has come to rest against it — is within their scope once expressed on the scale. A transient one is not: an impact delivers its force over an interval short enough that the framework's treatment of what happens during that interval does not apply, whatever it may say about the state left behind.

4.2 Definition

The C-value is the weight that direction places on force. Two things are defined here and they are not the same: a scale, and a quantity expressed on it.

$F_{\text{den}} = 480\,\text{N}$ is the scale (Section 4.3). Any force may be divided by it, and the quotient is that force *on the scale*. Externally supplied levels are read that way (§7); so are loads no person produces.

The C-value is one such quotient and it is a specific one. Let f_{ax} be the axial force a *single contributor* delivers at their contact — the component along the compression axis, after projection.

$$C \equiv \frac{f_{\text{ax}}}{F_{\text{den}}}, \qquad F_{\text{den}} = 480\,\text{N}, \tag{10}$$

so that $f_{\text{ax}} = C \cdot F_{\text{den}}$.

C is per contributor, and the ceilings are why that matters. What a thorax receives is the sum over the contributors bearing on it; Eq. (25) assembles that sum from C and the depth. The two are not interchangeable and the ceilings of Section 4.6 are where the difference bites: a received load of 1200 N is 2.5 on the scale, and 2.5 is a number no contributor's C reaches. A total divided by the scale is not a C-value, and reading it as one would put every deep chain above a ceiling that bounds only individuals.

Equation (10) is a division. It carries no commitment as to how the force arose: whether it was pushed, leaned, transmitted along a chain, or delivered by an object.

4.2.1 What counts as a force

It does, however, carry a commitment as to what a force is. Three conditions admit a quantity to f_{ax}, and they are conditions on the quantity rather than on its origin.

1. **An identifiable contact.** The force is transmitted through a definite point of bodily contact. No contact, no contribution.
2. **A normal component.** The force has a component directed into the body it bears on. A purely tangential graze delivers nothing to f_{ax}.
3. **A traceable path.** The force has a transmission path that can be followed from its source to the thorax under load.

The three are not fussiness. They are what makes f_{ax} a quantity an instrument could read — a pressure mat, a force gauge — rather than a quantity a model computes and only a model can see.

And they exclude a specific thing. A repulsion that acts between individuals who are not touching fails the first condition. Such a term is standard in crowd modelling and is a coherent way to represent anticipatory avoidance, but it is not a contact force and it does not enter here. An "equivalent force" derived from separation, density gradient, line of sight, or inferred intention fails the same condition, for the same reason: there is no contact at which it acts. §9.1 takes up what this means for models built that way.

Objects are not excluded; non-contacts are. A vending machine bearing on a thorax satisfies all three conditions and enters at $C = 4800/480 = 10$ (Section 4.6). The criteria say nothing about what the source is. They say that whatever it is, it has to touch.

4.3 The denominator is a scale

F_{den} is the axial force an ordinary adult can exert. It is fixed once and thereafter treated as a unit. One route arrives at it, and this section reports the nearest measurement to it, which is not of the same quantity.

Route 1: weight and traction. An individual pushing horizontally cannot transmit more than their feet will hold:

$$F_{\text{den}} = \mu W_0 = 0.7 \times 686.5\,\text{N} = 480.5\,\text{N} \;\to\; 480\,\text{N}, \tag{11}$$

with $W_0 = m_0 g$ the weight of the ICRP reference adult ($m_0 = 70\,\text{kg}$, $g = 9.806\,65\,\text{m}\,\text{s}^{-2}$) and $\mu = 0.7$ a dry footwear–floor friction coefficient.

Route 2: measurement, and the distance to it. Fattal and Cattaneo (1976) report guardrail loading experiments conducted for occupational-safety load rating: three adult men leaning on and pushing against a rail at 42 in, with the resultant force and its angle tabulated per run. The reported force-to-weight ratios run from 0.24 to 0.37.

Those figures are not F_{den}, and the distance is worth stating exactly, because this is the nearest the literature comes. Three things separate them.

1. **The quantity is a resultant.** The tabulated ratios are of the resultant to participant weight, and the tabulated angles run from 45° to 82° from the horizontal. A rail at hip height is leaned upon, and what it records is largely weight transferred downward. Equation (11) bounds a horizontal force.
2. **Resolved horizontally, the figures are far below it.** The three-participant runs give a horizontal component of 45 lb pushing and 78 lb leaning — 200 N and 347 N — for three people together, against $\mu W_0 = 480\,\text{N}$ for one.
3. **The reaction is supplied by the rail.** A participant braced against a fixed structure is not a standing adult whose only reaction is friction, which is the configuration Eq. (11) describes.

The measurement therefore neither supports nor contradicts the derivation, because it is not of the same quantity. It is reported because a reader is entitled to know what the nearest available figure is and how far from the question it sits — and because the second and third items are not defects of the experiment. They are what an experiment looks like when it was instrumented for a rail.

Route 1 is a derivation of the scale; it is not a statement that friction governs any subsequent quantity. Once Eq. (11) has been used, μ has done its work: F_{den} is a fixed number, and a change in floor conditions does not change it.

Two friction coefficients, and only one of them is here. The μ of Eq. (11) is the footwear–floor coefficient: it governs how much force an individual can generate *against the ground*. A second coefficient governs how much tangential force transfers *across a body-to-body contact*, and it is a different quantity — clothing against clothing rather than sole against floor, with static values in the region of 0.5–0.7.

The two are physically independent. A person on a slippery floor wearing a grippy jacket has a low active ceiling and a high inter-person transfer; the reverse combination is equally available. That their numerical ranges overlap is a coincidence of materials and not a relation, and reading one for the other is the most natural error the notation invites. Only the footwear–floor coefficient appears in this paper, and it appears once, in Eq. (11).

4.4 Decomposition

Force is a vector; the C-value resolves it twice.

$$C = C_{\mathrm{h}} \cdot C_{\mathrm{v}}. \tag{12}$$

- C_{v}, the *vertical factor*, fixes the magnitude an individual delivers at the point of contact: $C_{\mathrm{v}} \cdot F_{\mathrm{den}} = F_{\mathrm{exert}}$.
- C_{h}, the *horizontal factor*, is the fraction of the delivered force that survives projection onto the axis along which the thorax is compressed.

4.4.1 Why the two are reported separately

The factors share an output type and nothing else. Both are ratios on $[0, 1]$ and both enter Eq. (12) by multiplication, and there the resemblance ends.

C_{v} is an equivalent conversion of how hard an individual is pushing. What sets it is physical capacity: how much force the musculature can deliver, and how much body weight a lean converts into horizontal force. Its ceiling is $1/\mu$, and that ceiling comes from weight and traction.

C_{h} is the actual direction of the force. What sets it is where the crowd is going. Its ceiling is unity, and that ceiling comes from alignment.

Neither ceiling can be derived from the other, because they are consequences of different physics. A composite value therefore cannot be inverted: a reported $C = 0.3$ may be 1×0.3 — a crowd with a common target exerting lightly — or 0.6×0.5 — a crowd exerting hard in dispersed directions. The two are the same number and different situations, and the difference is exactly what the weights of Section 4.8 are computed on.

This is also why neither needs to be observed. Both are difficult to measure at a scene, and the framework does not measure them. It bounds them, and it bounds them separately, because the bounds have separate origins. What an implementation must supply is the two factors and the provenance of each ceiling; establishing that provenance for a given scene is the report's burden, not the framework's (§8.8).

4.4.2 Computing C_v: one route

An individual standing at an angle θ from the vertical, in contact with another body, transfers weight into that contact. Taking moments about the feet for a rigid body contacting at centre-of-mass height,

$$F_\mathrm{exert} = W_0 \tan\theta \qquad \Longrightarrow \qquad C_\mathrm{v} = \frac{W_0 \tan\theta}{\mu W_0} = \frac{\tan\theta}{\mu}. \tag{13}$$

$\tan\theta$ is thus the fraction of body weight converted to horizontal force. Two assumptions bound Eq. (13): contact at centre-of-mass height (real contact is higher, giving a smaller moment arm) and rigidity. Both make Eq. (13) an upper bound on C_v at given θ.

The lean route is one way to obtain C_v, not a definition of it. Where F_exert is known directly, Eq. (10) applies without reference to any angle.

4.4.3 Computing C_h: directional coherence

C_h is a property of the group, not of any individual. Let L contributors deliver forces f_i at angles φ_i to the compression axis. The receiving thorax takes the axial projection:

$$F_\mathrm{recv} = \sum_{i=1}^{L} |f_i| \cos\varphi_i, \qquad \tilde{C} = \frac{\sum_{i=1}^{L} |f_i| \cos\varphi_i}{L \cdot F_\mathrm{den}}. \tag{14}$$

The numerator is what arrives; the denominator is what L contributors would deliver at the scale force in perfect alignment. $\tilde{C}$ is therefore bounded by the alignment of the group and by nothing else: components off the axis cancel and do not arrive.

$\tilde{C}$ is the mean of Eq. (10) over the contributors: each term $|f_i| \cos\varphi_i / F_\mathrm{den}$ is one contributor's C, and $\tilde{C}$ averages them. Where the contributors are alike, $\tilde{C} = C$. What neither is, is the total divided by the scale — that quotient is $L\tilde{C}$, it passes every ceiling of Section 4.6 once the chain is deep enough, and it is not a C-value.

The axis belongs to the recipient. The φ_i of Eq. (14) are measured against the anteroposterior axis of the receiving thorax, not against any fixed direction of the venue. Writing $\hat{e}_i^\mathrm{torso}$

for the unit vector along that axis and $\vec{F}_i^{\text{net}}$ for the vector sum of contact forces on individual i,

$$F_{\text{recv}} = \vec{F}_i^{\text{net}} \cdot \hat{e}_i^{\text{torso}}, \tag{15}$$

which is Eq. (14) in vector form. The distinction is not notational. A force of given magnitude and given direction produces a different F_{recv} according to how the recipient is turned within it: anteroposterior compression obstructs respiratory excursion directly, whereas lateral compression is carried substantially by the shoulder girdle and the upper arms.

This yields a prediction requiring no additional parameter. An individual who turns within a compressive field reduces $\hat{e}_i^{\text{torso}} \cdot \vec{F}_i^{\text{net}}$, and therefore C, and therefore the load on their thorax — with no change whatever in the surrounding force field. Survival under sustained compression should correlate with body orientation independently of position. The framework does not model why anyone turns. It requires only that turning changes which surface the force acts upon.

4.5 A reference value for C_{h}

C_{h} ranges over $[0, 1]$. Fully random directions cancel and give $C_{\text{h}} \to 0$; forces all delivered along the compression axis give $C_{\text{h}} = 1$. For rapid assessment the framework fixes a single reference value,

$$C_{\text{h}}^{\text{ref}} = \cos 60° = 0.5, \tag{16}$$

where 60° is the half-angle of the field of attention, taken as the forward 120°. Equation (16) is exact: it is the projection factor of the contributor at the edge of that cone.

The value carries two readings, and the pair is the reason it is fixed here.

As an upper limit on the routine state. An ordinary crowd has no common target, so its exertions do not fall within a single 120° cone and C_{h} does not reach 0.5; observed routine values are closer to 0.3. Using 0.5 for a routine estimate therefore overstates the load.

As a guaranteed floor under a common visible target. Where a directional instruction is unambiguous and is seen, every contributor's exertion lies within 60° of the indicated direction. The contributor at the edge of the cone still delivers $\cos 60°$, so the group delivers no less than 0.5 toward that direction. The floor is set by the geometry of attention; it does not depend on how the crowd is disposed within the cone, and it cannot be argued down.

Where a boundary case is at issue, C_{h} is computed from Eq. (14) directly. Equation (16) is a reference for rapid comparison, not a quantity to be refined: the values it is used to compare differ by more than an order of magnitude, and precision beyond the first digit buys nothing at that separation.

4.6 Values above unity

F_{den} is a scale, not a cap, so C is not confined to $[0, 1]$. Three cases place it higher.

Weight transfer. An individual who leans fully or has toppled transfers body weight rather than a traction-limited push:

$$C_{\text{passive}} = \frac{W_0}{\mu W_0} = \frac{1}{\mu} = 1.43. \tag{17}$$

Individual variation. W_0 and μ are reference values. A heavier individual, or one with better traction, exerts more, and their C scales accordingly. The present chapter fixes the scale and supplies the arithmetic; the distribution of human capacity across a population is not treated here.

External masses. A force whose source is not a person is subject to no bound from Eq. (17). A 4800 N object bearing on a thorax gives

$$C = \frac{4800}{480} = 10. \tag{18}$$

4.7 Why the division

Everything above could have been written in newtons. The scale would then be a constant appearing in the derivations rather than a denominator in the definition, and C would not exist. This section says what the division buys, because the answer is not notational.

A force is a number without a context. Consider a reported load of 400 N on someone's thorax. It carries no information about how it arose: it is consistent with an audience pressing lightly toward a stage, with the accumulation along a queue several people deep, and with the edge of a convergence where two streams meet. Those are different situations calling for different responses, and the figure does not separate them. This is why a quantity of that kind is not, in practice, something anyone records: recorded, it would not answer a question.

The same figure, on the scale, is not silent. It is $400/480 = 0.83$, and the ceilings of Section 4.6 then say something the newton figure did not: 0.83 falls below $C_{\text{passive}} = 1.43$, so one person leaning could have produced it, and the load therefore does not establish that a crowd was involved at all. That inference was always available in the physics; what makes it legible is that the number is expressed against a scale which is itself a physical limit. The newton figure is the same quantity with the limit removed.

And it reads a threshold the same way. The property generalises, and does so before any threshold value is known. Whatever the load at which some outcome occurs turns out to be, expressing it as a C-value says immediately whether one person can produce it: a threshold below $C_{\text{active}} = 1.0$ is one an individual reaches by pushing; one above $C_{\text{passive}} = 1.43$ is one no individual reaches at all, and which therefore requires a crowd. Those are different kinds of hazard — one a person can inflict, one only a configuration can — and in newtons the two look like the same kind of thing.

The values catalogued in §7 fall on both sides of that line. Nothing in this section depends on which, and nothing in this framework selects among them.

4.8 Completeness

Any external factor whose effect on safety is mechanical acts by changing a force at a contact, and every force at a contact is a number on the scale. The conversion is a division: a factor that adds 240 N at a contact adds $240/480 = 0.5$ on the scale, whatever the factor was.

The claim is inclusive rather than exclusive. It does not assert that the list of such factors is closed — no such list can be established — but that any factor admitting a mechanical description can be expressed on the scale, and that quantities so expressed are commensurable and additive. Each new factor is therefore an occasion to demonstrate the conversion, not a threat to the framework.

Which is not the claim that everything acts through C. It does not. A route that narrows raises the depth L_i as well (§5), and a depth is not a C-value: a partition changes the length available to a chain rather than the exertion delivered into it (§8.2). What is claimed is narrower and it survives: there is no channel by which a mechanical factor reaches a thorax except through a contact, which is what Section 4.2.1 requires, and every such contact is on the scale.

Differences in $\tilde{C}$. Where the exerted forces and the received force are both known, Eq. (14) determines $\tilde{C}$, and the only quantity standing between the two is direction. Any measure or condition that changes the received force at fixed exerted force appears as

$$\Delta\tilde{C} = \tilde{C}_{\text{after}} - \tilde{C}_{\text{before}}, \tag{19}$$

negative where the received force falls and positive where it rises. The magnitude is a weight in the units of Eq. (10), and weights so defined are additive across factors.

5. Transmission

5.1 Depth

Everything in this section presupposes the solid phase of §2.2. Where contact is intermittent no path of simultaneous contacts exists, and where lateral relief is available the exertion is not oriented along one (Proposition 3). The relations below do not apply in the fluid phase and are not claimed to.

Let $G(t) = (V, E(t))$ be the instantaneous contact graph: V the individuals, and $(j, k) \in E(t)$ wherever two bodies are in mechanical contact. For individual i, let $U_i(t)$ be the set of *upstream* contributors — those whose orientation and motion place a net push toward i. The chain depth is the length of the longest upstream-directed path terminating at i:

$$L_i(t) = \max_{\pi \in \Pi_i(t)} |\pi|, \tag{20}$$

where $|\pi|$ counts nodes and the receiving individual is not among them. L_i is the depth of the deepest chain bearing on i, not the number of bodies pushing on i: three neighbours each backed by two more make a depth of three, not nine, because the load accumulates along a path and Eq. (25) runs on the longest one.

Three properties constrain what L_i can be. They are geometric rather than assumed, and they are why the contact graph is tractable.

- **Sparsity.** Human geometry limits simultaneous contact neighbours to about six — the close-packing bound for standing bodies. The edge count is $O(N)$, far below the $O(N^2)$ of a complete graph.
- **Locality.** An edge exists only between individuals separated by less than twice a body radius, so $E(t)$ is tied to the Euclidean structure of the scene. A chain is a path through space, not through a list.
- **Time variation.** Edges appear, vanish and change weight between timesteps.

From a count to a distance. L_i is a number of bodies. Where a result must be expressed as a length — a partition spacing, a run of corridor — the conversion requires the separation between successive load-bearing contacts. At the packing limit that separation is fixed by the density itself: bodies on a square lattice at $\rho_{\max}$ per unit area stand

$$p = \frac{1}{\sqrt{\rho_{\max}}} \tag{21}$$

apart.

What this quantity rests on. The pitch is not an additional input beyond the packing limit; it is that limit restated in units of length. But the packing limit is itself a declared quantity rather than a derived one, and what fixes it is the plan area an upright body occupies. Anthropometry therefore enters here — not through the pitch, which adds nothing, but through the limit it is computed from. This paper's two declared quantities (§1.4) are sufficient for every result expressed as a force; expressing a result as a length requires the packing limit in addition, and that is a third declaration rather than a consequence of the first two.

An independent route to the same quantity exists and is worth noting, because the agreement is a check rather than a restatement. The separation between successive contacts along an axis is also the mean extent a body presents to that axis, averaged over orientations — a quantity obtainable from published torso dimensions by the Cauchy mean-width theorem, without reference to any packing limit. Evaluated on general anthropometric data it returns a figure within a few per cent of Eq. (21). The two are not the same derivation and should not be presented as one: they are different routes that happen to agree, which is the more informative situation of the two.

The square arrangement is the conservative of the two regular packings: a hexagonal arrangement at the same areal density gives a larger separation and therefore fewer bodies along a given run. A smaller pitch shortens the run over which a nominated load is reached, and raises the load that a given partition spacing admits — conservative in both of the uses to which the conversion is put (§8.1).

5.2 Accumulation

Let T be the map taking the force arriving at a body to the force it passes on. Two channels act there and they are not the same kind of quantity: one removes a fraction of what arrives, the other an amount fixed by the body rather than by the load. Together,

$$T(F) = \max(0,\ \beta F - \delta)\,, \qquad \beta \in (0,1],\ \ \delta \geq 0, \tag{22}$$

with β the proportional residual and δ the fixed one.

The affine form is derived, not adopted for convenience. Affine maps compose, and the two parameters do not mix under composition:

$$(\beta_1, \delta_1) \circ (\beta_2, \delta_2) = (\beta_1\beta_2,\ \beta_1\delta_2 + \delta_1), \tag{23}$$

so proportional residuals multiply and fixed ones accumulate. Neither the purely proportional family nor the purely fixed family is closed under composition — inserting one kind into a chain of the other leaves the family — and the affine family is their closure. The single fact that at least one channel removes an amount fixed by the body therefore *forces* the affine form.

Writing $f = C \cdot F_{\mathrm{den}}$ for the exertion a single contributor delivers, and taking that exertion to be approximately uniform along the chain, the map acts once at each body — on the total arriving there, not on each upstream source separately — so

$$A_{k+1} = T(A_k) + f, \qquad A_0 = 0, \tag{24}$$

and the load borne by the receiving individual is $F_i = A_{L_i}$. Solving,

$$F_i = f \cdot \Phi(L_i) \;-\; \delta \cdot \Phi(L_i - 1), \qquad \Phi(L) = \sum_{k=0}^{L-1} \beta^k, \qquad \Phi(0) = 0, \tag{25}$$

with the closed form, where one is required,

$$\Phi(L) = \frac{1-\beta^L}{1-\beta}. \tag{26}$$

Three boundaries of Eq. (25). Each is a convention or a short argument rather than a restriction, but leaving them implicit invites the objection that the closed form has been applied outside its domain.

At unity. Eq. (26) is undefined at $\beta = 1$; the sum it abbreviates is not. Direct evaluation gives $\Phi(L) = L$, and that is the value used wherever the design baseline of §5.6 is in force.

At zero depth. An individual with no upstream contributor bears nothing, so $F_i = 0$ when $L_i = 0$, and Eq. (25) is applied for $L_i \geq 1$. The second term is then well posed: at $L_i = 1$ it is $\delta\,\Phi(0) = 0$, which is correct, since the fixed channel acts on what passes through a body and there is nothing passing through the first.

Order of truncation. Eq. (22) truncates at each step, while Eq. (25) sums first. The two agree, because the truncation can bind only at the first step. If $\beta f \leq \delta$ then $A_k = f$ for all k by Proposition 4 and the sum is not in use. If $\beta f > \delta$ then $A_2 > A_1$; the map is non-decreasing, so A_k is non-decreasing, so $\beta A_k - \delta \geq \beta f - \delta > 0$ at every subsequent step and the truncation never binds again. No piecewise treatment is required, and the intermediate case a reader might expect — growth for some layers followed by truncation — cannot occur.

The fixed channel acts once, not once per source. A form in which each contributor's force independently traverses k bodies and is debited δ at each would be wrong in three ways. It would represent as δ a loss that is proportional to passage count, which is what β already is. It would make the absorbed amount depend on how the arriving load is composed, contradicting the definition of δ as fixed by the body. And it would violate split invariance: dividing one contributor into two delivering half each leaves the generated force unchanged but would increase the total absorbed. Physical absorption cannot depend on how sources are partitioned.

The fixed residual is declared zero, not estimated. Both factors of the correction term in Eq. (25) are non-negative, so

$$F_i \;\leq\; f \cdot \Phi(L_i) \qquad \text{for every } \delta \geq 0, \tag{27}$$

with equality at $\delta = 0$. The first term is the largest member of the family, so setting the fixed channel to zero biases the reported load upward and the reported safety downward. It is declared rather than estimated.

5.2.1 Whether the chain grows at all

The recursion settles a prior question. Continuous contact makes transmission possible; it does not by itself make the borne load exceed what one contributor delivers.

Proposition 4 (Growth condition)**.** *The load borne exceeds a single contributor's exertion at some depth if and only if $\beta f > \delta$. If $\beta f \leq \delta$, then $A_k = f$ for every $k \geq 1$: each individual bears exactly one contributor's worth of force, however deep the chain.*

Proof. From Eq. (24), $A_1 = T(0) + f = f$ and $A_2 = f + \max(0,\ \beta f - \delta)$. If $\beta f \leq \delta$ then $A_2 = f = A_1$, and f is a fixed point of the map, so $A_k = f$ for all $k \geq 1$. If $\beta f > \delta$ then $A_2 > A_1$; the map is non-decreasing in its argument, so $A_{k+1} \geq A_k$ follows by induction and the sequence increases. For $\beta < 1$ it converges to $(f - \delta)/(1 - \beta)$. □

Both quantities in the condition belong to the framework. The exertion is the product of the C-value and the exertion ceiling (§4); the fixed residual is a property of the receiving body, bounded above by its friction budget. Neither is a physiological threshold, and the condition is accordingly evaluable without one — which is what §3.2 requires of it.

The routine crowd sits on the non-accumulating side, and not by much. The two sides of the condition are estimated by independent routes — the exertion from a sustainable fraction of maximum voluntary effort together with the alignment of §4.1, the residual from what a standing body absorbs without passing on — and at routine values they come out within a small factor of one another, with the exertion the smaller. That they are comparable at all is not a coincidence requiring explanation: the force a person can sustain and the force a person absorbs without noticing are properties of the same body, and there is no reason for them to differ by much.

Its consequence is the substantive one. An ordinary dense crowd does not sit somewhere along a continuum of accumulating load. It sits on the far side of a threshold, where no accumulation occurs at all, and the framework's account of what happens above the threshold begins there rather than continuing smoothly from what precedes it.

What it takes to cross. The margin is a small multiple rather than an order of magnitude, and the multiple is within reach of a single change. Alignment alone will supply it: the horizontal factor at its solid-phase value against its value where effort is dispersed differs by about the factor in question, so a crowd acquiring a common direction crosses the condition *without anyone exerting themselves any harder.* The scene has changed; the people have not.

This is the mechanical content of the observation in §9.2 that a dense crowd with no common target does not accumulate. What keeps such a crowd on the safe side is not restraint. It is that there is nothing for its exertion to be aligned toward.

Near the boundary, the residual dominates. Where the exertion only just exceeds the residual, Eq. (32) saturates at a value that is small however favourable the transmission, because the numerator is a difference of two comparable quantities. Loads of consequence require an exertion substantially greater than the residual, not merely greater than it. The regime therefore divides in three rather than two: no accumulation; accumulation to a negligible ceiling; and accumulation to a ceiling that scales with the exertion.

5.2.2 The condition is not stationary in time

The preceding treats the exertion and the residual as fixed. Neither is, over the duration of an event, and both drift in the same direction.

Active exertion decays. No intensity of muscular effort is sustainable without limit (§4.1), and the interval over which the reported endurance curves are still falling is precisely the duration of a public event. Whatever a crowd is exerting at the outset, it is exerting less an hour later.

But a body that stops holding itself up does not stop delivering force. It leans. And a lean delivers force whose source is weight rather than muscle, which is the second of the two ceilings discussed in §4 and lies well above the first. The transition is not from exertion to rest but from one source of force to another, and the second is the larger.

The angle required is small. A forward inclination of a few degrees delivers a horizontal force comparable to the whole of the routine active exertion, and a few degrees more exceeds the residual. Neither is a posture an observer would remark on.

But a lean has a direction, and the alignment factor still applies to it. What a leaning body delivers is a force along the direction in which it leans, and a person leans forward with respect to their own orientation. Where a crowd faces a common direction, the leans are aligned and accumulate as any other aligned exertion does. Where it does not, they are distributed as the orientations are, and the horizontal factor of §4.1 discounts them accordingly.

This is not a qualification added to rescue the account. It is the same factor operating on a different source of force, and it is what reconciles the drift described here with the observation of §9.2 that a dense crowd without a common target does not accumulate however long it stands. Passengers in a carriage tire and lean like anyone else; they face in every direction while doing so, and what they deliver does not add along any axis. The drift raises what is available to a chain. Whether there is a chain for it to be available to is settled elsewhere.

And the residual decays with the exertion. Part of what a body absorbs without passing on is active: a person braced against what is behind them is using their postural capacity to absorb rather than to transmit. That capacity fatigues on the same schedule as any other. A tired body absorbs less and transmits more, so δ falls as f rises.

Whether the proportional residual moves with it is less clear. Its muscular component declines for the same reason, but tissue that has been compressed transmits more readily than tissue that has not (§9.9), and the two effects oppose. The framework does not resolve this, and does not need to: the two movements it can sign already point the same way.

So the growth condition becomes easier to satisfy as time passes. The quantity compared rises, the quantity compared against falls, and neither movement requires the density to change or anyone to form an intention. A configuration that fails the condition at the start of an event may satisfy it later with no change an observer could point to.

5.2.3 Where the voids are, and why it is the same axis

One of the three channels by which the proportional residual falls below unity is a void terminating the path, and §5.3 notes that it is a property of the layout rather than of the body. In one common class of scene that channel has a direction, and the direction turns out to matter.

A converging crowd closes across the flow before it closes along it. Take a crowd approaching a target small compared with the region it occupies, so that each individual moves along the line joining them to it, and let the approach speed be uniform. Two spacings then behave quite differently. Radial motion leaves the angular coordinate unchanged, so the arc separating two neighbours at the same radius is $r\,\Delta\theta$ and falls linearly with r. Two individuals on the same radial line approach at the same speed, so the difference of their radii does not change at all.

The areal density rises as the inverse of the radius, and by the above the whole of that rise is carried by the lateral spacing. The longitudinal spacing is conserved.

So the voids that remain are longitudinal. A contact graph closing under convergence closes first across the flow and last along it. Until the packing limit is reached, a void terminating a contact path is therefore a void along the direction of travel and not across it.

And that is the axis the chain runs on. In a run whose width is below the depth at which a nominated load is reached and whose length is above it, no chain across the width attains that depth while a chain along it does (§5.1). The long axis is where the load accumulates, and by the preceding paragraph it is also where the voids persist.

Proposition 5 (The two axes coincide)**.** *In a converging crowd within a run, the axis along which voids persist and the axis along which the accumulating chain must form are the same axis.*

The consequence is a dependency, not a cancellation. Whether a path is interrupted and whether the crowd has closed up are not two questions to be answered separately. They are one question about one axis. A treatment that draws the probability of a void from an isotropic packing argument, and the chain length from the run's long dimension, has assumed an independence between the two that convergence denies.

Scope. This is a statement about convergence and about nothing else. In a straight run with no convergence and uniform speed neither spacing changes and the argument says nothing: what generates the anisotropy is the reduction of cross-section available to the flow, not the fact that the crowd is moving. Real walking speeds are not uniform, and a spread in speed spreads the longitudinal spacing about its initial value, so conservation is an idealisation bounding how far the anisotropy survives rather than whether it arises. And lateral closure saturates at the breadth of a body, beyond which continued convergence is absorbed longitudinally or by a fall in flow rate — which is the configuration §5.8 takes up.

5.2.4 What a fall does to the chain

The treatment above concerns occupants who remain upright. One of them may not, and what follows from that is not obvious in either direction.

A fall severs the path. An occupant who topples ceases to be an upright load-bearing node. The horizontal load arriving at their position loses its downstream continuation, and the floor space they vacate becomes available to their neighbours. Both effects act against propagation: the chain is interrupted where the fall occurred, and the interruption comes with room.

An isolated fall therefore *truncates* a disturbance rather than carrying it. This is worth stating because the opposite is often assumed — that a fall in a dense crowd necessarily initiates a spreading collapse.

Unless the surrounding load is aligned. Where the exertion reaching the fallen occupant's position is coherent rather than spread, the toppling body's weight is redirected along the shared direction onto the occupant ahead, and it biases that occupant's own loss of balance into the same alignment. The disturbance is then carried forward as a directed cascade rather than truncated.

Whether a fall stays local is accordingly governed by the alignment factor of §4.1 — the same quantity that governs whether a load accumulates at all. It requires no separate propagation variable, and it inherits the dependence on boundaries established there: a fall in an open region truncates, and the same fall in a walled run may not.

And this is where occupancy does its work. A fall is in itself a common and unremarkable event; most end with the person standing up again. What decides whether one becomes something else is whether recovery remains available — whether there is room to rise, and whether the load from upstream relents long enough for rising to be possible.

That is what occupancy governs. Not whether someone falls — a fall is a matter of the load reaching them and of what they can withstand, neither of which is a density. Occupancy decides whether the floor space a recovery requires still exists, which is the criterion of §2.1 read at the scale of one body.

Proposition 6 (The role of occupancy)**.** *Occupancy does not determine whether an occupant loses balance. It determines whether one who has can recover.*

Which locates a criterion that is otherwise intractable. Establishing that no occupant will fall is not a proposition a venue can discharge: falls arise from footing, from contact, from a dropped belonging, and from a hundred things no configuration excludes. Establishing that a fall remains recoverable is a proposition about floor area and about relief at a boundary, and both are properties of the configuration.

So a criterion may retreat from the first to the second without giving anything up that was ever obtainable.

5.2.5 The free front edge, and why a gap is worth more than an interruption

A partition interrupts a chain by taking the load into a structure. A gap interrupts it differently, and the difference is worth setting out because it is larger than it appears and because it costs nothing to produce.

An occupant with nobody within reach ahead delivers nothing. Exertion directed forward is transmitted only where there is a body to receive it. At a free front edge there is none, so whatever the occupant exerts is expended against the ground and does not enter the accumulation. They are not a contributor.

And the same body is then available to resist. An occupant not engaged in pushing something ahead of them can brace against what arrives from behind, and what they can hold without being displaced is a measurable quantity — the force at which a standing person must take a step, which is considerably larger than the exertion any one contributor delivers.

Proposition 7 (The value of a free edge)**.** *Relative to a continuous run, a free front edge removes one contribution f and adds an absorption equal to what the front occupant can hold without stepping. The swing across the edge is the sum of the two, not the first alone.*

Which is why dividing a queue works, and why it does not require a barrier. A run released in groups, with an interval between them, carries a free front edge at the head of every group. Each edge is constituted by the occupants themselves and is renewed at every release. No structure is installed, nothing is rigid, and nothing has to be strong.

The contrast with a partition is instructive. A partition takes the load into the ground through a structure and holds whatever that structure holds; it is stronger, and it must be built, kept in place, and made rigid enough — none of which is easy in a temporary installation. A gap holds only what a person can hold, and it appears wherever a group is released.

The bound on a group's length follows. An edge persists while what arrives at it stays below what its occupant can hold. Since what arrives is the accumulation over the group behind, a group is self-sustaining only while

$$f \cdot \Phi(L_{\text{group}}) \;<\; F_{\text{stance}}, \tag{28}$$

where F_{stance} is the force at which the front occupant must step. Beyond that length the edge closes, the groups merge, and the division is undone. Groups must therefore be short — and how short is computable from Eq. (28) once F_{stance} is measured, which it can be.

How large the interval must be is a separate and lesser question. What matters is that the edge exists, not that it is wide. An interval of one pitch already removes the contact; a wider one buys the number of steps the front occupant may take before the edge closes, and so buys time rather than changing the mechanics.

There is an informational effect as well, and the result does not depend on it. An occupant in a continuous queue has no way of knowing where it ends, and the head of the queue is, so far as they can tell, the front of everything. Advancing whenever a step is available is a reasonable response to that. An occupant who can see that the queue stops a few metres ahead has no such reason, since a step taken now is a step spent waiting a moment later.

This is worth recording because it acts in the same direction as the mechanical result. But the mechanical result stands without it: at the edge the queue has in fact stopped, and the front occupant's capacity is in fact available, whether or not anyone behind has understood why.

5.2.6 The fixed residual, and why the run length is sensitive

The accumulation above carries a proportional residual. A second kind of loss enters at each transmission and behaves quite differently, and the difference is what makes the results below sharp.

The two residuals are not the same quantity. A proportional residual scales with what arrives: a body receiving twice as much passes on twice as much, less the same fraction. A fixed residual does not: a body absorbs a given amount whatever arrives, and passes on the remainder. The first rescales the accumulation; the second offsets it.

The accumulation with both. At the design baseline, where the proportional residual is taken at unity, the load borne at depth L is

$$F(L) \;=\; Lf \;-\; (L-1)\,\delta \;=\; L\,(f-\delta) \;+\; \delta, \tag{29}$$

since each of the L contributors adds f and each of the $L-1$ transmissions absorbs δ. Setting $F(L)$ to a nominated level F_* and inverting,

$$L \;=\; \frac{F_* - \delta}{f - \delta}. \tag{30}$$

The denominator is the whole of the matter. It is $f-\delta$, not f. Where the per-contributor exertion is close to the residual the denominator is small and the required depth is large; where the exertion is several times the residual the denominator approaches f and the depth falls toward F_*/f. A change in f therefore produces a *larger* proportional change in L, and the effect is strongest exactly where f is closest to δ — which is to say, in the ordinary case.

This is worth stating because it inverts an intuition. A modest increase in the exertion — of a kind that would be dismissed as within the noise of any estimate — does not lengthen the chain modestly. It shortens the run required to reach a nominated load by a multiple of its own size.

And it means the geometry can cross the growth condition on its own. The condition of §5.2.1 is that a contributor deliver more than a body absorbs. Since the residual is fixed and the exertion is $C_h C_v F_{\text{den}}$, the condition is crossed by any change that raises the alignment factor sufficiently — and §4.1 establishes that removing the lateral directions raises it by about a factor of three.

Proposition 8 (Boundedness as a switch)**.** *Where the exertion in an unbounded region falls below the fixed residual and the multiple obtained by removing the lateral directions carries it above, the transition between non-accumulating and accumulating is produced by the boundaries alone.*

Which changes what a wall is. On the reading in which the hazard is a matter of density, a bounded run is a place where a crowd happens to get denser, and the boundaries are aggravating. On the present reading they are not aggravating: they are the condition. An open region and a bounded one at the same occupancy, with the same people exerting the same amount, are on opposite sides of the growth condition, and nothing about the people distinguishes them.

So the familiar statement has it backwards. A large space is not dangerous because it is large. In an open region the exertion is unaligned and the accumulation does not begin, whatever the extent. What a large space permits is that enough people be admitted to fill it, and whether they are is a question about admission rather than about the space. A narrow bounded run is unfavourable directly and without any question of numbers, because it supplies the alignment that the growth condition turns on.

Where the declaration of the preceding paragraph does not hold. Eq. (27) licenses $\delta = 0$ wherever the fixed channel is used as a *bound*. It does not license it where the channel is used as a *criterion*: Proposition 4 compares βf against δ, and at $\delta = 0$ that comparison is satisfied trivially. The two uses are distinguished wherever both appear, and the consequence for the framework — that a value is required where a bound would not be — is recorded in §3.2.

5.3 What the coefficient is

The proportional residual is not a mechanism and it is not the whole of the receiving side's effect. That effect is the pair (β, δ) of Eq. (22): β carries what is removed in proportion to the load, δ what is removed regardless of it. Neither alone is the effect, and setting $\delta = 0$ is a simplification of known sign, not a statement that the fixed channel is absent.

Several distinct channels contribute, and by the composition rule their proportional parts multiply:

$$\beta = \prod_i \eta_i. \tag{31}$$

Force that diverts laterally and is absorbed rather than returned to the axis; the fraction taken vertically; a void that terminates a path. These are not alternatives — they act together.

Unity is not degenerate. Absorption fixed by the body rather than by the load enters δ and leaves β at unity, so $\beta = 1$ is a coherent description of a scene: it states that no channel removes a fraction. What it costs is read off Eq. (25). At $\beta = 1$ the load is $Lf - (L-1)\delta$, which grows without bound unless $\delta \geq f$, so a scene at unity is bounded only where the fixed channel absorbs a whole contributor's exertion, or where the geometry bounds L. What cannot be eliminated is the pair.

The value is not known, and neither is its order. The framework carries β as an unquantified factor rather than a calibrated one. §5.4 and §5.6 state what that does and does not cost.

5.3.1 And it is not a C-value

The two are ratios of the same kind and they attach to the same bodies, which is why the difference is worth stating.

C is a property of the exertion: the vertical factor is what a body generates by leaning, the horizontal factor is where that exertion points. Both answer *what the contributors put in.* The η_i answer a different question — *what survived passage through the bodies between* — and every one of them acts on force some other body has already generated. Nothing in Eq. (31) describes an exertion.

The naming invites the error. The vertical factor is so called because it is obtained from a lean angle measured from the vertical, and the force it delivers is horizontal. The second channel of Eq. (31) is force taken *in* the vertical direction, and it is a loss rather than a source. The word is the same; the quantities have nothing in common.

Conflating them costs a double count. Lateral spreading is already carried by the horizontal factor: a contributor off the axis delivers its projection and no more. A β that also represented geometric spreading would discount the same force twice. The η_i are what remains after that projection has been taken: absorption, not projection.

One channel is of neither kind. A void terminating the path is not an exertion and not an absorption; no body is present at which it acts. It is a property of the layout and the density, and it is why β is a scene parameter rather than a constant of the human body.

5.4 The critical coefficient as a map

Where $\beta < 1$ *and* the growth condition of Proposition 4 holds, the chain converges. From Eq. (24) the fixed point is

$$F_{\text{sat}} = \frac{f-\delta}{1-\beta}, \qquad \beta f > \delta, \;\; \beta < 1, \tag{32}$$

and outside those conditions the expression is not defined as a saturation value: where the growth condition fails there is no accumulation to converge and the load is f at every depth, and the right-hand side may then be non-positive and carries no meaning. A queue of a hundred and a queue of a million deliver the same bounded load. Nobody is excluded from the sum; the sum simply stops rising.

Inverting Eq. (32) gives the coefficient at which an unbounded chain just reaches a nominated load F_*:

$$\beta_{\text{crit}}(F_*) = 1 - \frac{f-\delta}{F_*}, \qquad \text{and at } \delta = 0, \quad \beta_{\text{crit}}(F_*) = 1 - \frac{f}{F_*}. \tag{33}$$

Which form to use. The second is evaluated at the declaration of §5.2 and is the appropriate one wherever the coefficient is wanted as a *bound* — it is the more conservative of the two, since discarding the fixed channel raises the load the chain is credited with and therefore lowers the coefficient at which a nominated load is reached. Where the fixed channel is instead carrying a criterion, the general form applies and a value is required (§5.2.1).

The pair is the framework's deliverable on this point, and neither member contains a physiological value. F_* is supplied by the reader from §7; the map is supplied here.

Reading the pair. Below $\beta_{\text{crit}}(F_*)$ the chain saturates beneath F_* at any length, and no increase in depth alters that. Above it, sufficient depth carries the load past F_*. Because the catalogue of §7 spans a range of values, it maps under Eq. (33) to a corresponding interval of coefficients, and three statements can be made without adopting any one of them: below the lower edge the chain saturates beneath every catalogued value; above the upper edge it passes every catalogued value at sufficient depth; between the edges the answer depends on which value the reader adopts.

What this replaces. An earlier formulation carried the critical coefficient as a fixed constant. That constant was Eq. (33) evaluated at one particular threshold, and every argument resting on it inherited that threshold's provenance. Restoring the map removes the dependency without altering any of the arithmetic.

5.5 Two ceilings, two natures

The load is bounded in two quite different ways and the text must keep them apart, because "bounded, therefore safe" is true of one and false of the other.

Regime	Bound	Nature
$\beta < 1$	$f/(1-\beta)$	mechanical — unreachable in principle
$\beta \to 1$	$f\,\ell/p$	geometric — the space is not there

The first is the formal content of the observation that ordinary queues do not accumulate: the chain terminates itself. The second is the design guarantee of §8.2: the chain would accumulate, and is denied the length.

The first bound is strongly non-linear near unity. A change in the third decimal place of β near unity changes F_{sat} by a factor. Margin must therefore be expressed in force or in chain length — both linear in that region and both communicable — and never as a percentage of the coefficient, which reads as small and is not.

5.6 The burden of proof

$\beta = 1$ is the load-maximising member of the admissible interval, from which the allocation of proof follows directly.

	Design / certificate	Forensic / attribution
Conservative direction	overstate load	understate load
Coefficient taken at	unity	lower bound, or evidenced
Burden falls on	one asserting $\beta < 1$	one asserting exposure

The allocation is not novel. It is the strength-reduction logic of structural design codes: a reduction must be evidenced, and no reduction is the baseline. Nobody is asked to prove that steel is not stronger than its specification; anyone wishing to design to a higher figure supplies test data.

The cost of the baseline is small where it matters. Adopting unity overstates the load relative to a scene at some $\beta < 1$, and the overstatement grows with depth:

$$\frac{L}{\Phi(L)} = \frac{L(1-\beta)}{1-\beta^L}. \tag{34}$$

At the depths a design certificate is trying to establish — short runs, partitions closely spaced — the factor is a few per cent. It grows only at depths that fail the certificate for other reasons.

The baseline is therefore nearly free in the region where certificates are issued, which answers in advance the objection that the assumption is too conservative to be affordable.

The forensic side may not invoke the baseline. A retrospective claim that a coefficient approached unity must rest on the scene, not on the design convention: the filled-container condition makes it a consequence rather than an assumption. The two contexts may arrive at the same value by entirely different routes, and a report must say which route it took. The question an opposing party will ask is whether the coefficient was assumed or established, and there is no third answer.

5.7 Adverse-bound construction

The pattern of §5.6 generalises. Where a quantity is disputed and one direction of error is conservative, the framework adopts the value most favourable to the party it might otherwise be accused of prejudging. If the finding survives, that party must overturn its own premises to contest it.

The framework uses this twice, in different currencies. Egress rate is taken at the full value the governing guidance permits, so that a capacity shortfall found under it is a shortfall under the guidance's own arithmetic (§8.3). Injury thresholds are conceded entirely, so that a load finding is a finding against the reader's own value and not against one the framework selected (§7).

What remains load-bearing after both concessions is geometry, conservation, and monotonicity.

5.8 Recovery and the ratchet

The load in Eq. (25) is borne rather than merely applied, and this section states why the distinction is not rhetorical. What makes a disturbance persist is the same geometric fact that makes it transmit.

Recovery consumes the quantity the phase criterion counts. An individual displaced from equilibrium restores it by stepping: the support base is repositioned beneath the centre of mass. That step requires floor area not occupied by another body — which is the quantity φ of Definition 1, counted for exactly this purpose. Recovery is available in the fluid phase and unavailable in the solid one, and nothing further need be assumed to obtain that: it is the third consequence listed in §2.1, restated where it bears on transmission.

A perturbation therefore changes kind at the crossing. Below it, a displaced individual steps, recovers, and the disturbance is dissipated by the person who received it. Above it, the step is unavailable and the only support remaining is contact with neighbours — which is to say that the only available support is the mechanism of §5.2. The displaced individual neither recovers nor falls; they lean, and what they lean on carries the load onward. A disturbance that would have been a transient becomes a boundary condition.

The asymmetry between advancing and retreating is geometric. Within the solid phase the two are not mirror operations. Advancing is *local*: the space ahead becomes available as soon as the single individual occupying it has moved, and one person's displacement suffices. Retreating is *collective*: the space behind is occupied by the upstream run, and it becomes available only if that entire run withdraws together. One requires a single neighbour to move; the other requires coordinated simultaneous motion along an extended path.

No psychological premise enters this. The asymmetry follows from which space is occupied by how many, and would hold for any medium with the same occupancy structure.

The contact channel carries magnitude, not configuration. A supplementary observation, mechanical rather than behavioural. What propagates through Eq. (22) is a force: a scalar magnitude along an axis. It carries no representation of the state of the bodies it has passed through. An individual well upstream therefore receives no signal distinguishing a downstream neighbour who is standing from one who is not, because the channel reaching them transmits only how hard, never what. The collective withdrawal that retreating would require is thus not merely difficult to coordinate: the medium supplies nothing on which to coordinate it.

The ratchet is the irreversibility of recovery. The name is a metaphor for a monotonicity, not a description of a mechanism: no latch engages, nothing holds the configuration where it is, and no part of what follows requires one. Assembling the preceding: in the solid phase a disturbance cannot be locally undone, and the collective operation that would undo it is not available. Occupation is therefore non-decreasing — not because anything drives it upward, but because the operation that would reduce it has no local form.

This is worth stating as the mechanism rather than as its consequence. Monotone compaction is what an observer sees, and it is easily mistaken for the cause; but a crowd that could recover would exhibit no such monotonicity even under identical inflow, and one that cannot exhibits it without any inflow surplus at all. The ratchet requires no net inflow, no urgency, and no exertion beyond what is already present. It requires only that the space needed to undo a displacement be occupied.

Bearing on the growth condition. The ratchet does not itself determine whether the borne load rises. That is Proposition 4, whose two quantities are properties of exertion and of the receiving body rather than of recovery. What the ratchet establishes is that whatever load is borne is *sustained*: the configuration does not relax of its own accord, and relief, where it comes, comes from outside the mechanism.

5.8.1 Formation, maintenance, and dissolution are three different operations

The preceding paragraphs describe what holds a configuration in place. It is worth separating that from what put it there and from what would remove it, because the three have different costs and the difference is the practical content of the phase distinction.

Formation requires work. Closing interstitial space is done against the resistance of the bodies being compressed, and the work is supplied by whatever is driving the crowd forward.

Maintenance requires none. Once the space is closed it is held by geometry rather than by continued effort. Compaction persists whether or not the upstream continues to push, because the operation that would reverse it is unavailable to any individual: recovery is collective (§5.8), and the collective operation has no local form.

Dissolution requires space, and space has one source. It can appear only at a free boundary, and from there it propagates inward at the rate at which individuals successively recover. It follows that the condition for dispersal is not that the driving force falls below whatever formed the configuration. It is that a free boundary sustains net outflow. Any residual inflow at the upstream end suppresses the propagating relief, since what the boundary clears the upstream replaces.

Nor can the downstream end reverse it beyond a certain depth. An individual at the front cannot relieve the load by retreating alone: what they vacate the chain immediately occupies. To create space they would have to push the whole upstream run back, and the force resisting them is the accumulated load they are already bearing. Their own capability is bounded by their ceiling, so reversal from the downstream end is available only where

$$\Phi(L) \; < \; \frac{1}{C}, \tag{35}$$

which at the design baseline is $L < 1/C$. At the reference values of §4.5 this admits a run of about a dozen bodies and no more. A deeper run cannot be undone from its downstream end however hard those at the front try, and the factor separating what maintains the configuration from what would reverse it is $\Phi(L)$ — the same quantity that makes the chain load large in the first place.

And holding position is lost sooner than reversing it. Eq. (35) compares the accumulation against what an occupant can generate *momentarily* — the ceiling itself, which is unity on the scale. Merely holding position against the accumulation, indefinitely rather than for one attempt, is bounded instead by what can be sustained, which on the same scale is the vertical factor. So

$$\Phi(L) \; < \; \frac{C_v}{C} \; = \; \frac{1}{C_h} \tag{36}$$

is the condition for holding, and it is a great deal tighter. Where alignment approaches unity it admits a single contributor: an occupant with more than one person behind them cannot hold position indefinitely against the accumulation, whatever their intention. What remains available beyond that is not resistance but a single attempt.

Three thresholds of the same form, and the band between the last two. Eqs. (36) and (35) together with the condition for reaching a nominated load F_* divide the depth behind an occupant into four regimes:

Depth behind	Available to the occupant	Borne load
$\Phi(L) < 1/C_h$	holding, indefinitely	below F_*
$1/C_h \le \Phi(L) < 1/C$	one attempt, not sustainable	below F_*
$1/C \le \Phi(L) < F_*/(CF_{\text{den}})$	**nothing**	below F_*
$\Phi(L) \ge F_*/(CF_{\text{den}})$	nothing	at or above F_*

The third row is the one worth naming. The lock closes before the load arrives: reversal from within becomes unavailable at a depth well short of the depth at which any nominated level is reached. In that band the configuration can no longer be undone from inside and the load has not yet reached the level, so what happens next is determined entirely by what acts from outside the run.

That is where an intervention can still change the outcome, and it is the only band in which one can. Its width is a property of the scene — it is the interval between $1/C$ and $F_*/(CF_{\text{den}})$, so it narrows as the exertion rises and as the nominated level falls.

All three are upper bounds. Each compares an exertion against the accumulation by static balance, which is the weakest form of the requirement: an occupant who merely balances the accumulation has not moved anything. Displacing the upstream run requires in addition overcoming its inertia and the ground friction of each body in it. The true depths are shorter than those tabulated, and the direction of the error is toward locking sooner.

So removing the cause does not remove the load. After compaction the accumulated force is borne by whatever bounds the configuration — a wall, a barrier, a structure — and no longer by the party whose advance produced it. That party may reduce its exertion to very little and the boundary load is unchanged, because the load was never being held up by their continued effort. A configuration is therefore not dispersed by persuading its upstream to stop.

5.8.2 What does and does not change the occupation

The preceding says that compaction persists. Two further statements follow from the same argument, and each is commonly assumed to be false.

A bookkeeping quantity. Index the occupants of a run $k = 1, \dots, N$ in order along it, let each present an extent p along the axis, and write $g_k \ge 0$ for the clearance between occupant k and the one behind. The *slack* of the run is $S = \sum_k g_k$, and its occupied length is $Np + S$. Contact throughout is $S = 0$. The occupation of a run of fixed membership is therefore fixed by S alone, and everything below is a statement about S.

The dimensions of the space do not enter. Whether a given clearance closes depends on the exertion of the one occupant behind it and on nothing else. The width of the run does not appear, its length does not appear, and the number of occupants does not appear. A wide run closes up at the same rate as a narrow one.

Corollary 9 (Locality)**.** *The closing of any clearance is independent of the dimensions of the region and of the total occupancy.*

Which inverts a common reading. A wider approach and a longer one both hold a given population at lower density, and both are ordinary provisions for a large crowd. Neither is disputed here. What Corollary 9 says is that neither slows the closing: a wide run arrives at the same contact state, and takes more people to do it.

What the dimensions do determine is the *extent* over which that state obtains once reached — and extent is what sets the depth of a chain (§5.1), which is what sets the load. Provision of space therefore governs the severity of the terminal configuration without governing whether it is reached.

This is not an argument against providing space. It is a statement that providing it does not address this mechanism, and that where it is provided for other reasons the consequence for this one runs the wrong way.

The state belongs to the run, not to the people in it. S is a function of the occupants' positions and of no property of the occupants themselves. Replacing one occupant by another at the same position leaves it unchanged.

Proposition 10 (Net removal)**.** *A release admitting one arrival for each departure holds N constant and leaves S where it stands, so the occupation of the run is unchanged by it. Only a release in which departures exceed arrivals lowers the occupation.*

Proof. Immediate from the preceding paragraph and from occupied length $= Np + S$. □

And an arrival admitted as space appears advances it. The newcomer arrives with clearance around them, and by §5.8.1 that clearance is consumed and not returned. Releasing forward as room appears at the head therefore moves the run toward contact rather than away from it.

Whether that matters in a given case depends on how close the run already is to the contact state, which is a question about how long it has been held rather than about the rate at which people are being let through.

5.8.3 The two phases in one region, and what follows for the upstream part

A region under pressure is rarely in one phase throughout. The part against the boundary is compacted; the part behind it, further from whatever is holding the front, still has floor area to step into. The two are in different phases and, by §2.1, have opposite properties.

The upstream part is both the safe part and the supply. It is safe in the specific sense this framework can state: the relations of §5.1 require continuous contact, and where a recovery step remains available they do not apply. It is the supply because the accumulation downstream is fed by advance from behind.

Three properties of acting on it, and they do not all point the same way.

1. **It cannot trigger the mechanism there.** Whatever a disturbance does to a crowd in the fluid phase, it does not produce a chain-borne load, because the phase in which such a load accumulates has not been reached. The hazard this framework describes is not among the things that can be caused by acting on that part.
2. **It stops the growth.** Formation requires work, and the work is supplied by advance from upstream (§5.8.1). Removing the advance removes what is still being added.
3. **It does not dissolve what exists.** Maintenance requires nothing, so the compacted part does not relax when the supply stops. This is the same asymmetry, and it bounds the benefit as firmly as the first property bounds the risk.

So the downside is bounded and the upside is real. A measure acting on the upstream part cannot produce the mechanism in the part it acts on, and it removes what is still being added to the part it does not act on. Neither statement is a prediction about how much: the framework does not supply the rate at which the supply would fall. Both are statements about direction, and the directions are not symmetric.

And the two timescales are not comparable. Compaction proceeds by the consumption of clearance, and clearance is consumed as occupants exert — so the state of a region is a function of how long it has been held (§5.8.2). That is a period measured in the tens of minutes and hours over which an admission runs, not in seconds.

Once a load is borne, the interval that matters is physiological, and the figures catalogued in §7 for sustained compression are measured in minutes. The interval over which a configuration forms and the interval over which it becomes fatal therefore differ by more than an order of magnitude, and it is the first that is available for a measure to be decided upon and carried out.

This is worth stating because the two are easily conflated. A compressive configuration is often described as though it arose suddenly, and the suddenness belongs to the consequence rather than to the formation.

What this does not authorise. Nothing here licenses acting on the compacted part. The first property above holds because the region acted on is in the fluid phase, and it does not transfer: a disturbance applied where contact is already continuous acts on a configuration in which recovery is unavailable, and this framework says nothing reassuring about that. The distinction between the two parts is the whole of the content, and a measure that does not distinguish them has not used it.

And a criterion in density cannot make the distinction. Density does not record whether floor area for a step remains, so it does not separate the part in one phase from the part in the other. A reading taken over a region containing both returns one number, and the two halves of that region have opposite properties. Whatever is decided from such a number is decided about both halves at once.

5.8.4 Why the order of two measures is not interchangeable

Two measures are available where a region has compacted against a boundary: stopping the supply from upstream, and opening a relief boundary so that occupants can leave. They are usually treated as independent, to be taken in whatever order circumstances allow. They are not independent, and the order is forced.

Relief dissolves only at net outflow. Opening a boundary permits departure, and departure lowers the occupation of a region only in excess of arrival (§5.8.2). Where the supply from upstream continues, each departure is replaced, the count is unchanged, and by the same argument the slack the newcomer arrives with is consumed rather than returned. The relief is then not dissolving the configuration; it is passing the population through it.

And it is spending something finite. A relief route has a capacity of its own. Used while the supply continues, that capacity fills at the supply rate, and what it buys is the throughput of the interval rather than a reduction in the occupation it was opened for. When it is full it is no longer relief: it is a second occupied region, and it cannot be opened a second time.

Which may leave a worse configuration than the one it relieved. A route provided for emergency egress is commonly narrower than the region it serves and is bounded on both flanks. If it fills to contact it satisfies the conditions of §5.1 on its own, and its geometry may admit a longer coherent run than the region it was opened to relieve. Nothing prevents a relief measure from producing a second instance of the configuration it was taken against.

Proposition 11 (Forced order)**.** *Where both measures are available, opening a relief boundary reduces the occupation of a compacted region only after the supply from upstream has stopped. Taken in the reverse order it consumes the relief capacity without reducing that occupation.*

Proof. Immediate from Proposition 10 and from the finiteness of the relief route's own capacity. □

The reverse order is not a weaker version of the same measure. It is a different outcome. The measure taken second in the correct order is still available; the measure taken first in the wrong order is not available again. What distinguishes them is not degree but whether the option is still held afterwards.

And the first step is the one with the bounded downside. Stopping the supply acts on the upstream part, which is in the fluid phase, and §5.8.3 states what follows: it cannot produce this mechanism where it acts, and it removes what is still being added. Opening a relief boundary acts on the compacted part, about which this framework offers no such assurance. The step whose consequences are bounded is also the step that must come first, which is a convenience rather than a necessity but is worth recording.

None of this says which measures a controller should take. It states a property of two of them: that their effects do not commute. A plan listing both without an order has not specified either.

Whether the two conditions coincide is not established. The free volume a recovery step requires and the free volume whose exhaustion removes the lateral option of Proposition 3 are counted by the same expression (Definition 1), but they have not been shown to agree in magnitude: a step and a sideways displacement of a neighbour need not occupy equal area. Where they do agree, the alignment property and the recovery gate become two readings of one condition, and the containment criterion of §6.5 becomes an identity rather than a correspondence. The question is taken up in §9.9; it requires no input the framework does not already hold.

Compaction acts on the load through the contact graph. The ratchet is stated above as a claim about occupation, and the load of Eq. (25) depends on depth rather than on occupation directly. The connection runs through §5.1. An edge exists between two individuals when they are in contact; closing interstitial space brings bodies into contact that were not, so the edge set grows as compaction proceeds. Longer upstream-directed paths become available, and L_i is by Eq. (20) the longest of them.

So the borne load rises with time in the solid phase for a reason that requires no change in exertion: the same individuals exerting the same force deliver more of it to whoever is at the end, because the path reaching them has lengthened. This is what makes $F(t)$ in §3.3 a trajectory rather than a constant, and it is why the mechanism does not require anyone to push harder as the event proceeds.

Where that account stops. It holds while interstitial space remains to be closed. Once it is exhausted the contact graph is complete in the relevant neighbourhood and no further path is available, at which point compaction can no longer be what accumulates. The two stages are those of §2.1, and the question of a single quantity accumulating continuously across the boundary between them is the open item recorded there.

6. The R Indicator

6.1 Why a rate

A quantity that answers “how much has been borne” is a conclusion and belongs to §3. What an operator needs while an event is running is different: whether the load is rising, and how fast.

The distinction is not one of convenience. Treating the exposure budget as fixed makes the tolerance time appear unstable, because it is not in fact a budget. For a level F_* supplied from outside,

$$\tau_{\mathrm{tol}} = \frac{F_* - F}{\dot{F}}, \tag{37}$$

which is a ratio of two things that move. An event that raises the exertion — a fall, an external stimulus, a change in what the crowd is trying to reach — raises $\dot{F}$ and raises F together, and the tolerance time collapses faster than either moves. Nothing anomalous has occurred; the quantity is simply a quotient, and both of its arguments have changed.

The three-way division. Eq. (37) also shows where each term is owned:

	Belongs to	Owner
$\dot{F}$, F	the scene	this framework
F_*	the body	the reader (§7)
τ	the coupling	neither alone

The framework supplies the state and its rate. It does not supply the level they are compared against, and R is accordingly defined without one.

6.2 The state, not the rate alone

A rate alone is ambiguous in a way that would be dangerous in an operational readout. A vanishing rate occurs in three configurations, and they are not equally safe.

Condition	F	$\dot{F}$	Reading
$\varphi \geq 1$ (fluid)	≈ 0	0	never engaged
$\varphi < 1$, $\beta f \leq \delta$	f	0	engaged, not growing
$\varphi < 1$, saturated	F_{sat}	0	fully loaded

The first and the third are opposite conditions reported by the same number. An indicator that carried only the rate would read the most loaded configuration a scene can reach as indistinguishable from one in which the mechanism has not engaged — and would do so precisely when the reading mattered.

Definition 12 (Operational state)**.** *R is the pair $\left(F(x,t),\, \dot{F}(x,t)\right)$ over the occupied domain.*

The remedy is standard rather than ad hoc: position and velocity together, as in any controlled system. It is also why §3 is built on an integral — an integral retains the position that a derivative discards, so the two readings are separable by construction.

The middle row is a certificate, not a warning. The second configuration is Route 2 of §3.4: contact is continuous, a chain exists, and it does not grow. Each individual bears one contributor's exertion at any depth. This is a state a scene can be designed to occupy, and it is distinguishable from the other two only because the state carries F as well as $\dot{F}$.

6.3 Observability

The rate formulation has a practical advantage that the earlier one did not, and it concerns instrumentation rather than theory.

An indicator built on the C-value requires measuring the C-value, for which no direct method exists (§4). An indicator built on $\dot{F}$ requires something else: whether load is accumulating in a region, which manifests as sustained net inflow into a space that is not clearing. That quantity has a proxy, and the proxy is usually already instrumented — entry and exit counts at controlled points exist at most venues, and are not commonly differenced.

The proxy yields a sign, not a magnitude. Sustained net inflow establishes $\dot{F} > 0$; it does not establish how large. This suffices for the use the framework puts it to, since the operational question is whether accumulation is occurring and the design question is whether it can occur at all — both binary. It does not suffice for reading a rate off a display, and the distinction must be stated wherever the quantity is reported.

What it does not resolve. The proxy locates accumulation in a region, not in an individual. It cannot recover F for a particular person, and nothing in this section should be read as claiming otherwise.

Nor can the individuals supply it. A separate observation, and an unwelcome one. The force at which a standing person first becomes aware of being pushed, and the force at which they become aware of pushing, are of the same order as the residual of §5.2 — which is to say, of the same order as the threshold at which transmission begins. Below it a person neither notices a push nor notices pushing; a body between two others registers nothing while the load on it is still being absorbed rather than passed on.

The consequence is that the onset of accumulation is not perceptible to those in whom it is occurring. By the time it can be felt, transmission is already under way. Reports from within a crowd, and the impressions of staff standing in it, cannot therefore serve as early indications of this mechanism — not because people are inattentive, but because the quantity becomes perceptible at about the same point at which it begins to propagate. Detection has to come from outside the crowd, which is what §6.3 provides.

6.4 Prediction refused, measurement delivered

A rate-based indicator appears to conflict with a position taken elsewhere in this framework, namely that the rate at which occupation rises is not supplied. The quantities that would be required for it — the frequency of exertion, the interstitial space consumed per event, the initial distribution of that space — are not available from geometry and are not carried here.

The conflict is only apparent, and resolving it improves the position rather than weakening it. None of the three phases in which R is used requires predicting the rate.

Phase	What is asked of the rate	How it is obtained
Design	that it is identically zero	proved (§3.4)
Operation	whether it is positive	measured (§6.3)
Forensic	its realised history	reconstructed

A structural zero does not require knowing what the rate would be; it requires showing that no mechanism produces a non-zero one. An operational readout does not require a magnitude; it requires a sign. A reconstruction does not predict; it recovers. The refusal to supply a predicted rate is therefore preserved intact, and becomes a statement about what the framework is careful not to claim rather than a gap in what it can do.

6.5 Design philosophy

The architecture of §3.4 and the present section together express a position that is worth making explicit, because it determines what a report is for.

A configuration is to be designed so that the mechanism does not engage, not so that its consequences can be managed once it has. This follows from three independences the framework has already established, each for its own reason:

- **Not rescue.** Lateral access into a loaded region is itself bounded by the mechanics of this chapter, so rescue capability cannot be a premise of a finding that the same mechanics produce.
- **Not egress rate.** It is an unverified behavioural input, and its failure transfers the shortfall upstream rather than merely withholding a benefit (§8.3).
- **Not a threshold.** It is not the framework's to supply (§7).

What remains is containment: an isolated loss of balance that does not propagate is not an incident. This is the standard posture of structural design, which does not require a structure to survive collapse but to not collapse.

And there is a mechanical reason, not merely a prudential one. The three operations of §5.8.1 are not inverses. Forming the configuration takes work; holding it takes none; undoing it takes space, which appears only at a free boundary and propagates inward one recovery at a time. A response that reduces the driving force is therefore not the reverse of the process that produced the state, and does not undo it. Beyond the depth of Eq. (35) the downstream end cannot assist either.

So the case for designing against occurrence is not that management is unreliable. It is that the operation management would have to perform is not the one that produced the situation, is slower than it by a mechanism the framework can name, and has a rate set by the geometry of the free boundary rather than by anything a responder controls.

Containment and the recovery gate are the same condition. An isolated perturbation fails to propagate exactly when the individuals around it can recover, and recovery requires the free volume of Definition 1. A configuration in which recovery remains available is one in which $\dot{F} \equiv 0$, because a perturbation that is absorbed adds nothing to a chain. The design criterion and the containment criterion are therefore not two requirements but one, stated in two vocabularies.

7. External Thresholds

7.1 This chapter carries no weight

The framework establishes what force a configuration can produce and how that force accumulates. It does not establish what force injures a person. Every value in this chapter is supplied from outside it, catalogued for the reader's selection, and used by the framework in no derivation whatever.

The claim is testable, and the test is mechanical. Delete this chapter. No equation elsewhere in the paper changes, no proposition loses a premise, and no result requires restatement. What breaks is a set of cross-references, together with the support for one observation about the state of the literature (§9.8) — and neither of those is a result of the framework.

Nothing outside this chapter is conditioned on any value inside it. The maps of §5.4 take a nominated load as an argument; the certificate of §3.4 is asserted for every choice of dose function; the operational state of §6.2 is defined without reference to a level. A reader who rejects every value catalogued here retains the whole of the framework and loses only the ability to convert its output into a verdict — which is the conversion the framework declines to perform on their behalf in any case.

No endorsement is offered. The framework does not assert that any catalogued value is correct, that the set is complete, or that any particular entry suits any particular population. Selection is the reader's, and §7.4 states what that selection carries with it.

7.2 The catalogue

Values are grouped by what they describe, because the groups are not interchangeable. Evidential character is stated for each: an experiment on living subjects, an output of a biomechanical model, or a retrospective estimate from an incident.

Value	Condition	Provenance	Character
Structural failure of the ribcage			
1000 N	static, prone, no fracture observed	Michalewicz et al. 2007, via Kroll et al. 2017	experiment
2550 N	static, flail chest (±250)	Kroll et al. 2017	model
4050 N	dynamic, flail chest (±320)	Kroll et al. 2017	model
Sustained compression — lower bounds on time to death			
1112 N	death may occur within 4–6 min	UK Home Office, via Hopkins et al. 1993 and Madea 2020	estimate
6227 N	death may occur within 15 s	as above	estimate

Table 1: External values, grouped by what they describe. The framework uses none of them and endorses none of them. Values mirror the project constants database; where the two differ the database governs.

Values examined and not carried. The catalogue is shorter than it might be. Tolerance figures for chest loading in conscious subjects — the force a person can bear against a bar, with and without the ability to push back — circulate widely in this field and are among the values a reader might expect to find here. They have been examined and are not included.

The reason is what this section is otherwise about. Those figures reach the present literature by relay: they are quoted in a paper that quotes another, the primary work could not be confirmed, and its date could not be established with confidence. Carrying a value whose origin cannot be named would be inconsistent with the treatment given to every entry that is included, where the chain is written out precisely so that a reader can judge it. A value is not made usable by being familiar.

This is not an assertion that the figures are wrong. It is a statement that the framework has no basis on which to represent them, and would rather catalogue five values whose provenance it can state than seven of which two it cannot.

The last pair are lower bounds, not typical values. The governing formulation is that death *may occur within* the stated interval — the earliest at which it is reported, not a median. An individual enduring longer is therefore not a counter-example; what would displace such a bound is an earlier death. Any use that reads these as expected survival times misreads the syntax.

The chain of transmission matters for the last pair. They are commonly attributed to the 1993 conference paper in which they appear. That paper is a conduit: the values are there credited to work by the United Kingdom Home Office, and the primary work has not been located by the present author. Readers adopting either value should be aware that its provenance terminates in an unexamined source, and that this is a property of the value rather than of the paper that carries it.

A measured field value, for scale. Load cells mounted at a stage barrier during an eighty-five-minute concert recorded a mean of 418 N at chest height across seventeen subjects, ranging from 116 N to 774 N. It is not a threshold and appears here only to indicate the magnitudes an ordinary event produces. It is moreover a panel-level aggregate, from which individual accumulation cannot be recovered; no inference about individual exposure should be drawn from it.

7.3 The gap between the groups

The catalogue has a feature that is easy to misread and that motivates the whole of the framework's treatment of these values.

Two entries sit barely a tenth apart and carry opposite labels: one is reported as a load under which no fracture was observed, the other as a load under which death may follow in minutes. They are not two labels for one quantity. The first bounds *acute structural failure* — whether ribs break. The second bounds *sustained hypoxia* — whether breathing is restricted for long enough. Different mechanisms, different timescales, and their proximity is a coincidence of arithmetic rather than mutual corroboration. Reading either as evidence for the other is an error this framework has made in earlier revisions and does not repeat.

That the second mechanism is distinct is documented, even where its force is not. A reader might reasonably ask whether sustained hypoxia is a separate mechanism at all, or merely a lower point on the same curve. The forensic literature answers the first question without answering the second. Autopsy series of crush asphyxia report the petechial and cyanotic signs of restricted respiration *without* rib fracture (Byard et al., 2006; Gill and Landi, 2004) — an outcome the structural-failure mechanism does not produce. Burial series, in which the applied pressure can be estimated from depth, describe the same presentation (Maron et al., 2007).

What this establishes is that the mechanism exists and is separable by its signs. What it does not establish is a force. None of these sources reports a measured load, and the estimates that can be reconstructed from them are not of a kind this catalogue would carry.

The region between them is acknowledged to be empty. The modelling source states in terms that its results do not cover the long-duration moderate compression characteristic of crowd loading, and refers the reader onward — to the retrospective pair whose own provenance terminates unexamined. The interval that matters most for the mechanism this framework describes is therefore the interval for which quantitative human data are least available, and the only figures within it are estimates rather than measurements.

The two preceding paragraphs are not in tension. The mechanism is documented and the force is not, and a catalogue of forces is the wrong place to look for evidence of the first.

What follows for the framework. Not that the values are wrong. That a framework resting a derivation on any one of them would inherit a provenance it cannot verify, and would produce findings that could be contested by disputing the value rather than the mechanics. Externalising them is the response, and the dual parameterisation of §3.3 — over both level and duration — is not a technical preference but the only honest form given that the distinguishing variable between the two groups is time.

7.4 Limits of applicability

Every value in Table 1 carries the population of the study that produced it.

The biomechanical model is derived from measurements on adult males, and its authors state that it is not validated for women, for children, for older adults, or for those with reduced bone strength. The retrospective pair carry no stated population at all.

Crowds at public assembly are mixed in age, stature and physical condition. The applicability of every catalogued value is therefore narrower in the setting this framework addresses than in the setting that produced it.

The obligation passes downstream. Any report adopting a value from this catalogue restates this limitation. A value selected without it is a value applied outside its validated population, and the fact that the selection was the reader's does not alter the arithmetic.

The same will apply to any future calibration. Should the framework later adopt externally measured values for quantities it currently carries unquantified — the fixed residual of §5.2, or a stability margin — the population question arises identically. Existing balance-recovery experiments use young adults under controlled perturbation in a laboratory, and the coverage gap is of the same kind. That gap is also, for a framework that intends to measure, an indication of what is worth measuring.

8. Applications

8.1 The design certificate

The certificate of §3.4 is issued against a configuration, not against an event. Its input is the geometry of the occupied domain together with the arrangement of any partitions; its output is a bound.

Two routes are available and a configuration may qualify under either.

Route 1 — bounding the chain. Where partitions maintain a clear spacing, no coherent path exceeds that length, and the load is bounded by $f\,\Phi(\ell/p)$ evaluated at the transmission coefficient's upper end. The bound is geometric: it holds because the space a longer chain requires is not present.

Route 2 — bounding the growth. Where the growth condition of Proposition 4 fails, the load is f at every depth. This route constrains no geometry at all, and a configuration certified under it remains certified however long its runs.

The deliverable is a bound with its premises attached. A certificate states that the load borne cannot exceed a stated value, under stated premises, and that the reader may compare that value against whatever level they adopt. It does not state that the configuration is safe. The word has no meaning within this framework, which possesses no quantity against which safety could be assessed, and a certificate that used it would be asserting something its own arithmetic does not support.

Premises are the substance. What distinguishes one certificate from another is not the bound alone but how difficult its premises are to violate. A bound resting on a partition that must be kept clear is weaker than the same bound resting on a wall, and a report that does not distinguish them has not reported the thing that matters.

Route 2 describes ordinary conditions, not a rare case. It would be reasonable to suppose that a configuration in which the growth condition fails is unusual — something a designer must work to achieve. The reverse appears to hold. The exertion of an ordinary dense crowd and the residual a standing body absorbs are of the same order (§5.2.1), so a crowd with no common target is already on the non-accumulating side, and the everyday configurations of §9.2 are instances of it that have been running for decades. What a certificate under this route establishes is therefore not an unusual property but the persistence of an ordinary one under the conditions the event will impose.

8.2 Partition spacing

The function of a partition in this framework differs from the function usually assigned to it, and the difference changes what a designer is choosing.

A partition does not block a dangerous chain. It caps the load a chain can deliver, by denying it length. The design variable is therefore continuous rather than binary: each spacing corresponds to a guaranteed ceiling, and the designer selects a ceiling rather than satisfying a rule.

This is stronger than the alternative framing. A guarantee expressed as a load contains no physiological premise, and can be compared against any value a regulator, an insurer or a court adopts — including values that did not exist when the design was made. A guarantee expressed as compliance with a spacing rule is only as durable as the rule.

The guarantee is conditional on clearance. A partition that is kept clear divides the domain; a partition that is stood upon does not exist. The bound therefore holds only while the clearance holds, and the premise belongs in the certificate rather than in a footnote. This is the characteristic point at which as-built departs from design, and reporting it as a premise is what allows a later reader to establish whether the premise held.

8.3 Capacity assessment

Where a configuration must accommodate a stated population, two quantities are available: the area it occupies, and the rate at which occupants can be removed from it. They are not equivalent inputs and the framework does not treat them as such.

The established practice this bears on. Venue assessment is a developed discipline with its own methodology, resting on admissible density together with the conservation of flow through the routes a site provides; Still (2000, 2014) set out the form in which it is generally applied. This paper does not offer a replacement for it and is not in a position to: that methodology addresses admission, routing, timing and the sizing of exits, none of which the relations here speak to.

Where the two meet is narrower and is worth isolating. Both produce a figure that bounds occupancy, and the finding of this section is that the two figures bound different things — one an evacuation requirement, the other nothing at all, since the compressive relations contain no headcount. That is a statement about which variable a criterion controls, not about whether the methodology is well executed, and it leaves everything that methodology does properly to it.

Static capacity is established before flow is considered. Two reasons, and the second is the stronger.

The first is epistemic. Area is geometric and can be verified by measurement; egress rate is a behavioural quantity taken from guidance, and whether a given configuration achieves it on a given day is not established by the guidance saying so.

The second is mechanical. When flow does not achieve its assumed rate, the deficit does not vanish. Every person admitted on the strength of flow credit is a person who must actually leave; where that does not happen, the shortfall

$$\Delta N = (R_{\text{assumed}} - R_{\text{realised}})\, T \tag{38}$$

appears upstream as occupation. Flow credit therefore does not merely fail to deliver a benefit — it transfers a hazard, and it does so at the moment the assumption fails, which is the moment the hazard matters.

Flow capacity is an admission with a liability attached. Expressed plainly: flow does not create safety. It converts margin into admitted numbers. That conversion may be entirely appropriate, but a finding that records only the benefit has recorded half of the transaction.

Adverse-bound construction applies here. A shortfall computed with the governing guidance's own rate, over its own permitted period, is a shortfall under that guidance's arithmetic. A party wishing to contest it must contest the guidance rather than the framework (§5.7).

A seated audience changes phase when it stands. There is a particular case in which an assumed rate is applied to a configuration that cannot deliver it, and it is common enough to state separately.

A seated audience is partitioned: each seat divides its occupant from the next, so no connected path forms and the relations of §5 do not apply. Standing removes the partitions. The row then fills, and where it fills to the condition of Definition 1 it is in the solid phase for as long as it takes to empty.

What follows is the asymmetry of §5.8.1. A row in that state does not empty at a flow rate, because flow presupposes that occupants can move relative to one another, which is what the phase removes. It empties from its free boundary — the aisle — inward, at the rate at which successive occupants recover the space to move. Nor can the far end assist: by Eq. (35) a run deeper than about a dozen cannot be reversed from its downstream end at routine exertion, and rows are commonly longer than that.

Rates taken from guidance are measured on crowds that are moving, which is to say in the fluid phase. Applying such a rate to a row that has entered the solid phase overstates what the row can deliver, and does so at the moment of the event rather than in the plan. This is a specific instance of the transfer described above: the deficit does not remain at the row, it accumulates upstream of it.

What is claimed here, and what is not. The claim is a condition on validity: a rate calculation presupposes the fluid phase, and where the receiving configuration has left it the calculation is not conservative but inapplicable. That much follows from §2 and §5.8.1 and requires no quantity the framework lacks.

What is not claimed is a corrected rate. Supplying one would require the rate at which recovery propagates inward from a free boundary, and that is precisely the quantity §6.4 declines to predict. A treatment of egress under these conditions is a separate undertaking of comparable size to the present one, and is not attempted here.

Headcount is not the variable this mechanism depends on. It is worth stating plainly, because it runs against the form in which capacity criteria are conventionally written.

Nothing in Eq. (40) contains a number of people. Two hundred occupants at a given partition spacing and two hundred thousand at the same spacing produce the same guarantee, because what the relation bounds is the length of a connected run. The same holds for the growth condition, which compares two mechanical quantities, and for the reversal bound of Eq. (35), which is again a length. A criterion expressed as an admissible headcount is therefore controlling a variable that the compressive mechanism does not depend on.

The practical consequence is a favourable one. Consider a square occupied area twenty metres on a side. Undivided, its longest run admits $\lceil 20/p \rceil = 53$ bodies in depth. Divided at two metres, the figure is 6. The guarantee falls by a factor of nearly nine, while the capacity lost is the area the partitions themselves occupy — nine lines of a tenth of a metre across a twenty-metre span, or some four and a half per cent of the floor.

The conventional remedy is to reduce density, which purchases a lower load by giving up occupancy directly and in proportion. Partitioning purchases a larger reduction at a small fraction of that cost. The difference is not a refinement of the same trade; it is a consequence of the two approaches acting on different variables.

A headcount bound exists, but not from this mechanism. What bounds occupancy is the requirement that the population be accommodated somewhere once it leaves — refuge area, and what the links carry away within the permitted period. Those are the quantities of the opening paragraphs of this section, and this framework supplies neither the holding density they require nor the rate (§9.7).

A venue assessment therefore carries two findings rather than one: a load finding, determined by the partitioning and stated as a guarantee; and an occupancy finding, determined by refuge and egress and drawn from the governing guidance. They are not substitutes, and a criterion that combines them into a single admissible number obscures which of the two is binding.

8.4 Expressing a standard in mechanical terms

A use distinct from assessment is available, and it is in practice the one a venue operator asks for first. Existing guidance defines admissible configurations in terms of density, effective width, permitted period and admitted headcount. Those configurations have a mechanical image, and computing it requires nothing this framework has not already supplied.

The map. The input is what the guidance specifies; the output is the load per person that those specifications imply, by way of the depth the configuration permits:

$$\text{admissible density, width, period} \longmapsto \text{depth } L \longmapsto F = f\,\Phi(L)$$

Nothing in this is new to the paper. It is Eq. (25) read in the direction a regulator's quantities happen to supply.

The result is configuration-specific, and that is the finding. It is not a property of the standard. Two venues complying identically may differ substantially in the load implied, because the depth available to a chain differs. This is not a limitation of the conversion: a standard expressed in density and width does not control the quantity this framework shows to be governing, and the size of the divergence between two compliant configurations is itself the evidence for that.

The framework adjudicates nothing here. The conversion permits the statement that a given standard, applied to a given configuration, corresponds to a stated load per person, and it permits comparison — between standards, and between configurations under one standard. It does not constitute a judgement that the standard is adequate, and none is offered: whether an admissible load is acceptable is a question for the body that set the standard. The conversion makes a quantity visible; it does not evaluate it, and no evaluation should be read into it.

8.5 A worked configuration

The relations of this paper are short enough to be applied without software, and the following runs all of them on the simplest configuration that exercises each: a straight run of corridor, uniform in width, divided by partitions at a spacing s. Width does not enter — the chain runs along the corridor, and what a body bears depends on the depth of the run bearing on it rather than on how many run in parallel.

Everything below follows from the two declared quantities of §1.4 together with s. No further input is used.

Step 1 — depth, from the geometry alone. By Eq. (21) the spacing between successive load-bearing contacts is $p = 1/\sqrt{\rho_{\max}} = 0.378\,\text{m}$, so a run of length s admits $\lceil s/p \rceil$ bodies in depth. The ceiling convention applies because the quantity is being used as a guarantee (§5.6).

s (m)	s/p	$L = \lceil s/p \rceil$
2	5.29	6
5	13.23	14
10	26.46	27
20	52.91	53

No parameter has yet been used. A reader disagreeing with anything later cannot disagree with this.

Step 2 — does the run accumulate at all? By Proposition 4 it does only if $\beta f > \delta$. Dividing through by the scale, and writing $C_\delta = \delta/F_{\mathrm{den}}$ for the residual expressed on it, the condition is simply

$$C > C_\delta/\beta, \tag{39}$$

which at the design baseline $\beta = 1$ reduces to $C > C_\delta$. Where it fails, the certificate of §3.4 is available by Route 2 and the remaining steps are not needed: the load is f at any spacing, and the partitions are doing nothing.

This test is worth performing first because it is the cheapest, and because where it passes the answer does not depend on the geometry at all. Its weakness is stated in §5.2.1: at routine values the two sides of Eq. (39) are of the same order, so the verdict turns on a declared quantity near its own boundary. A certificate resting on it is not robust to a scene in which the exertion might rise. Route 1 is.

Step 3 — the guarantee, where the run does accumulate. At the design baseline $\Phi(L) = L$, so

$$F_{\mathrm{guar}} = C \cdot F_{\mathrm{den}} \cdot \lceil s/p \rceil. \tag{40}$$

The result is one multiplication, and it is linear in both arguments.

		F_{guar} (N), declared C		
s (m)	L	0.075	0.15	0.30
2	6	216	432	865
5	14	505	1009	2018
10	27	973	1946	3892
20	53	1910	3820	7640

The three columns span the routine value and two multiples of it. No comparison against any level is made here; that is the reader's, using §7.

Step 4 — the same relation, inverted. A designer does not have a spacing and want a load. They have a load they are willing to guarantee and want the spacing that delivers it. Inverting Eq. (40),

$$s_{\max} = p \left\lfloor \frac{F_*}{C\, F_{\mathrm{den}}} \right\rfloor, \tag{41}$$

with the floor convention because the spacing is now the thing being bounded.

	$s_{\max}$ (m), declared C		
F_* (N)	0.075	0.15	0.30
600	6.05	3.02	1.51
1000	10.21	4.91	2.27
2000	20.79	10.21	4.91

Step 5 — what survives the uncertainty. The C-value has no direct measurement (§9.7), so the columns above are not three candidate answers among which one is correct. They are the sensitivity, and its structure is the finding:

- The guarantee is *exactly* linear in the declared exertion. Doubling it doubles the load and halves the admissible spacing — not approximately, and with no interaction between the two.
- The admissible spacing is therefore inversely proportional to the exertion, so an error of a given proportion in the declared value produces an error of the same proportion in the spacing, and no more.
- Consequently a report need not settle the exertion in order to be useful. It can state the spacing required at each of several declared values and leave the declaration where it belongs — with the party who knows what the event is.

That last point is what the worked example is for. A configuration is not certified by choosing a number for an unmeasured quantity; it is certified by a spacing, together with the range of exertion over which that spacing holds.

What this configuration does not establish. It does not establish how many people the corridor may admit. That question requires the static capacity, the flow, and the allocation between them (§8.3), and it requires an occupancy density that this framework does not supply. Nor does it establish that any tabulated value is acceptable: every figure above is a load, and whether a load is acceptable is settled elsewhere and by someone else.

8.6 A comparison that needs no dose function

The scalar of §3.7 is available without any dose–response function, and it is worth showing what that buys. Consider two reconstructions on the corridor of §8.5, with the trajectories taken as established rather than predicted (§6.4).

	Spacing	L	F (N)	Duration	J (kN s)
A	2 m	6	216	30 min	389
B	10 m	27	973	5 min	292

Every figure follows from the declared quantities and the two stated durations. No physiological input has been used.

The two obvious summaries disagree. Configuration A accumulates the larger J, by a third. Configuration B reaches four and a half times the load. A reader asking which is worse has been given two answers.

Where one trajectory dominates, the question is settled without a dose function. If $F_A(t) \le F_B(t)$ at every instant, then $g(F_A(t)) \le g(F_B(t))$ at every instant for any increasing g whatever, so $D_A \le D_B$ regardless of which dose function a reader adopts. The comparison is decided by the trajectories alone, and the framework can report it as decided.

The trajectories above do not dominate: A is the lower for longer, B the higher for less. So this comparison is *not* decided by anything the framework holds, and J reporting a definite ordering would misrepresent that. What J is for here is the opposite service — it makes visible that the two configurations trade load against duration, which is precisely the trade no threshold-free quantity can adjudicate.

What the pair does bound. J and the duration together give the mean load, and for any *convex* dose function — which is what the superlinearity of injury data implies — Jensen's inequality gives

$$D \;\ge\; \tau\, g\!\left(\frac{J}{\tau}\right). \tag{42}$$

The framework supplies J and τ; the reader supplies g; and the bound holds for every convex g without the framework selecting one. It is a lower bound only: an upper bound would need the peak, and a peak is not recoverable from a mean.

Evaluated on the two configurations, Eq. (42) reverses the ordering that J alone suggests as soon as any convexity is admitted. That is the expected result and it is the reason J is not reported alone: the convexity belongs to the reader's function, and the framework has no standing to assume it.

So the scalar earns its place, but narrowly. It is computable from the framework's own quantities; it decides comparisons in which one trajectory dominates another; it identifies those in which none does; and together with the duration it bounds the dose from below under an assumption about the shape of the dose function rather than its values. What it does not do is rank two configurations that trade load against duration, and a report using it for that purpose has exceeded it.

8.7 Forensic reconstruction

Applied after an event, the framework answers a narrower question than is usually asked of it. It does not establish cause. It establishes whether the configuration could produce a stated load along the direction in which convergence occurred, and for how long that load persisted.

Both filters must be established separately. The load condition (§3.5) requires a chain aligned with the convergence that occurred — not the longest chord the geometry admits in some direction. The exposure condition requires that relief did not arrive. Neither implies the other, and a reconstruction establishing only the first has established that a load was available, not that it was borne.

The design baseline is not available here. A retrospective claim that the transmission coefficient approached its upper end must rest on the scene. Where the configuration was filled to the point that no relief boundary remained, that is a finding about the configuration and the coefficient follows from it; where it was not, the claim requires other evidence. Adopting the design convention in a forensic setting would be adopting a load-maximising assumption in order to establish a load-maximising conclusion, and the burden allocation of §5.6 forbids it.

What a reconstruction can and cannot deliver. It can deliver the trajectory of §3.3 and the scalar of §3.7. It cannot deliver a conclusion about injury, because the dose function required to reach one is not the framework's. Where a reader supplies that function, the resulting figure is the reader's finding computed on the framework's trajectory, and a report should say so in those terms.

8.8 Reporting: three states, not two

A finding under this framework takes one of three values, and the third is not a variant of the first.

Established	the criterion is shown to hold under the stated premises
Not established	the criterion is shown to *fail* under those premises
Undetermined	**the available inputs do not decide it**

The second and third are not degrees of the same thing. The naming invites the confusion, so it is worth stating flatly. *Not established* records a determination that was reached: the inputs sufficed and the criterion failed. *Undetermined* records that no determination was reached: the inputs did not suffice and the criterion was not evaluated to a conclusion. The first is a finding about the configuration. The second is a finding about the assessment. Reporting either as the other misdescribes what was done, and the two call for different responses — the first for a change to the configuration, the second for an input.

Undetermined must not be reported as established. The two are distinguishable only if the report distinguishes them, and a reader encountering an absence of adverse findings will otherwise read one. Where an input is missing, the correct report names the input, states what it would decide, and returns no verdict on that criterion.

The three states are outcomes; the two routes of §8.1 are not among them. A reader meeting both structures in one chapter may take them to correspond, and they do not. Route 1 and Route 2 are alternative ways of reaching *Established* — one by bounding the chain, the other by failing the growth condition — and a configuration qualifying under either is in the same state as one qualifying under the other. Neither route leads to the second state or the third.

Findings are conditional by construction. Every criterion in this framework is evaluated under premises — the geometry as drawn, the exertion as declared, the transmission coefficient at a stated value, the comparison level as selected by the reader. A finding stated without them is not a stronger finding but an incomplete one, and the premises are what allow a later reader to determine whether it still applies.

One premise is easily overlooked: which sources of load the configuration admits. The ceilings on C derived in §4 bound what a person generates by pressing against the ground. They do not bound a body that has fallen, a barrier that has given way, or an agent that is not a person at all (§4.1). A report that carries those ceilings forward without saying which sources it has admitted is asserting something stronger than it has established, and the gap will not be visible in the arithmetic.

The declaration is short and belongs with the other premises: the sources of load treated, those excluded, and the reason for each exclusion. Where a source is excluded because it is judged not to arise, that is a judgement about the scene and its author should be identified along with it.

Why no named incident appears in this paper. A reader may expect a reconstruction of a known case, and the absence is deliberate rather than an omission.

A reconstruction is an assertion about a particular event, and assertions of that kind are properly made where the inputs can be established: with access to the geometry as it stood, to the record of what occurred, and with the standard of care that a finding about identifiable people requires. This paper establishes relations. It is not the document in which to apply them to an event, and applying them here would attach to the relations a claim they do not need and cannot support at this remove.

There is a further reason and it is worth being direct about it. Several of the events a reader might have in mind remain subject to proceedings or to continuing disagreement about cause. Naming such an event in a paper that offers a mechanism is an implicit assertion that the mechanism applies, whatever disclaimer accompanies it. The framework makes no contested causal judgement and names no unresolved case, and the worked configuration of §8.5 is fictional for that reason.

What the framework can be asked to do after an event is set out above. What it yields when so asked belongs in a document addressed to that event.

9. Discussion

9.1 Relation to existing models of dense crowds

The account given here does not compete with the established microscopic models of pedestrian crowds, and the reason is worth stating carefully because a coarser version of it would be wrong.

Purely repulsive formulations cannot reach the solid phase. Where inter-personal interaction is represented by a bounded repulsive potential alone — the form introduced by Helbing and Molnár (1995), whose gradient is bounded above by a constant — two agents can always be brought closer at finite cost. A system of such agents remains compressible at every density, and no choice of parameters produces the phase of §2.1, which is defined by the unavailability of further approach.

Contact-force formulations are a different case, and the distinction is not that they fail to jam. The extension of Helbing et al. (2000) carries a body force linear in overlap together with a tangential friction term, and is not bounded in this way. Such formulations do develop connected load-bearing structures — arching at constrictions among them. Nothing in this paper disputes that. A claim that they cannot exhibit chain-borne loading would be false, and is not made here.

The difference is identifiability. In those formulations the psychological term and the mechanical term act on the same variable — the separation between two people — and enter the equation of motion additively. An observed spacing is therefore consistent with many combinations of the two: a crowd keeping its distance because it prefers to and a crowd unable to compress further produce the same configuration, and no measurement of position or velocity distinguishes them. The parameters are fitted jointly to trajectories, and the split between wanting space and lacking it is not recoverable from the data the fit uses.

This matters for the present purpose specifically. A borne load must be attributed to one of two things: force a person is exerting, and force a person is transmitting because they cannot yield. The C-value carries the first and the transmission map the second, and they are estimated by different routes — the exertion from what a body can sustain and where it is directed, the transmission from what a body absorbs. Whether that is the right separation is a fair question. That it is a separation, rather than a single fitted quantity, is what allows the question to be asked at all.

And the relations here are derived rather than simulated. A simulation reports what its parameters imply for the case it was run on. The relations of §5 give a load for a geometry in closed form, so a reader can check them against a configuration without running anything, and can see which quantity a conclusion depends on. That is a difference in what the two things are for, not a claim that one is more accurate.

9.1.1 *Crowd pressure and the present indicator*

One empirical quantity in this literature addresses the same practical question as §6 and deserves to be set beside it, since a reader familiar with the field will otherwise supply the comparison unaided and may draw the wrong conclusion from it.

Helbing et al. (2007) proposed a *crowd pressure*, defined as the local density multiplied by the local variance of velocity, and used it to locate where and when a dense crowd enters the turbulent state that precedes injury. The quantity is measurable from video, requires no instrumentation of the crowd, and has been applied at scale. On the practical question — where is this crowd becoming dangerous — it is farther advanced than anything offered here.

It is not a mechanical pressure, and its authors say so. The quantity is gas-kinetic in construction. Johansson et al. (2008) state explicitly that it is to be distinguished from the mechanical pressure experienced in a crowd, and that a monotone relationship between the two is only *likely*, and then only after averaging over the fluctuations of the force. This paper takes them at their word. The framework does not dispute the indicator; it supplies the quantity the indicator was distinguished from.

The two apply on opposite sides of the crossing. Crowd pressure is built on velocity variance and therefore requires motion: Yu and Johansson (2007) note that turbulence disappears where flow does, so a non-vanishing flow is essential to it even at high density. That places the indicator in the fluid phase of §2.1 — or, more precisely, at its upper edge, where movement is violent but has not yet ceased. The relations of this paper begin where movement has ceased.

The division is therefore not between two answers to one question. It is between two questions on either side of a transition: which region is about to jam, and what load the region bears once it has. Neither answers the other, and a practitioner has use for both.

They share one difficulty, and for the same reason. Crowd pressure vanishes in a calm crowd and vanishes again in a fully jammed one, since neither has velocity variance. The rate of §6.2 vanishes in the same two conditions, for the analogous reason. In both cases the resolution is the same: the indicator must be read together with the state it is a rate of, not alone. That this appears independently in a kinetic construction and in a mechanical one suggests it is a property of the transition rather than of either formulation.

The same holds for the empirical record. Controlled experiments on dense crowds are conducted in the fluid phase, for sound reasons: the solid phase cannot ethically be produced in volunteers. The absence of chain-borne loading from experimental data is therefore a property of what can be measured. The literature is silent on this regime because the regime is inaccessible, not because it has been examined and found empty.

And there is a second reason, which is stronger. Even where the solid phase is reached — and everyday crowding reaches it — the accumulation does not begin. Proposition 4 fails at routine exertion, so an experiment conducted on an ordinary dense crowd would measure no accumulation however carefully it was instrumented, and would be right not to. The mechanism is not merely hard to observe; under routine conditions it is not operating.

This has a consequence for how the framework should be read. A quantity that is zero in every accessible condition and non-zero only in conditions that cannot be produced deliberately will not be discovered by measurement. It has to be derived, and then its predictions checked against the conditions that do occur — which is the argument of §9.2.

Two arguments this paper declines to use. Both are available and both would be convenient; neither is consistent with the framework's own commitments, and declining them explicitly is preferable to declining them silently.

The first is the observation that increased urgency can reduce throughput. In its simulation form this arises from friction and arching terms within force-based models — terms this framework treats as non-identifiable artefacts of that formulation. Invoking the result would be invoking the very construction the framework declines to rely on elsewhere.

The second is the charge that two-layer contact formulations double-count the force transmitted between bodies. They do not, because they are not counting force: such models carry no mass, no friction and no force dimension, their potentials being specified as accelerations. The accurate statement is narrower and harder to contest — that they possess neither the dimension nor the accounting structure required to quantify borne load, and were not constructed to.

9.2 A discriminating observation

The account given here and the density-based criteria it is offered alongside make predictions that differ in kind rather than in degree, and one configuration separates them.

What each must say about a dense crowd with no common target. A criterion expressed in density says that risk rises as density rises. It does not assert that risk is zero at any density at which people are present, and could not: the quantity it is expressed in does not distinguish a crowd that is merely dense from one that is also loaded.

This framework asserts something stronger and narrower. Where exertion is unaligned, the contribution reaching the chain is small; where it is small enough that Proposition 4 fails, the borne load does not exceed what one neighbour delivers, at any depth and for any duration. Not a low risk — an absent mechanism.

The configuration exists and is enormously sampled. Commuter rail at peak occupancy is dense, is entered voluntarily, is repeated daily, has been so for decades, and involves no common target: the occupants of a carriage are not trying to reach anything, and there is nowhere for them to go until the doors open. If the per-exposure risk of spontaneous compressive injury were non-zero, an exposure of this magnitude would have made it a statistical fact rather than a matter of argument.

The claim is restricted to the stationary case, and the restriction is not incidental. Everything above concerns a carriage at rest or in steady motion. Acceleration and braking are a different matter, and by this framework's own reasoning they are a serious one: an inertial force acts on every occupant at once, in the same direction, for as long as the acceleration lasts. That is alignment — not a common target, but a common direction, which is the quantity that actually enters the chain. The horizontal factor of §4.1 is not small during such an event, and the argument of this section does not extend to it.

Nor can the present treatment be extended to it as it stands. The relations of §5 are quasi-static (§9.7), and an acceleration event is transient by construction: what matters is the magnitude, the duration and how the crowd's own compliance responds to both. Those require dynamics the framework does not carry.

The case is nonetheless an attractive one to study, and worth marking as such rather than leaving unmentioned. It supplies a controlled, repeated, scheduled alignment event of known magnitude, applied to a dense crowd, in a configuration that is otherwise the non-accumulating one — which is a combination that does not occur elsewhere and cannot be produced deliberately. It is deferred here because treating it properly requires the dynamic extension, not because it is uninteresting.

The absence is selective, which rules out the obvious alternative. It might be said that such systems are simply well managed. But management suppresses hazards generally, and other mechanisms in the same environment are not suppressed: entrapment at moving stairways, falls, and injuries at closing doors are all documented. The environment, the population and the management are common to all of these. One mechanism is missing from an environment in which comparable mechanisms are present, and that pattern is not what management produces.

Which model predicts what is seen. A model that predicts non-zero risk must account for an exposure that has not produced events at the implied rate. This framework predicts the observation directly, and predicts it as a structural consequence rather than as a small number — which is the distinction §3.4 rests on. A configuration that is rarely dangerous is one perturbation from being dangerous; a configuration in which the accumulating mechanism does not engage is not.

Scope of the claim. The three parts of such a system are not equivalent and only one is conclusive.

- **The carriage** is the conclusive case: dense, contained, and without a common target.
- **The platform** is not near the onset condition at all, so it is consistent with the framework without testing it.

- **Stairways** have a distinct problem. What initiates an incident there is loss of footing — a spilled drink will do it — and that is an initiating mechanism this framework does not carry. It may nonetheless be the trigger of a subsequent compressive event, and where it is, the relations of §5 apply to what follows. The division is the one already drawn in §9.5: initiation and load are separate questions, and the framework answers the second.

The form of the claim. It is not asserted that no such incident has ever occurred. That form of claim obliges its author to reconstruct every case another party cares to propose, and the obligation is unbounded. The claim made here is that the observed rate is disproportionate to what a density-based prediction implies at this exposure. A party disputing it must reconstruct a case far enough to identify a decisive factor the present account does not carry — which is the same allocation of burden this framework applies to its own claims (§5.6).

9.3 Why inertia is negligible

A recurring proposal is to infer contact loads from density and velocity fields. The proposal fails on magnitude, and the arithmetic is short enough to give.

An individual brought to rest over a distance d delivers an impulsive force of order $mv^2/2d$. Braking distances in a crowd are of the order of a single step, and crowd speeds at the densities of interest are a fraction of free walking speed. Evaluated at reference body mass, the resulting force is an order of magnitude below the sustained exertion an individual delivers by leaning — and remains a fraction of it even at free walking speed, which is not attained at the densities where the mechanism operates.

The conclusion is not that inertia is absent. It is that a formulation resting the load on inertia is attempting to explain a large term with a small one. Where the dominant contribution is sustained exertion and the inertial contribution is an order of magnitude smaller, a density–velocity mapping cannot recover the quantity of interest even approximately, and its agreement with any measurement would be coincidental.

This restates a qualitative position already in the literature. The observation that crowds do not conserve momentum, that individuals stop and start at will, and that they are not reducible to the equations governing particles in a viscous medium, was advanced as a modelling principle some decades ago. The contribution here is to render that principle as an arithmetic comparison that a reader can check, which strengthens rather than displaces it.

9.4 A consistency requirement for simulation

The framework derives borne load from the contact graph. It requires no motion model to do so, and this yields a check that can be applied to models which do have one.

The dependency runs in one direction. Contact forces are determined by which bodies touch; which bodies touch is determined by where bodies are; where bodies are is determined by how they moved. Force is downstream of motion. A model whose motion produces the wrong configuration therefore cannot produce the right contact loads, however carefully its contact formulation is constructed — the contact formulation is acting on the wrong graph, and refining it does not help.

The check. For any configuration a model produces, its own geometry determines a chain depth. The load it reports as borne must be consistent with that depth under Eq. (25), at some admissible transmission coefficient. Where it is not, the model is internally inconsistent: its reported forces and its reported positions cannot both be right.

What this is and is not. It is a necessary condition, not a sufficient one — a model can satisfy it and still be wrong in other respects. It is also not a claim of priority over any model's dynamics. The framework offers no motion model of its own and takes no position on whose is better. What it offers is a relation that any of them must satisfy, derived independently of all of them, and the independence is what makes it usable as a check.

9.5 One relation, two geometries

Compressive incidents of quite different appearance share an initiating condition, and the relation governing the load turns out to be shared as well — evaluated, in the second case, at parameters the geometry fixes rather than the scene supplies.

Initiation is common. In each case a perturbation exceeds an individual's postural margin while the space required to recover from it is occupied. That is the condition of §2.1, and it holds whether the subsequent loading is borne horizontally or vertically.

A column fixes both parameters of the transmission. Where bodies come to rest one upon another, the exertion is not a push but a weight, so the C-value sits at the passive ceiling; and the proportional residual is unity, because each of the channels by which force leaves the axis is absent. A column offers no lateral component to lose. The fraction taken vertically is not a loss here, because the axis *is* vertical — force taken along it is not taken off it. And no void terminates a path, because gravity maintains the contacts. Substituting into Eq. (25) at $\delta = 0$,

$$F = C_{\text{passive}} \cdot F_{\text{den}} \cdot n = \frac{1}{\mu} \cdot \mu W_0 \cdot n = n\, W_0. \tag{43}$$

The cancellation is not a discovery. The passive ceiling was *defined* as body weight expressed on the scale, so it unwinds when multiplied back by the scale. What Eq. (43) shows is not a result about columns but a property of the framework's constructs: they return static equilibrium in the geometry where equilibrium is the whole of the answer. Each body bears its own weight and everything above it, and nothing else was ever available to it.

Why this matters for the framework's scope. A treatment that marked stacking as a different mechanism would be conceding that its relations describe one geometry only, and would have no answer to the obvious question about a pile — where is the chain. The answer is that the chain is vertical, that its depth is read off the configuration rather than inferred, and that the case therefore carries no free parameter at all. What differs between the two geometries is not the law but which of its inputs the scene is free to vary.

The word "vertical" is doing two jobs. In a horizontal chain, force taken vertically is a loss: it passes through the bearing individual's feet into the ground and does not continue along the axis. In a column it is the load path itself. The two usages are unrelated, and the coincidence of vocabulary has caused error before (§5.3.1).

9.6 Time, and what it does to a density criterion

The drift described in §5.2.2 has a consequence that bears directly on the form in which criteria are written, and it is sharper than the objection raised in §1.1.

At fixed density the load is not fixed. A crowd held at one density for an hour is not delivering at the end of that hour what it delivered at the start. The exertion has fallen and the leaning has replaced it, and the second is the larger. Nothing about the density records this, because the density has not changed.

The qualification of §5.2.2 applies here as everywhere: what a lean delivers is subject to the same alignment factor as any other exertion, so the rise is a rise in what is available to a chain rather than in what any chain necessarily bears. Where the crowd is oriented, the two coincide. The point stands either way, because the density records neither.

The objection to a density criterion is therefore not only that density is a proxy for the quantity that matters. It is that the proxy is stationary while the quantity it stands for is not. A relation in which one side moves and the other does not is not an imprecise relation; it is not a relation.

Which raises the share attributable to geometry. If the load at a given density depends on how long the configuration has persisted, and how long it persists depends on whether a free boundary can clear it (§5.8.1), then the arrangement of the space is doing more of the work than a density-based account can attribute to it. Geometry sets the run length, sets where relief can enter, and sets how long the configuration lasts. Density sets none of these.

And it explains a pattern that otherwise puzzles. Serious incidents follow prolonged crowding rather than brief crowding. Read through active exertion alone this is difficult to account for: people tire, so a crowd that has been standing for an hour should be pushing less hard than one that has just arrived, and the risk should fall rather than rise.

The resolution is that the source of the load changes rather than the amount of effort. Muscular exertion does fall. What replaces it is weight, which does not fatigue and which delivers more.

Duration is therefore not incidental to these events; it is the interval over which one source of load is exchanged for a larger one.

A consequence for models that compute contact forces. An agent-based or discrete-element treatment obtains its forces from the overlap between bodies, and can in principle represent a crowd that leans. What such formulations do not carry is any mechanism by which a body ceases to hold itself up. Their agents are as capable at the end of an hour as at the beginning, and pursue their assigned velocity with undiminished capacity throughout.

Results from such models therefore describe a crowd that does not tire. That is a fair description of the opening minutes of an event and a poor one of its later hours, which is when the incidents this framework is concerned with occur. The point is not that the contact mechanics are wrongly computed; it is that they are computed for a body in a posture the crowd has by then left.

9.7 Limitations

1. **Domain.** The framework treats the solid phase. Nothing here applies below the crossing of Definition 1, and no claim is made about that regime.

2. **Chain pitch.** The spacing between successive load-bearing contacts is derived from assumed body dimensions taken from outside this work. Every chain-length result scales linearly with it.

3. **The C-value has no direct measurement.** It is declared rather than measured, and the operational proxy of §6.3 establishes whether accumulation is occurring, not how fast. Design-side results depend on it only through an upper bound (§3.4); forensic results depend on it directly.

4. **The holding density is undefined.** Where a capacity finding requires the density at which a population may be held rather than merely contained, this framework supplies no value, and findings are accordingly parameterised over it.

5. **The transition itself is not modelled.** The framework describes the upright configuration (§5) and the stacked one (§9.5), and gives the condition under which recovery from a perturbation is unavailable (§5.8). It does not describe the passage from one configuration to the other: how many individuals leave the upright state, in what order, or over what interval. A reconstruction requiring that passage requires something this framework does not supply.

6. **The treatment is quasi-static.** Impulsive contributions from a moving crowd meeting a stationary front are not represented. §9.3 argues that they are small relative to sustained exertion at the relevant speeds; it does not argue that they are zero, and a formulation covering sudden arrivals would need to carry them.

7. **The accumulating quantity changes within the solid phase.** While interstitial space remains, the ratchet advances by closing it; once it is exhausted there is none left to close and further advance appears as force (§2.1). An integral formulation over the

whole phase requires a single quantity defined across that boundary, and the framework does not supply one. The stored compressive load is a candidate, being defined in both stages; a piecewise definition is the alternative and is not adopted here, since a quantity that changes identity midway through an integration is not one a reader can be asked to trust.

8. **Which mechanism bounds self-generated exertion is unsettled.** The ceiling used here is derived from purchase against the ground (§4), so that the exertion limit and the limit on the C-value coincide by construction. A single measurement suggests that at the contact heights which occur in crowds the binding constraint is postural instead, and lies below it: equilibrium against a reaction applied at chest height requires a forward displacement of the centre of mass that the body cannot supply, and the subject stepped rather than slid well short of the frictional limit. If that holds generally, the framework's worst case for this source overstates the load. The direction is the characteristic one — conservative where a guarantee is being issued, anti-conservative where exposure is being asserted — so the two uses should not share a single value until the question is settled. One self-measured subject is not grounds for changing the definition; it is grounds for not relying on the present one in the forensic direction.

9.8 The evidential situation

It is worth stating plainly what evidence exists for the account given here, because the answer is unusual and a reader is entitled to weigh it.

There is very little, and what there is cuts both ways. Quantitative data on crowds treated as a mechanical system are scarce, old, and in several cases traceable only through relays whose terminus cannot be confirmed (§7). No body of measurement supports the relations of this paper. Equally, and for the same reason, none is capable of refuting them. The deficit is symmetric, and it would be a mistake to present it as though only one side of the argument suffered from it.

What the deficit does not do is leave the question open in every direction. A framework can be assessed on its internal consistency, on whether its derivations follow from their premises, and on whether its predictions agree with what is observed in the conditions that do occur. Those assessments do not require the measurements that are missing, and this paper has tried to make all three available to a reader.

What breaks the symmetry is not data but disagreement. The relations here are not merely compatible with the observed record. They conflict with several conclusions that are widely held — that risk is a function of density; that a dense crowd is dangerous in proportion to how dense it is; that reducing the driving force disperses a compacted configuration; that an assumed egress rate can be applied to a configuration regardless of its state. Each of those is contradicted here, and the contradiction is derived rather than asserted.

A prediction that conflicts with an established one is worth more than a prediction that agrees with it, because it can be wrong in a way the reader can check. §9.2 identifies one such conflict and compares it against a body of ordinary experience large enough to have settled the matter if the established view were right.

So the position is a research programme rather than a settled result. The claim made is that the derivations are sound, that they follow from two declared quantities and geometry, and that where they disagree with received criteria the disagreement is checkable. That is what a framework can offer before the measurements exist. What it cannot offer is the measurements, and this paper does not pretend otherwise: the quantities it lacks are named in §9.7 and the ways of obtaining them in §9.9.

9.9 Further work

Five directions follow from the material above. None requires an input the framework does not already contemplate, and two of them — the transition density and the recovery identity — require no measurement at all.

Ground slope. The exertion ceiling depends on the friction available, which a gradient reduces; a gradient also introduces a gravitational component aligned with the chain and requiring no exertion to sustain. Both effects are computable from the quantities already carried here. Their treatment is deferred to a companion paper, pending a determination of whether the gravitational component enters the chain directly or acts only by reducing the margin.

The recovery identity. Whether the free volume required for postural recovery and the free volume whose exhaustion permits growth are the same quantity is stated in §5.8 as a candidate theorem. A proof would collapse the alignment property and the recovery gate into one result and would require no new external input.

Transmission under progressive compaction. The proportional residual is treated here as fixed within a scene. If tissue compaction raises it as loading proceeds, then during filling both the chain depth and the transmission efficiency increase together, and the load rises faster than either alone accounts for. Whether that is so is an empirical question about the receiving side.

The transition density. Definition 1 states the criterion separating the two phases but does not evaluate it. The density at which the free volume required for a step is exhausted has not been derived, and it is the single quantity whose absence is felt most widely.

It is wanted for two reasons. It is the only quantity in this framework that would bound an occupancy: the load relations are expressed in run length and contain no headcount (§8.3), whereas the phase criterion is a density and a density multiplied by an area is a number of people. And it is what would let the area-valued questions be answered at all — how much receiving space a configuration requires, whether a boundary can clear as fast as its interior fills — which at present must be parameterised over an undetermined density rather than resolved.

The derivation appears to require no external input. Writing the step requirement as a fraction of the area reachable from a given support base, the transition density is the packing limit reduced by that fraction, and both terms are geometric. What it needs is a defensible statement of the area a recovery step occupies, which is an anthropometric question of the same kind as those already carried in the pitch.

The transmission coefficient admits direct measurement. The framework carries the proportional residual without a value (§5.3), and the per-contact map that defines it is affine, so a controlled load applied to one side of a standing person and measured on the other yields both of its parameters as the slope and the intercept of a line. Repeating the arrangement with two, three and four people in series tests the composition rule, whose closure property is what justifies the affine form in the first place (§5.2). That the derivation is sound does not establish that each contact is in fact affine, and the experiment would decide it.

Experimental work on the propagation of controlled impulses through standing crowds bears on the same quantity and should be read against any such determination. The present author has not assessed it, and no claim is made here about what it shows.

The fixed residual. Proposition 4 makes this quantity a criterion rather than a bound, and a criterion requires a value (§3.2). It is mechanical rather than physiological and admits direct measurement under controlled loading. Obtaining it is the most consequential single measurement the framework currently lacks.

10. Conclusion

A crowd before and after the exhaustion of its free volume is not one system at two parameter values. Below the crossing a disturbance is absorbed by whoever receives it; above it the step that would absorb the disturbance has nowhere to land, and the disturbance is passed on instead. This paper treats the second regime and makes no claim about the first.

The restriction is narrower than it appears and it carries an explanation with it. Models representing contact by a bounded repulsive potential admit interpenetration and therefore never jam; experiments on dense crowds are conducted below the crossing because the regime above it cannot ethically be produced in volunteers. That the mechanism described here is absent from both is a property of their domains rather than evidence about this one.

What the framework establishes is the load a configuration can produce and how that load accumulates along a coherent path. What it does not establish is which loads injure. Every physiological value appears in one chapter, enters no derivation, and can be deleted without disturbing any result elsewhere. Once that concession is made, and once egress rate is conceded at whatever value the governing guidance permits, what remains load-bearing is geometry, conservation and monotonicity — and a party wishing to contest a finding must contest its own premises rather than the framework's.

Three quantities carry the analysis and they answer different questions. The C-value is the input: what the scene supplies. The state and its rate are the process: what accumulates, and whether it is still accumulating. The exposure integral is the conclusion: what was borne, and for how long.

Separating them is what allows the third to be constructed without assuming an answer to the question it is asked to settle. Before an event the deliverable is not a value of that integral but a demonstration that it vanishes — either because the geometry denies a chain the length it would need, or because the growth condition fails and no length suffices. Both routes are evaluated at a bound rather than an estimate, so the demonstration holds whatever the transmission proves to be, and it holds for every choice of dose function. After an event the deliverable is the trajectory, to which a reader applies the dose function they accept. The framework supplies the argument of the integrand; it does not supply the integral.

Underlying these is a single constraint, applied three times:

> An indicator must not depend on a quantity the framework does not possess.

It is why the injury thresholds are catalogued rather than adopted; why rescue capability is not a premise, the framework's own relations bounding the access on which rescue depends; and why egress rate carries no weight in a capacity finding, its failure transferring the shortfall upstream rather than merely withholding a benefit. The three are not separate positions arrived at separately. They are one requirement, and the architecture of this paper is what satisfying it costs.

Data and Code Availability

This paper is a specification and a derivation. It analyses no primary data, and every numerical value it reports either follows from the two declared quantities of §1.4 or is catalogued in §7 with its provenance stated. The relations are given in closed form so that a reader may reproduce any figure in the paper without access to code.

Conflict of Interest

The author is the founder of PneumaTheorem, Inc., a Delaware corporation which offers commercial services in crowd-risk assessment. This framework is published open-access under CC BY 4.0 and its use is not conditioned on any commercial relationship.

This paper is not a commercial engagement, an expert-witness statement, or a regulatory certification, and it does not assert any certifying status. It offers no judgement as to whether any particular configuration is acceptable; §7 states why that judgement is left to the reader and to the body setting the applicable standard.

References

M. Björkstén and B. Jonsson. Endurance limit of force in long-term intermittent static contractions. *Scandinavian Journal of Work, Environment and Health*, 3(1):23–27, 1977.

Roger W. Byard, Regula Wick, Ellie Simpson, and John D. Gilbert. The pathological features and circumstances of death of lethal crush/traumatic asphyxia in adults—a 25-year study. *Forensic Science International*, 159(2–3):200–205, 2006. doi: 10.1016/j.forsciint.2005.08.003.

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

S. G. Fattal and L. E. Cattaneo. Investigation of guardrails for the protection of employees from occupational hazards. Technical Report NBSIR 76-1139, National Bureau of Standards, Washington, DC, jul 1976. 114 pp. Prepared for the Occupational Safety and Health Administration.

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

John J. Fruin. The causes and prevention of crowd disasters. In R. A. Smith and J. F. Dickie, editors, *Engineering for Crowd Safety*. Elsevier, Amsterdam, 1993. Originally presented at the First International Conference on Engineering for Crowd Safety, London, March 1993.

James R. Gill and K. Landi. Traumatic asphyxial deaths due to an uncontrolled crowd. *American Journal of Forensic Medicine and Pathology*, 25(4):358–361, 2004. doi: 10.1097/01.paf.0000147316.62883.8b.

Mats Hagberg. Muscular endurance and surface electromyogram in isometric and dynamic exercise. *Journal of Applied Physiology*, 51(1):1–7, 1981.

Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282–4286, 1995.

Dirk Helbing, Illés Farkas, and Tamás Vicsek. Simulating dynamical features of escape panic. *Nature*, 407(6803):487–490, 2000.

Dirk Helbing, Anders Johansson, and Habib Zein Al-Abideen. Dynamics of crowd disasters: An empirical study. *Physical Review E*, 75(4):046109, 2007.

International Commission on Radiological Protection. Basic anatomical and physiological data for use in radiological protection: Reference values. ICRP Publication 89, ICRP, 2002.

Anders Johansson, Dirk Helbing, Habib Zein Al-Abideen, and Salim Al-Bosta. From crowd dynamics to crowd safety: A video-based analysis. *Advances in Complex Systems*, 11(4): 497–527, 2008.

Andrea J. Liu and Sidney R. Nagel. Jamming is not just cool any more. *Nature*, 396:21–22, 1998.

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

Bradley A. Maron, Tammy S. Haas, and Barry J. Maron. Sudden death from collapsing sand holes. *New England Journal of Medicine*, 356(25):2655–2656, 2007. doi: 10.1056/NEJMc070913. Correspondence (research letter).

Svend Erik Mathiassen and Elisabeth Åhsberg. Prediction of shoulder flexion endurance from personal factors. *International Journal of Industrial Ergonomics*, 24(3):315–329, 1999.

Corey S. O'Hern, Leonardo E. Silbert, Andrea J. Liu, and Sidney R. Nagel. Jamming at zero temperature and zero applied stress: The epitome of disorder. *Physical Review E*, 68(1): 011306, 2003.

Walter Rohmert. Ermittlung von erholungspausen für statische arbeit des menschen. *Internationale Zeitschrift für angewandte Physiologie einschließlich Arbeitsphysiologie*, 18:123–164, 1960. doi: 10.1007/BF00698869.

Walter Rohmert. Problems in determining rest allowances. part 1: Use of modern methods to evaluate stress and strain in static muscular work. *Applied Ergonomics*, 4(2):91–95, 1973. doi: 10.1016/0003-6870(73)90082-3.

G. Keith Still. *Crowd Dynamics*. PhD thesis, University of Warwick, 2000.

G. Keith Still. *Introduction to Crowd Science*. CRC Press, 2014.

Wenjian Yu and Anders Johansson. Modeling crowd turbulence by many-particle simulations. *Physical Review E*, 76(4):046105, 2007.

Acknowledgements

The author is grateful to Dr G. Keith Still for extensive correspondence during 2026. Three contributions shaped the development of this work: first, the identification of the interval between the end of loading and the physiological outcome as a distinct phase requiring its own representation, rather than one collapsed into the loading event itself; second, the sourcing of published threshold figures to the mechanisms they describe, together with the distinction between live-subject and post-mortem tolerance data that such sourcing makes visible; third, the framing of injury mechanisms in which the trajectory is established before external loading ends.

Beyond these, Dr Still directed the author to material the work required, and to questions it had not asked. What has been made of that material, and what has not yet been made of it, are the author's own.

The framework has been substantially restructured since that correspondence, and the present formulation has not been put to Dr Still. The exchange took place as private correspondence. Any errors of interpretation, and any extension of the framework beyond the points Dr Still raised, are the author's own. Dr Still bears no responsibility for the conclusions drawn here, nor for any downstream application of the framework.

Author: Yixuan Tao

Correspondence: kevin@pneumatheorem.org

ORCID: 0009-0005-8136-3886

Web: pneumatheorem.org