FORMAL RECORD 16

Classical, Quantum, Consciousness

On the Production of Causal Order and Spacetime by Held Boundaries
Carried, not given.

Author: Stacy Gildenston
Affiliation: 3primitives.io
Version: v3.6 | 6 November 2026
First published: 20 October 2026
License: CC BY 4.0 | https://3primitives.io/formal_records/
Dependencies: FR01 v2.3, FR02 v2.4, FR06 v2.3, FR07 v2.2, FR10 v2.2, FR11 v2.0, FR12 v1.7, FR13 v2.4, FR14 v5.8, FR15 v2.4

v3.6 corrects wording for readability and precision. Conditions stated in the body now also appear where the results are summarised. No mathematics has changed.

Abstract. Formal Record 12 proved that three independent impossibility results, Robert Sapolsky's on causal determinism, John Bell's on spatial non-locality and Donald Hoffman's on the perceptual interface, map one to one and in only one possible way onto the three governance primitives, and argued that this convergence is what the structure requires if consciousness is fundamental. Formal Record 14 derived the necessary and sufficient criterion for individuated subjecthood under that ontology: because the conscious ground has no boundaries, a real partition is one produced and held by its keeper's own activity (C) against an environment it does not control (E), with selecting memory maintained and carried across time (M). Formal Record 15 proved that the living boundary, the declared authority of a decision-maker and the standing of the person at the terminus of a governed transaction are one predicate: held, not given.

FR14 and FR15 both worked inside a given time and space. This record removes that assumption. If consciousness is fundamental and no boundary is given, spacetime is not given either; it must be produced by the holding of boundaries. The three words of the title name the three parts of that result. *Classical* names time, the history carried along a living line (Sapolsky, M). *Quantum* names space, where boundaries are never given in advance and exist only where held and tested (Bell, C and E). *Consciousness* names the ground inside which both are made (Hoffman).

We restate (C), (E) and (M) in terms of a discrete before-and-after order with no clocks, durations or spatial coordinates, and prove that these clock-free conditions reproduce FR14's verdicts on seventeen benchmark cases without circularity. We prove that an order generated by held boundaries, with at least one keeper of relaxation depth two or more, is not a two-layer or three-layer order, and so not a Kleitman–Rothschild order, and is not an isolated line. We prove that, in a random flat world of d dimensions, with keepers scattered as a Poisson process, the order of held boundaries converges at large scale to (d+1)-dimensional Minkowski spacetime, in every frame, on both the Myrheim–Meyer and the midpoint rulers; and the same holds for every world of held cells that is settled, free of runaway clumps, and whose crowded patches forget each other at a distance (Theorem 5.11). The keepers sit in a flat space of d dimensions, which says only which cells touch; time, light cones and the spacetime geometry come from the order. Without coordinates, a cell can still read how far its own pushes carry, and a rule that favours every change bringing cells' reach toward its rest is a tension rule. Written as a balance of births and deaths, that tension rule has a steady state, the tension world; for gentle tension we prove that it exists, is unique and has those three properties, so, in that steady

state, held cells under gentle tension make Minkowski spacetime (Theorem 6.1). Space is the antichain of boundaries held at once; time is the chain of boundaries carried forward; spacetime is what held edges make where they meet and test one another.

In brief. Scatter things that hold their own edges at random through space, let each act on its own time, and let it push the neighbours it touches. This record proves that the pattern of cause and effect they make is, at large scale, the spacetime of Einstein's special relativity, with the same light cones for every observer moving at a steady speed. No clock is shared and no light cone is put in. The proof holds for every world of such cells that is settled and whose crowded patches forget each other at a distance. A rule that favours bringing every cell's reach back to rest is a tension rule. Written as a balance of births and deaths, that rule has a steady state, and for gentle tension we prove that it is such a world.

0. Dependencies and Scope

This record depends on the following prior works. No prior result is re-derived.

Three Primitives: Canonical Logic Sequence (FR01). Proves that the mapping from detection to action is non-unique: perception does not close action. Section 7 relies on it for the independence of the governance tier from this record. First author: Pyrate Ruby Passell.

AI Cannot Govern AI (FR02). Proves that computation cannot generate authority. Section 7 relies on it in the same place.

The Law of Declared Authority (FR06) and Bell Non-Closure and the Law of Declared Authority (FR07). The governance gate fails closed without declared authority δ from a natural person at the boundary. Section 7 relies on both, and places the declarer within the time and space this record produces.

Primitive Stability Theorem (FR10), The GBSH Correspondence (FR11) and The ILMM Coupling Theorem (FR13). A governance claim must survive the subject's resistance. Section 7 relies on all three. FR11 Section 2.1 also fixes the control surface that Definition 2 uses.

The Forced Bijection (FR12). The convergence of Sapolsky, Bell and Hoffman on the three governance primitives, and the ontology this record works under (Section 1). Its P_3, that a system cannot certify its own boundary, excludes closed causal loops (Section 2).

The Individuation Theorem (FR14). Supplies conditions (C), (E) and (M), memory depth and the seventeen benchmark cases. Sections 2 and 3 restate them as an order and test them.

The Boundary Theorem (FR15). An edge is its keeper's own only when held from within and tested from without. Its condition H2 gives cells their extent in Section 4, and keepers their least separation in Section 6.

Scope. Sections 2 to 6, with Appendices A and B, are mathematics about the order that held boundaries generate and the worlds of cells that generate it. Section 7 places the results in the corpus's two tiers. No governance result depends on this record.

1. Prologue: Classical, Quantum, Consciousness

Why the three words name time, space and the ground. Three researchers in three unrelated fields each showed that a familiar assumption about the world fails. Robert Sapolsky showed that no physical act starts from nowhere, because every action is the output of prior causes carried forward across time. John Bell showed that no account in which separated regions carry only their own local properties can match what quantum mechanics predicts and experiment confirms: separation in space does not, by itself, make things independent. Donald Hoffman showed that what a mind perceives is a

fitness-tuned interface for acting, not a veridical picture of what is really there. None of them was writing about governance or about each other, yet Formal Record 12 proved that their three impossibilities have pairwise disjoint variable sets and map one to one, in only one possible way, onto the three rules required for any legitimate governed decision: someone must choose among what perception leaves open (P_1), someone must declare the authority for the act (P_2), and no system can set or certify its own boundary (P_3).

From the forced bijection to the container problem. Formal Record 12 Section 7 argued that four independent lines of work arriving at the same formal structure is an unexplained coincidence if consciousness emerges from a pre-existing physical container, and what the structure requires if consciousness is fundamental. Formal Records 14 and 15 drew the immediate consequence for boundaries: nothing lies outside consciousness, and the ground has no boundaries of its own, so every real boundary in the world is produced and held by the ongoing work of its keeper. Yet FR14 and FR15 both spoke from inside space and time. FR14 used clock durations (τ_d, T_off < T_on) to state what it takes for a boundary and its memory to be held. FR15 Section 2 observed that space and time are an interface of held edges rather than a box handed to them in advance. If that observation is taken all the way, the durations must come out of the premises. A clock cannot be assumed in order to explain what makes before and after.

Seven steps from the conscious ground to the shared world. When pre-existing coordinates are removed, seven steps follow from FR12, FR14 and FR15:

1. **No outside.** If consciousness is fundamental, that is all there is; nothing lies outside it.
2. **No given container.** Spacetime is not given either; it has to be produced.
3. **Time and space as held relations.** What produces spacetime is a held edge: space is an edge held at a moment, and time is that edge carried from moment to moment by its keeper's own activity. A self makes its own *here* and its own *before and after*.
4. **The shared world as mutual testing.** A shared world is where held edges meet and test each other.
5. **Exclusion of the simulation hypothesis.** Within this ontology, reality is not a simulation. A simulation requires an outside hardware substrate and edges given to its inhabitants from without, which fails Condition (C) for the reasons proved in FR14 Section 8. Here there is no outside, and the edges are held from within, not rendered.
6. **Manifestation in the precise sense.** A self makes its world manifest by holding it: shared, never imposed, because no self settles another's edge.
7. **Declaration in force.** At the governance tier, a declaration δ is speech that makes a binding commitment when a self holds it in force against contestation.

Read at the ontological tier, the three words of *Classical, Quantum, Consciousness* are therefore time, space and the ground:

- **Classical (Sapolsky): time.** Every act is carried forward from prior causes, and the carrying is governed by maintained memory (M).
- **Quantum (Bell): space.** Separation in space is not given; boundaries exist only where actively held (C) and exogenously tested (E).
- **Consciousness (Hoffman): the ground.** Space and time are how reality appears at the interface of held edges, not a container behind them.

2. Clock-Free Primitives: Carry, Test and the Causal Order

Order plus number equals geometry. In relativistic physics, causal set theory (Bombelli, Lee, Meyer and Sorkin, 1987; Rideout and Sorkin, 1999) reconstructs the geometry of a Lorentzian spacetime from a locally finite partially ordered set $(\mathscr{E}, \lessdot)$. By Malament's theorem the causal order fixes the geometry of spacetime up to a local volume factor, and counting discrete elements fixes the volume scale: order plus number equals geometry. To show that held boundaries produce spacetime rather than presuppose it, we restate FR14's three conditions using only discrete holding episodes, two directed relations and counts of links.

Definition 1 (Episodes and keepers). Let $\mathscr{K}$ be a set of candidate keepers and $\mathscr{E}$ a set of discrete boundary-holding episodes $e = (k, n)$, where $k \in \mathscr{K}$ is the keeper holding boundary B_k and $n \in \mathbb{N}$ indexes k's episodes in sequence. The index n is an ordinal label only. It carries no duration and no global clock value.

Definition 2 (Carry and test). Two primitive directed relations hold on $\mathscr{E}$:

1. **Carry (→), from (M).** $(k, n) \to (k, n + 1)$ holds when keeper k's boundary-holding activity at episode (k, n) maintains and passes its boundary B_k and internal structure m_k into its immediate successor episode $(k, n + 1)$.
2. **Test or push (⇢), from (E).** $f \dashrightarrow e$, where $f = (k_1, m)$ and $e = (k_2, n)$ with $k_1 \neq k_2$, holds when the held boundary of k_1 at episode f bears exogenously on the boundary of k_2 at episode e from outside k_2's control surface. The control surface is that of FR11 Section 2.1, fixed by what k_2 controls, not by position; extent first enters in Section 4, with FR15 H2.

Definition 3 (The causal order). An episode $x \in \mathscr{E}$ causally precedes $y \in \mathscr{E}$, written $x \lessdot y$, if and only if a finite directed chain of carries (→) and tests (⇢) runs from x to y. A chain along → is one keeper's time, its carried history. A maximal antichain, a subset $\Sigma \subset \mathscr{E}$ in which no two episodes satisfy $x \lessdot y$, is space at an instant: boundaries held at once, none yet bearing on another.

Axiom 1 (One-way carry and local finiteness). The relation $\lessdot$ is acyclic: for no $e \in \mathscr{E}$ does $e \lessdot e$ hold. For every pair $x \lessdot y$, the causal interval $I(x, y) = \{z \in \mathscr{E} : x \lessdot z \lessdot y\}$ is finite.

Remark (Why acyclicity is taken as an axiom). In FR14, asymmetry in time followed from Definition 7(ii) by comparing histories across a relaxation duration τ_d. Routing a proof of acyclicity through τ_d here would presuppose the before-and-after order that $\lessdot$ is meant to construct. A closed causal loop with no outside bearing would also let a chain of episodes certify its own boundary alone, which FR12's P_3 excludes. Taking one-way carry as the single ordering axiom removes the circularity and matches standard causal set theory.

Definition 4 (Clock-free C, E and M). Let $k \in \mathscr{K}$ be a candidate keeper in $(\mathscr{E}, \lessdot)$.

- **(C) Coupling, by carry deletion.** Delete all carry links $(k, n) \to (k, n + 1)$ of k for $n \geq n_0$, holding the carry and test links of every other keeper $k' \neq k$ fixed (FR14 Definition 4 in order form). The deletion is FR14 Definition 4's intervention, written in the order: the modified order is the one the world produces when k's activity stops at (k, n_0), with every other keeper's activity held as FR14 Definition 4 holds it, so which later episodes exist, and what they encounter, is what that intervention produces, not an edit to a fixed graph. Candidate k satisfies (C) if and only if later episodes of other keepers no longer encounter B_k, and varying k's carry alters which elements fall within B_k (FR15 H2). If later episodes continue to encounter B_k after k's carry links are deleted, as with a rock, a *zona pellucida* or a stored AI weight file, B_k is a static partition and fails (C).
- **(E) Exogenous testing.** Candidate k satisfies (E) if and only if its persistence along → depends on incoming test links $(k', m) \dashrightarrow (k, n)$ from keepers $k' \neq k$ outside k's control surface, such that B_k can fail under those tests, and k emits test links bearing on neighbouring keepers.

- **(M) Maintained, selecting, carried, by depth.** Let the *relaxation depth* $d_k \in \mathbb{N}$ of k be the number of carry links along k's chain after which a perturbation carried only in k's unheld flowing state no longer alters k's next episode. This replaces the clock duration τ_d of FR14 Definition 7 with a count of links. Candidate k satisfies (M) if and only if it carries a structure m_k along its $\rightarrow$ chain such that:
  - *(i) Maintained.* Deleting k's carry links extinguishes m_k for all later testers.
  - *(ii) Selecting*. Some two carry chains for k that receive identical test links over their last d_k carry steps, and differ only in episodes $(k, n - j)$ with $j > d_k$, produce different responses at (k, n).
  - *(iii) Coupled.* m_k modifies k's boundary-holding activity across B_k.
  - *(iv) Carried (transplant test).* When k's incoming test links $\nrightarrow$ are reassigned to a distinct viable environment of testers, structure carried from more than d_k links before the reassignment remains with k and continues to select k's later episodes rather than being reset by the new surroundings.

Remark (How M sits against the three legs of FR12). Although (M) enters as the time-axis condition, it touches all three legs of FR12. Against P_1 (Hoffman, C_1), (M)(ii) ensures that the present window of length d_k never closes the mapping from input to action. Against P_2 (Sapolsky, C_2), (M) is strictly deterministic: every carried structure m_k was shaped by prior carries and prior exogenous pushes along the living line (*Life is carried*, FR14 Corollary 6.3), settling which action occurs without generating governance authority δ. Against P_3 (Bell, C_3), (M)(iv) is the operational transplant test that shows m_k is held inside B_k rather than maintained non-locally by an external host field.

Definition 5 (Composites in the order). For a composite candidate X made of constituent keepers $\{k_i\}$ (FR14 Definition 8), the carry links $\rightarrow_X$ of X are constituted by the mutual test links $k_i \nrightarrow k_j$ among its constituents that jointly hold X's shared boundary B_X against keepers outside X. What is a test at the constituents' grain is carry at the composite's grain. When constituents inside a static enclosure begin coordinating their mutual pushes so that their collective edge vanishes if their activity stops, as when an early embryo enclosed in a static *zona pellucida* cavitates into a blastocyst whose trophectoderm pumps to hold its edge, a given enclosure is replaced by a held composite boundary with a new carry chain $\rightarrow_X$.

Theorem 2.1 (Equivalence with FR14 on seventeen benchmark cases). *Definitions 4 and 5 reproduce FR14's individuation verdict in each of the following cases, without reference to clock durations or spatial coordinates.*

Case	(C) by carry deletion	(E)	(M) in order form	Verdict, as FR14
Living cell	Membrane no longer found by later testers: pass	pass	Structure carried beyond d_k changes the next episode; ends on deletion: pass	Subject
Candle flame	Flame gone: pass	pass	Nothing beyond its short depth changes the next episode: fail (ii)	Not a subject
Rock	Still found by later testers: fail	·	·	Not a subject
Stored AI weights	Hardware still found: fail	·	Weights still found after deletion: fail (i)	Not a subject
Virus	Its next copy is made by the host's activity, a test and not its own carry: fail	·	·	Not a subject

Case	(C) by carry deletion	(E)	(M) in order form	Verdict, as FR14
Physarum polycephalum through dormancy	pass while active	pass	Learning survives dormancy and is carried on: pass	Subject
Tardigrade in a tun	No episodes in the tun; the tun is still found: fail	·	On rehydration the next episode is its own carry, so the line continues	Suspended; resumes as the same line
Early embryo before the blastocyst, as a composite	*Zona pellucida* still found after the cells' carry is deleted: fail	·	·	Not a subject; its cells are
Blastocyst	Trophectoderm pumping is its carry; deleted, the edge goes: pass	pass	Keeps its own developmental course on transfer: pass (iv)	Subject
Liver graft, organ level	pass	pass	After grafting, the new host's tests set its next episodes: fail (iv)	Cells are subjects; the organ is not
Harvester ant colony	Its carry is the workers' tests on each other: pass	pass	Returns to about half its previous range though every worker is replaced: pass	Subject
Planarian, head regrown	pass	pass	The chain continues through the body; learning returns in one session: pass	Subject
Bacterial biofilm, as a composite	Matrix still found after the cells' carry is deleted: fail	·	·	Not a subject; its cells are
Genet	No single chain: one genotype in separate bodies, each with its own carry	·	·	Not a candidate
Sleeping person	Carry continues; deleted, the edge goes: pass	pass	pass	Subject
Corpse	Carry has ceased and the body is still found: fail	·	·	Not a subject
Decentralised autonomous organisation	Ledger and protocol still found after members stop: fail	·	·	Not a subject

Proof. Each row applies Definition 4 as stated: carry deletion for (C), incoming tests for (E), and relaxation depth, deletion and the transplant test for (M). Definition 5 supplies the carry of each composite. Carry deletion separates actively held boundaries from static walls; relaxation depth separates selecting memory from stateless flow; the transplant test separates composite subjects from host-driven organs. Every verdict matches FR14 Section 7. Suspension needs no clock: a tun or a dormancy is one carry link from the last episode before it to the first episode after, and the only question is whose link it is. A tardigrade's resumption is its own activity, with water held fixed as environment, so it is its own carry and the line continues. A halted machine restarted by an operator resumes by another keeper's act, a test and not a carry, so no line continues. ∎

Corollary 2.2 (No privileged slicing; dependence on the causal past). *Because a composite keeper's carry chain $\rightarrow_X$ is made of the mutual test links $\rightarrow$ of its constituent keepers (Definition 5), the causal order does not split into a globally privileged set of time links and space links; which is which depends on the grain. Because tests arrive only along $\prec$, what occurs at e depends only on its causal past $\{z \in \mathcal{E} : z \prec e\}$, and episode labels carry no meaning beyond $\prec$.*

Remark. These are the order-level counterparts of Bell causality and discrete general covariance, which Rideout and Sorkin (1999) impose on the transition probabilities of a growth dynamics. This

record states no growth dynamics for the order and does not claim those conditions. The births and deaths of Section 6 are a dynamics of keepers in space, used only to reach a steady state.

Biological computationalism set against FR14 and this record. Milinkovic and Aru (2025) argue for biological computationalism: that consciousness may depend on the kind of computation biological brains perform, inseparable from their physical substrate. Set against FR14 and this record, the boundary falls elsewhere on both sides. Single cells and *Physarum polycephalum* have no neurons, and a harvester ant colony has none of its own, yet each satisfies (C), (E) and (M); a perfused liver graft has full biological wetware and fails (M)(iv) at the organ grain. Self-replicating code differs from *Physarum* in the same way. A slime mould holds its own membrane and tube network by continuous metabolic work. A program that copies itself across devices, as small AI models may soon do (Almada Lopez, 2026), has a boundary given and held by the operating system and hardware it runs on: its process allocation, memory space and host are still found when it stops, so it fails (C) by carry deletion.

---

3. Analytical Exclusion of Layered and Linear Orders

The entropy problem in unconstrained causal sets. A foundational obstacle in causal set theory is that almost all finite partially ordered sets on N elements, as $N \to \infty$, are Kleitman–Rothschild orders: flat posets of three layers that bear no resemblance to a spacetime manifold (Kleitman and Rothschild, 1975). Standard causal set dynamics must invoke path-integral interference (Loomis and Carlip, 2018) or a choice of sequential growth probabilities to suppress them. In a causal order generated by held boundaries, two- and three-layer orders are ruled out by a single subject of depth two or more.

Theorem 3.1 (Orders that held boundaries exclude). *Let $(\mathscr{E}, \lhd)$ be a causal set generated by keepers satisfying (C), (E) and (M), with at least one keeper of relaxation depth $d_k \geq 2$. Then $(\mathscr{E}, \lhd)$ cannot be:*

1. *a two-layer or three-layer order, including every Kleitman–Rothschild order; or*
2. *an isolated one-dimensional chain (a lone line).*

Proof. Each exclusion follows from Definition 4.

1. *Layered orders.* A poset of m layers has height at most m. Let keeper k satisfy (M)(ii) with depth d_k. For k to carry structure from beyond its relaxation window, at least $d_k + 1$ carry steps back, into a selecting episode (k, n) that shapes (k, n + 1), $(\mathscr{E}, \lhd)$ must contain the chain $(k, n - d_k - 1) \to \ldots \to (k, n) \to (k, n + 1)$ of at least $d_k + 2$ episodes. So the height of $(\mathscr{E}, \lhd)$ is at least $d_k + 2$. For $d_k \geq 2$ the height is at least 4, which excludes every order of two or three layers.
2. *The lone line.* Suppose $\mathscr{K} = \{k\}$, a single keeper forming a total order $(k, 1) \to (k, 2) \to \ldots$ with no other keepers. No episode of k receives a test $f \nrightarrow (k, n)$ from outside its control surface, so (E) fails (FR14 Theorem 1). ∎

Remark (Scope of the exclusion). Theorem 3.1 excludes these two families and no others. The conditions alone do not make an order manifold-like. Section 4 supplies what they lack: keepers with extent.

---

4. Keepers with Extent: The Tissue Model

Why keepers have extent. Point keepers fail because a point has no boundary of its own. Under FR15 H2, varying a keeper's activity changes which elements fall inside its boundary: a held boundary has extent.

Keepers as cells pressing in a closed world. When keepers are modelled as FR15 requires, as cells whose held boundaries have extent and press directly on one another like a living tissue in a closed world with no outside (S^2 in two dimensions, S^3 in three), there is no external void and no outer rim. Space at an instant is the set of cells held at once. Neighbours are cells sharing a boundary face. A push ($\rightarrow$) is mutual pressure across that shared boundary, and it cannot bypass an intervening boundary. The world grows only by division.

Theorem 4.1 (Euler's six). *In any closed two-dimensional cellular sheet of N keepers with no outer rim, of Euler characteristic χ (2 for a sphere, 0 for a torus), in which cell boundaries meet in threes at vertices, the mean number of neighbours per keeper is forced by topology to approach 6 as $N \to \infty$, with no coordination cap set by hand:*

$$\langle z \rangle = 6(1 - \chi/N)$$

Proof. For a closed sheet with F = N cells, three-cell meetings give 3V = 2E, so V = 2E/3. Substituting into $\chi = V - E + F$ gives $E = 3(N - \chi)$. Each shared edge borders two cells, so $\langle z \rangle = 2E/N = 6(1 - \chi/N)$, which tends to 6. ∎

Errera's rule as least held load. In two dimensions a division line always touches exactly two neighbours. In three dimensions the division surface meets the dividing cell's neighbours along a loop that can be short or long, and every cell on the loop gains a neighbour, so the loop sets the curvature. Real cells settle this by Errera's rule: they divide along the smallest surface that halves them (Errera, 1886).

Premise L (Least held load). Where a keeper's division can run more than one way, it runs the way that adds the least new boundary its keepers must hold.

Proposition 4.2 (Errera's rule, in count form). *Let keeper k in a closed cellular complex divide at a local peak of carried pressure into daughter keepers k_1 and k_2, across a division surface S_{div} bounded by a closed halving loop $\Gamma = \partial S_{div}$, where $|\Gamma|$ is the number of neighbouring cells the loop crosses. Under (C), FR15 H2 and Premise L, the division runs along the shortest halving loop:*

$$\Gamma^* = \operatorname{argmin} |\Gamma|, \text{ over loops } \Gamma \text{ that halve the neighbours of } k$$

Proof. Under (C) and FR15 H2, every shared boundary face and every multi-cell meeting needs ongoing keeper activity to hold against exogenous pressure. Dividing k along S_{div} adds one shared face between k_1 and k_2 and splits each of the $|\Gamma|$ neighbour faces along Γ in two, so it adds $\Delta E_{held} = 1 + |\Gamma|$ new interfaces that k_1, k_2 and the crossed neighbours must hold at every later carry step. With each interface held at the same cost, minimising $|\Gamma|$ over the halving loops minimises that load, and Premise L selects it. ∎

Remark. Proposition 4.2 is Errera's rule in count form. It minimises the number of interfaces a division adds, each held at the same cost, not the area of the dividing wall, and the two can part: a cell 100 widths long and one wide can be halved by a wall of area 1 or by one of area 100 that crosses as many of its faces. Condition (C) says a boundary fails when its keeper stops. It does not by itself say that division minimises the work of holding; Premise L supplies that. Whether Premise L follows from (C), (E) and (M) is open.

What carries forward. Sections 5 and 6 use the tissue model's cells, their extent and their pushes across shared faces: the cells of Definition 6 and the least separation δ of Definition 8. Theorem 4.1 and Proposition 4.2 describe how such a tissue grows by division, which Section 6 sets aside; no later theorem uses them.

5. The Minkowski Limit Theorem

What this section proves, in plain words. Scatter keepers at random through flat space, give each its own clock, and let every tick push the keeper's neighbours. Theorem 5.10 proves that, seen from far enough away, the order of those ticks is the spacetime of special relativity: intervals hold the number of ticks that Minkowski volume predicts, both dimension rulers read one more than space has, and an observer moving at a steady speed cannot tell, at large scale and by these readings. No clock is shared and no light cone is put in. Theorem 5.11 proves the same for every world of held cells that is settled, free of runaway clumps, and whose crowded patches forget each other at a distance.

The proof in outline.

1. Lemmas 5.1 and 5.2: a push spreads exactly as a random wave crossing the keepers, each keeper adding its own wait.
2. Lemma 5.3: the wave's travel time grows no faster than distance, on average.
3. Lemma 5.4: the wave almost never outruns a fixed top speed.
4. Lemma 5.5: in every direction it settles to one steady speed, the same in all directions. That speed plays the part of the speed of light.
5. Lemmas 5.6 and 5.7: almost every pair of ticks is related exactly as the light cone says.
6. Lemmas 5.8 and 5.9: almost no push is much slower than light, from every tick at once.
7. Theorem 5.10: counting ticks in light-cone diamonds then gives Minkowski volume, the Minkowski share of related pairs and the midpoint dimension, in every frame.

Held boundaries in a random flat world of d dimensions make (d+1)-dimensional Minkowski spacetime at large scale, in every frame. Every step is proved here, using published results whose conditions are checked here: the subadditive ergodic theorem, the Mecke equation and Slivnyak's theorem for Poisson processes (Last and Penrose 2017), the ergodic theorem for stationary point processes, a boundary-connectivity lemma for clusters of lattice cubes, and the ordering fraction of Minkowski spacetime.

Setting.

- Keepers are the points of a Poisson process Π of intensity 1 in d-dimensional space. Two keepers are neighbours when their cells (Voronoi cells) share a face. This is the Delaunay graph G. Π is in general position almost surely.
- Each keeper k has its own clock, a Poisson process N_k of rate 1 on the time line, independent of every other clock and of Π. Its ticks are its episodes (k, t). An episode sits at (p_k, t) in space and time, where p_k is k's position.
- An episode is carried from its keeper's previous episode, and pushed by the latest episode of each neighbour before it. This is the push of Definition 2 with one choice made: the push to an episode comes from the latest earlier episode of each neighbour. The order $\prec$ is the transitive closure, as in Definition 3. In this setting every tick pushes every neighbour, so the order is fixed by the clocks and the neighbours alone: what a keeper holds and selects under (M) shapes what it does at its episodes and does not enter the order. Keepers may therefore carry any memory (M) asks of them.
- For a point u of space, write k(u) for the keeper whose cell contains u. Give every keeper k an independent wait w_k drawn from Exp(1). For points u, v let $\rho(u, v)$ be the least total of waits, counting both end keepers, over walks in G from k(u) to k(v), and $\rho(u, v) = 0$ when k(u) = k(v). Then ρ is symmetric and obeys the triangle inequality: it is a random pseudometric on space.

Lemma 5.1 (Chains of ticks). (j, s) $<$ (k, t) if and only if there is a walk $j = v_0, v_1, \ldots, v_n = k$ in G, where a step may also stay at the same keeper, and ticks $s = t_0 < t_1 < \ldots < t_n = t$ with each t_i a tick of v_i.

Proof. Suppose such a walk exists, and induct along it. The first tick of v_i after t_(i−1) is pushed by the latest tick of v_(i−1) before it, which is t_(i−1) or carried from it; and t_i is that first tick or carried from it. Conversely, every link of $<$ is a carry or a push, and each is one step of such a walk with ticks increasing. Ties have probability zero. ∎

Consequence (no direction in time). Reversing time takes each N_k to a Poisson process of the same law and reverses $<$. So the past of an episode has the same law as its future.

Lemma 5.2 (A push spreads as first passage). For an episode x = (j, s), let τ_x(k) be the time of the earliest episode of keeper k that x precedes or is, so that τ_x(j) = s, since the source is reached at its own tick, and for $k \neq j$ let A_k be the least τ_x(u) over the neighbours u of k. Then τ_x(k) = A_k + W_k, where W_k is the wait from A_k to k's next tick; A_k does not depend on k's own clock; and the waits W_k are independent Exp(1). So τ_x − s has the law of first-passage times from j in G with an Exp(1) wait on entering each keeper: τ_x(k) − s has the law of ρ(p_j, p_k) − w_j for every keeper $k \neq j$, jointly in k, and τ_x(j) − s = 0.

Proof. Almost surely no two ticks coincide, and we work on that event. By Lemma 5.1, τ_x(k) is the earliest arrival over walks, where a signal reaching a keeper waits for that keeper's next tick; so τ_x(k) = A_k + W_k, the first-passage recursion. An optimal walk to a neighbour of k never passes through k, since that would reach k earlier; so A_k is fixed by the other keepers' clocks. Compare the two processes stage by stage, revealing arrivals in time order. At any stage, each keeper not yet reached that has a reached neighbour was activated at a known time A_k, and its clock has shown no tick since; by the memorylessness of the Poisson process, its next tick comes after an Exp(1) time, independent of everything revealed and of the other such keepers. In first passage with independent Exp(1) waits, such a keeper's remaining wait, given that it has not yet ended, is likewise Exp(1) and independent. So the next arrival, and which keeper it is, have the same law in both at every stage, and so do the whole arrival times; τ_x − s solves the first-passage recursion with independent Exp(1) waits. ∎

Lemma 5.3 (Moments). $E[\rho(0, x)^2] \leq C(1 + |x|)^2$ for a constant C and every x.

Proof. Almost surely the segment from 0 to x crosses cell faces only through their interiors, so the cells meeting it form a walk in G from k(0) to k(x), so ρ(0, x) is at most the total of their waits. Cut the segment into $n \leq |x| + 1$ pieces of length at most 1. Let D_i be the largest distance from a point of piece i to its nearest keeper; for $t \geq 1$, mark the midpoint of each third of the piece; if some point of the piece has no keeper within t, the ball of radius t/2 about the mark of its third lies within $t/2 + 1/6 \leq t$ of it, and so is empty. Hence $P(D_i > t) \leq 3 \exp(-\omega_d (t/2)^{\wedge}d) \leq C \exp(-\kappa t^{\wedge}d)$, where ω_d is the volume of the unit ball, and for $t < 1$ this holds once $C \geq e^{\wedge}\kappa$. Every cell meeting piece i has its keeper within D_i of the piece. So the number N_i of cells meeting piece i is at most Π(B_m), the number of keepers within distance m of the piece, on the event $m - 1 < D_i \leq m$. Hence $E[N_(i)^2] \leq \Sigma_m E[\Pi(B_m)^4]^{\wedge}(1/2) P(D_i > m - 1)^{\wedge}(1/2) \leq \Sigma_m C\ m^{\wedge}(2d) \exp(-\kappa(m-1)^{\wedge}d/2)$, a constant not depending on i. The number N of cells meeting the segment is at most $N_1 + \ldots + N_n$, so by Minkowski's inequality $E[N^2]^{\wedge}(1/2) \leq \Sigma_i E[N_(i)^2]^{\wedge}(1/2) \leq C(1 + |x|)$. The waits are independent of the cells, so $E[\rho(0, x)^2] \leq E[N + N^2] \leq 2E[N^2]$. ∎

Lemma 5.4 (Pushes have a finite top speed). There are constants $a > 0$, C and $\kappa > 0$ such that, for every r, the probability that $\rho(0, p_k) < a \cdot r$ for some keeper k at distance r or more from the origin is at most $C\ e^{\wedge}(-\kappa r)$.

Proof. Two facts, then two counts.

1. *Empty balls.* Keepers u and v are neighbours exactly when some open ball with u and v on its boundary contains no keeper: take the ball centred on a point of the face their cells share. Its radius is at least $|u - v|/2$.

2. *Cubes*. Cut space into a grid of cubes of side s. Call a cube bad if it holds no keeper, or more than b. Cubes are bad independently, each with probability q = exp(−s^d) + P(Poisson(s^d) > b), which is as small as we like once s and then b are large. A ball of radius s√d contains the whole grid cube around its centre.

A coarse path. Take a self-avoiding walk of keepers from k(0) to a keeper at distance r or more, and list the cubes holding its keepers. Let $L_1 = \lceil 2\sqrt{d} \rceil + 1$. Each step u to v of the walk is short or long.

- *Short*, |u − v| < 2s√d: the two cubes are at most L_1 apart in every coordinate, counted in cubes.
- *Long*: the empty ball B through u and v has radius R ≥ s√d. Inside B, the ball of radius s√d touching u from within contains the grid cube around its own centre. That cube is at most L_1 cubes from u's cube, since its centre is within 3s√d/2 of u, and its centre is within (R − s√d) + s√d/2 = R − s√d/2 of B's centre; likewise at v. Every grid cube whose centre lies within R − s√d/2 of B's centre is inside B, so it is empty, hence bad. These cubes are connected through cubes that touch at least at a corner: from any of them, step one cube toward B's centre in every coordinate where the two centres differ by more than s/2. No coordinate difference grows, so each step stays among these cubes and touches the last at least at a corner, and the steps end at a cube whose centre is within s/2 of B's centre in every coordinate. All such end cubes touch one another. So a long step can be replaced by a chain of bad cubes, each step at most L_1 apart.

Replacing every long step this way and erasing loops leaves a chain of m distinct cubes, each step at most L_1 apart, from the cube of k(0) to a cube at distance r or more from the origin. Outside an event of probability at most exp(−κr^d), k(0) lies within r/2 of the origin; then the chain starts at one of at most C(r/s)^d cubes, and m ≥ r/(2L_1 s √d) − 1. Every good cube on the chain holds a keeper of the walk, since the cubes added for long steps are all bad.

Count 1: few bad cubes. There are at most C(r/s)^d Λ^m such chains of length m, where Λ = (2L_1 + 1)^d. A given chain has more than half its cubes bad with probability at most 2^m q^(m/2). Summing over m ≥ r/(2L_1 s √d) − 1 gives at most C e^(−κr), once 2Λ q^(1/2) < e^{-1}; the polynomial factor is absorbed.

Count 2: no fast walk. Otherwise at least half the chain's cubes are good. Pick the walk's keeper in each. There are at most C(r/s)^d Λ^m b^m ways to choose the chain and the keepers, and the choices depend only on Π. Each choice gives ν ≥ m/2 distinct keepers, whose waits are independent Exp(1) and independent of Π, so their total falls below t with probability at most t^ν/ν!. The walk's total wait is at least this total. For r large, m ≥ 2, so t = a·r ≤ 2a L_1 s √d (m + 1) ≤ 3a L_1 s √d·m, and since ν! ≥ (ν/e)^ν, t^ν/ν! ≤ (et/ν)^ν ≤ (6e·a·L_1 s √d)^(m/2) once the bracket is below 1. The sum over m is then at most C e^(−κr), once a is small enough that 6e·a·L_1 s √d·$(\Lambda b)^2$ < e^{-2}. The constants are chosen once, in this order: s and then b so that 2Λq^(1/2) < e^{-1}, where Λ depends on d alone; then a. ∎

Lemma 5.5 (The front along each ray). There is a constant c with 0 < c < ∞ such that, for every q in space, ρ(0, Lq)/L → |q|/c in probability and in mean as L → ∞ through the real numbers. Since the law of ρ(0, x) depends only on |x|, E|ρ(0, x) − |x|/c| = o(|x|) as |x| → ∞.

Proof. Fix q ≠ 0. For whole numbers m < n put X_(m,n) = ρ(mq, nq). The triangle inequality gives X_(0,n) ≤ X_(0,m) + X_(m,n). The model's law is unchanged by the shift by q, so the array is stationary. A Poisson process with independent marks is mixing under any nonzero shift, so the shift is ergodic. X ≥ 0 and E X_(0,1) < ∞ by Lemma 5.3. The subadditive ergodic theorem (Kingman 1968, in Liggett's form, 1985) gives ρ(0, nq)/n → γ(q), a constant, almost surely and in mean. For real L, |ρ(0, Lq) − ρ(0, ⌊L⌋q)| ≤ ρ(⌊L⌋q, Lq), which has the law of ρ(0, (L − ⌊L⌋)q) and is bounded in mean by Lemma 5.3; divided by L it tends to 0. With Lemma 5.3's second moment, ρ(0, Lq)/L → γ(q) in probability and in mean. Scaling gives γ(tq) = tγ(q) for t ≥ 0, and the model's law is unchanged by rotations, so γ(q) = |q|/c for a constant c. Since k(nq) lies within a distance of nq that is bounded in probability, Lemma 5.4 gives γ(q) ≥ a|q|, so c ≤ 1/a; and γ(q) ≤ E ρ(0, q) < ∞, so c > 0. ∎

Remark. Ziesche (2016), extending Boivin (1990), proves a stronger shape theorem for random pseudometrics of this kind, uniform over directions. The proof here needs only Lemma 5.5.

Lemma 5.6 (One added keeper changes little). Add a keeper at a point p, with its own clock and wait w_p, to get Π ∪ {p}, and write ρ′ for ρ computed in Π ∪ {p}. Its cavity is the set of keepers of the Delaunay simplices whose circumscribed balls contain p. Let H_p be the total of the waits of p and of its cavity. Then, for every point x:

- ρ′(p, x) ≤ w_p + w_(k(p)) + ρ(p, x), where k(p) is p's nearest keeper in Π;
- ρ(p, x) ≤ ρ′(p, x) + H_p, when x's cell in Π ∪ {p} is not p's.

H_p has an exponential tail, the same for every p.

Proof. Inserting p removes exactly the Delaunay simplices whose circumscribed balls contain p, and joins p to every keeper of the cavity (Bowyer 1981; Watson 1981); k(p), the nearest keeper to p, is one of them. Every edge of the new graph not at p is an edge of the old one, since its empty ball is empty for Π too; and every old edge with an endpoint outside the cavity survives, since every simplex at that endpoint does. A point's cell can change keeper only to p. For the first bound, take an optimal walk in Π from k(p) to k(x); its waits total ρ(p, x), or w_(k(p)) when k(x) = k(p) and the walk is k(p) alone, where ρ(p, x) = 0. Take its last cavity keeper. Then p, that keeper and the rest of the walk form a walk in Π ∪ {p}, since every later edge has an endpoint outside the cavity, and it costs at most w_p + w_(k(p)) more than ρ(p, x). (If x's cell is now p's, the left side is 0.) For the second, take a walk in Π ∪ {p} from p, which visits p only at its start. Replace p by a walk in Π from k(p) to the walk's second keeper along the edges of the removed simplices, which connect the cavity since their union is star-shaped from p, at a cost of at most the cavity's waits. Every cavity keeper lies on the boundary of an open ball that contains p and holds no keeper of Π, so within twice the radius of the largest such ball. If one has radius at least t, it contains a ball of radius t/2, empty of Π and centred within t/2 of p: centre it on the segment from p to the large ball's centre, t/2 from p, or at that centre if it lies within t/2 of p. That ball contains a ball of radius t/4 centred at a point of a fixed net of at most C points within t/2 of p, spaced so that every point there is within t/4 of the net, with C depending on d alone. So the largest such ball has radius at least t with probability at most C exp(−ω_d(t/4)^d) ≤ C exp(−κt^d). So the cavity's size has an exponential tail: with t = c m^(1/d) and c small, P(more than m cavity keepers) ≤ P(that radius exceeds t) + P(Π(B(p, 2t)) > m) ≤ C e^(−κm). H_p is a sum of that many independent Exp(1) waits, independent of Π, so it has an exponential tail too. Since k(p) is a cavity keeper, w_p + w_(k(p)) ≤ H_p, and both bounds give |ρ′(p, x) − ρ(p, x)| ≤ H_p wherever the second applies. ∎

Lemma 5.7 (Few misjudged episodes). Fix ε > 0 and a bounded region A of space-time. Say an episode w = (k, t) is misjudged from an episode z = (j, s) when z ≺ w although t − s < |p_k − p_j|/c − εL, or z ⊀ w although t − s > |p_k − p_j|/c + εL. Let x be the first episode after time 0 of k(0). Then, as L → ∞:

1. the expected number of episodes in L·A misjudged from x is o(L^(d+1));
2. the same holds with x replaced by y, the first episode after time L·t_0 of k(L·u), for misjudgements of y's past;
3. the expected number of pairs (z, w) of episodes in L·A with w misjudged from z is o(L^(2(d+1))).

Proof. (1) By Lemmas 5.1 and 5.2, w = (k, t) follows x exactly when t ≥ τ_x(k) = A_k + W_k. Write Y_k = |τ_x(k) − s − |p_k − p_0|/c|, where p_0 is the position of k(0). The misjudged episodes of k are ticks of k lying between s + |p_k − p_0|/c ± εL and τ_x(k), and there are none unless Y_k > εL. k's ticks before A_k are a Poisson process of rate 1 independent of A_k, and after τ_x(k) they start afresh, so the expected number misjudged at k is at most E[(1 + Y_k) · 1{Y_k > εL}]. In detail: if τ_x(k) comes more than εL late, the misjudged ticks lie before A_k, since k has no tick in (A_k, τ_x(k)), and given A_k they are Poisson with mean at most Y_k; if it comes more than εL early, they are τ_x(k) itself and the ticks after it, a fresh Poisson process, at most 1 + Y_k in mean. By Lemma 5.2, τ_x(k) − s differs in law from ρ(0, p_k) by the wait at k(0) and by |p_0|/c, both with bounded moments. Sum over keepers in L·A by the Mecke equation and Slivnyak's theorem: the expected sum equals the integral over positions p of the same expectation in Π ∪ {p}. By Lemma 5.6, applied from p and using that ρ is

symmetric, adding p changes the passage time between p and p_0 by at most H_p, since p_0 is a keeper and so keeps its cell. Adding p also leaves k(0) and so x unchanged, unless p is nearer to 0 than p_0 is, which has probability at most exp(−ω_d|p|^d). Split the integral. Where |p| < δL the integrand is at most C(1 + δL) by Lemma 5.3, contributing at most Cδ^(d+1)L^(d+1). Where |p| ≥ δL the integrand is o(L), uniformly, by Lemma 5.5 and Markov's inequality, since δL ≤ |p| ≤ C_A L on the bounded region, contributing o(L^(d+1)); the event that p takes 0's cell adds at most C L exp(−ω_d(δL)^d) there. Let L → ∞, then δ → 0.

(2) By the consequence of Lemma 5.1, reversing time about $L \cdot t_0$ turns y's past into a future with the same law, started from the keeper k(L·u). The model's law is unchanged by the shift taking L·u to 0, so (1) applies.

(3) By the Mecke equation, applied to the keepers with their clocks and then to each clock in time, the expected number of such pairs is the integral, over z = (q, s′) in L·A, of the expected number of episodes in L·A misjudged from an episode at time s′ of a keeper added at q. By Slivnyak's theorem, shifting space by −q and time by −s′ leaves the law of the model with that added keeper unchanged, apart from moving the keeper to 0 and the episode to time 0. The shifted region L·A − z lies in L·A′ for the fixed bounded region A′ = A − A. So the integrand is at most the expected number of episodes in L·A′ misjudged from an episode at time 0 of a keeper added at 0, the same for every z. That number is o(L^(d+1)), by the proof of (1) with p_0 = 0 and A′ for A: Lemma 5.2 holds from an added keeper, and Lemma 5.6, applied once for each of the two added keepers, changes the passage time between 0 and p by at most H_0 + H_p, except when 0 is p's nearest keeper. That has probability exp(−ω_d|p|^d), where ω_d is the volume of the unit ball, which is negligible where |p| ≥ δL; where |p| < δL, Lemma 5.3 bounds the integrand as in (1). Lemma 5.6 applies with Π ∪ {0} in place of Π, since a ball empty of Π ∪ {0} is empty of Π, so the cavity's tail is unchanged. Integrating over L·A gives o(L^(2(d+1))). ∎

Lemma 5.8 (No slow pushes). There are constants C and κ > 0 such that, for all points u and v of space, P(ρ(u, v) > C(1 + |u − v|)) ≤ C exp(−κ|u − v|^(1/d)).

Proof. In d = 1 the walk from k(u) to k(v) is the run of keepers between them, and the bound is immediate. Let d ≥ 2. Cut space into a grid of cubes of side s, and let $L_0 = \lceil \sqrt{d} \rceil + 1$. Call a cube fine if it holds a keeper and every cube within L_0 of it, in every coordinate and counted in cubes, holds at most b keepers.

1. *Few cells per fine cube.* Every point of a fine cube Q is within s√d of a keeper of Q, so the cell containing it has its keeper in a cube within L_0 of Q. So at most B = $(2L_0 + 1)^d$ b cells meet Q.
2. *Short walks between fine cubes.* For fine cubes sharing a face, the cells meeting the segment between their centres form a walk in G of at most 2B cells.
3. *Bad cubes are rare and short-range.* A cube fails to be fine with probability at most p = exp(−s^d) + $(2L_0 + 1)^d$ P(Poisson(s^d) > b), which is small once s and then b are large. Cubes more than $2L_0$ apart fail independently.

Hulls. Group the cubes that are not fine into clusters of cubes touching at least at a corner; they are finite almost surely, by the count in Count 1. The regions a cluster leaves are the components of the other cubes, joined through shared faces. A cluster's hull is the cluster with every bounded region it leaves. Two facts follow.

- (a) A cube of a bounded region shares faces only with cubes of that region or of the cluster. So a chain of cubes sharing faces that starts outside a hull first meets it at a cube of the cluster; and a cube outside a hull that shares a face with it touches the cluster, and so is fine, or it would belong to the cluster.
- (b) Two clusters do not touch, so each lies in a single region left by the other: two touching cubes, neither in a cluster A′ and the first not touching A′, lie in a block of 2^d cubes that all touch the first, so none is in A′, and the block joins them through shared faces. If A′ lies in a bounded region left by A, the unbounded region left by A misses A′, so it lies in the unbounded region left by A′,

and A′'s hull lies in A's. If each lies in the unbounded region left by the other, their hulls are disjoint: a common cube would lie in a bounded region W left by A and W′ left by A′; W misses A′, so W lies in W′, and likewise W′ in W; then the cubes bordering W would belong to both clusters. So two hulls are disjoint, or one lies in the other.

A walk. Let $\ell = |u - v|$. The segment from u to v passes through a chain of at most $M = d(\ell/s + 2)$ cubes, each sharing a face with the next. If some hull holds both u's cube and v's cube, the walk is the cells the segment meets, counted in Count 2. Otherwise let U be the largest hull holding u's cube, the hull of a cluster A_u, and V likewise at v, of A_v; by (b), U and V are disjoint. If there is no U, u's cube is fine; let e_u be u's cube, and otherwise the first cube after the chain last leaves U, which is fine by (a). Let e_v be v's cube, or the last cube before the chain next meets V. The walk starts with the cells met by the segment from u to where it last leaves U, then by the straight line on to the centre of e_u. Between e_u and e_v, take the clusters whose hulls the chain meets and that lie in no other such hull. None of these hulls meets U or V, or holds e_u or e_v. A hull meeting U would, by (b), contain U, since U is the largest hull holding u's cube, and the chain between e_u and e_v does not meet U. A hull of a cluster A holding e_u would hold it in a bounded region W left by A, since e_u is fine; A_u touches e_u and not A, so by (b) A_u lies in W too, and the hull would contain U. Likewise at v. So, by (a), the chain first meets each of them at a cube of its cluster. Splice as follows, on the current chain: while it holds a cube that is not fine, take the first such cube, the largest hull containing it, and that hull's cluster A, and cut the current chain from the last cube before it first meets A's hull to the first cube after it last leaves it. Both ends touch A from its unbounded region, by (a). The cubes that touch A at least at a corner and lie in its unbounded region are connected through shared faces (Timár 2013, Theorem 4, in its form for the lattice of cubes: for a finite set of cubes connected through touching at least at a corner, the cubes outside it that touch it at least at a corner and can be reached from far away without entering it are connected through shared faces), and they are fine, because A is a whole cluster. So the cut can be closed through them, adding at most $3^d|A|$ cubes, all fine. Each splice removes every cube of one hull from the chain and adds only fine cubes, so the number of cubes on the chain that are not fine falls each time, and the splicing ends in a chain of fine cubes, each sharing a face with the next, from e_u to e_v. Every hull spliced is one of the hulls above, since its first cube on the current chain is a cube of the original chain. Item 2 joins them into a walk in G, and the end at v is as at u. Let S be the total size of these clusters.

Count 1: small clusters. The chain meets each of these clusters, so the chain's cubes with these clusters form a set of cubes connected at least through corners, containing u's cube, with at least S cubes not fine. There are at most C_1^j such sets of j cubes. Among j′ cubes that are not fine, at least j'/K_0 are pairwise more than $2L_0$ apart, where $K_0 = (4L_0 + 1)^d$, so all j′ fail with probability at most p to the power j'/K_0. A set of $j > M$ cubes holds at least $j - M$ not on the chain, and those are cubes of the clusters, so not fine. Summing C_1^j p to the power $(j - M)/K_0$ over $j \geq KM$ gives $P(S > KM) \leq C\, e^{(-\kappa M)}$, with $K = 2$, once p is small.

Count 2: the ends. Let S_u be the size of A_u, or 0 if there is no U. A_u, with the cubes from u's cube to it along one coordinate direction, is a set of at most $2S_u$ cubes connected at least through corners, containing u's cube, at least half not fine; so $P(S_u > t) \leq C\, e^{(-\kappa t)}$ by the same count. U lies in the smallest box of cubes holding A_u, so every point of the segment from u to where it last leaves U is within $s\sqrt{d}\cdot S_u$ of u, and e_u holds a keeper within $s\sqrt{d}(S_u + 1)$ of u. So the cells met before the centre of e_u have their keepers within $3s\sqrt{d}(S_u + 1)$ of u. Their number N_u is at most the number of keepers within that distance, so $P(N_u > t) \leq C \exp(-\kappa\, t^{(1/d)})$. The same holds at v, and when one hull holds both ends.

Count 3: the waits. On the event $S \leq KM$, $N_u \leq \ell$ and $N_v \leq \ell$, the walk from k(u) to k(v) has at most $N_u + N_v + 2B(1 + 3^d K)M \leq C'(1 + \ell)$ cells, with C′ fixed. The walk is chosen from Π alone, so its waits are independent Exp(1), independent of where the keepers are, and their total exceeds $2C'(1 + \ell)$ with probability at most $e^{(-\kappa\ell)}$. Adding the failure probabilities gives the bound. ∎

Lemma 5.9 (Rare slow pushes, from every episode). Fix δ > 0 and a bounded region A of space-time. Let F_L be the number of episodes z = (k, t) in L·A for which some keeper k′ within distance δL of p_k has τ_z(k′) − t > C(2 + 4δL), with C from Lemma 5.8. Then E[F_L] → 0. The same holds for the past of every episode.

Proof. By the Mecke equation, applied to the keepers and then to each clock, and Slivnyak's theorem, E[F_L] is at most the integral, over z = (p, t) in L·A and positions p′ within δL of p, of the probability that a keeper added at p′ is reached late from an episode at time t of a keeper added at p. By Lemma 5.2 it is enough to bound ρ(p, p′) with both keepers added. By the first bound of Lemma 5.6, applied once for each added keeper and using that ρ is symmetric, that is at most ρ_Π(p, p′) plus four waits: those of p and p′ and of the nearest keeper to each when it is added. In d = 1 no detour is needed: ρ_Π(p, p′) is the total wait of at most 2 + Poisson(δL) keepers, the end keepers and those between, which exceeds C(1 + 3δL) with probability at most C e^(−κδL) once C is large. For d ≥ 2, pick a point w at distance δL from p, perpendicular to the segment from p to p′. Then ρ_Π(p, p′) ≤ ρ_Π(p, w) + ρ_Π(w, p′). Both legs have length between δL and 2δL, so by Lemma 5.8 they total more than C(2 + 3δL) with probability at most 2C exp(−κ(δL)^(1/d)). The four waits total more than CδL with probability at most 4e^(−CδL/4). So E[F_L] ≤ C L^(d+1)(δL)^d exp(−κ(δL)^(1/d)), with new constants, which tends to 0, since a stretched exponential beats any power. The past follows by the time reversal of Lemma 5.1. ∎

Theorem 5.10 (The Minkowski limit). Let x be the first episode after time 0 of k(0), and y_L the first episode after time L·t_0 of k(L·u), where |u| < c·t_0. Rescale space and time by 1/L, and let D be the causal diamond from (0, 0) to (u, t_0) in Minkowski spacetime with light speed c. Then, in probability as L → ∞:

1. the interval I(x, y_L) holds L^(d+1)(|D| + o(1)) episodes, where |D| is the volume of D, which depends only on the proper time between its tips;
2. the share of related pairs in I(x, y_L) tends to f(d+1), the ordering fraction of (d+1)-dimensional Minkowski spacetime, so the Myrheim–Meyer dimension of I(x, y_L) tends to d + 1;
3. both limits depend on u only through the proper time between the tips, the share not at all, so at equal proper time an interval moving through the keepers reads as one at rest, and motion through the keepers cannot be detected by these readings;
4. the midpoint dimension of I(x, y_L) also tends to d + 1.

$$\mathbf{f(n) = \Gamma(n + 1)\,\Gamma(n/2) / (2\,\Gamma(3n/2))}$$

Proof.

1. *Density.* Episodes form a point process in space and time, with one episode per unit volume, whose law is unchanged by every shift of space and time. It is ergodic under these shifts, since shifts of space alone are mixing; shifts of time alone are not, as they leave the keepers in place. By the ergodic theorem for stationary point processes (Daley and Vere-Jones 2008), applied along the convex averaging sequence L·B, the number of episodes in L·B is L^(d+1)(|B| + o(1)) almost surely, for every box B with a corner at the origin. Adding and subtracting such boxes gives the same for every box, and approximating from inside and outside by finite unions of boxes gives it for every bounded B whose boundary has volume zero. The number of pairs in L·B × L·B′ is the product of the counts.
2. *The interval.* By Lemma 5.4, every episode of I(x, y_L) lies in a fixed bounded rescaled region A, with probability tending to 1; Lemma 5.7 is applied with this A. With probability tending to 1, x and y_L lie within εL/2 of the tips of L·D, since their offsets from the tips have bounded moments. Then an episode z of L·D more than εL inside D's boundary is more than εL/2 inside the cones from x and y_L, so follows x and precedes y_L under the cone rule, and one more than εL outside fails one of them; any other outcome is a misjudgement counted in Lemma 5.7, parts (1) or (2), with ε/2 for ε. Enlarge A if needed so that it contains D, and let G_L be the event that every

episode of I(x, y_L) lies in L·A and x and y_L lie within εL/2 of the tips. On G_L, I(x, y_L) and the episodes of L·D differ only by episodes of L·A misjudged from x or y_L and by episodes within εL of the boundary of L·D. The misjudged episodes of L·A number o(L^(d+1)) in expectation by Lemma 5.7, whatever happens off G_L, so o(L^(d+1)) in probability by Markov's inequality; the boundary strip holds a share O(ε) by step 1; and P(G_L) → 1. Letting ε → 0 gives item 1. The volume of a diamond depends only on the proper time between its tips.

3. *Pairs*. A pair of episodes in I(x, y_L) is related exactly as the cone rule with speed c says, except for pairs counted in Lemma 5.7, part (3), pairs within εL of the cone, a share O(ε) by step 1, and pairs involving the o(L^(d+1)) episodes where I(x, y_L) and L·D differ. The pairs of D × D within ε of the cone, and the pairs related under the cone rule, form regions whose boundaries have volume zero; covering them by finitely many products of boxes, step 1 counts the episode pairs in them. So the related share tends to that of pairs of uniform random points of D under Minkowski causality with speed c. Measuring time in units of 1/c makes this standard Minkowski spacetime, where the share is f(d+1) in every diamond (Myrheim 1978; Meyer 1988). The Myrheim–Meyer dimension is a continuous function of the share, giving item 2. Every diamond is a boost and rescaling of every other, so the share is the same in every diamond, and by item 1 the count depends on the tips only through the proper time between them, giving item 3. ▮

Proof of item 4. Write n = d + 1. The midpoint dimension of an interval with m episodes is $\log_2(m / h)$, where h is the largest, over episodes z of the interval, of the smaller of |I(x, z)| and |I(z, y_L)|. By item 1, m = L^n(|D| + o(1)). The proof of item 1 applies to any two endpoints chosen as x and y_L are, by stationarity. When the second endpoint lies on or outside the future cone of the first, the closed diamond between them is a null segment or empty, and has volume zero. Every episode of the interval is then misjudged, counted in Lemma 5.7, or lies within εL of both cones; the set of rescaled points within ε of both cones decreases, as ε → 0, to that closed diamond, so its volume tends to zero, and by step 1 the interval holds o(L^n) episodes. In Minkowski spacetime the proper time from 0 to a point of D plus the proper time from it to (u, t_0) is at most the proper time across D, with equality only on the line between the tips. Diamond volume grows as proper time to the power n, so the smaller of the two sub-diamonds has volume at most |D|/2^n, with equality at the centre.

Lower bound. Let z *be the first episode after time L·t_0/2 of k(L·u/2), at the centre. By item 1 for the pairs (x, z)* and *(z, y_L), both sub-intervals hold L^n(|D|/2^n + o(1)) episodes, and z* lies in I(x, y_L) with probability tending to 1, by Lemma 5.5 and its time reversal, since the centre lies strictly inside D. So h ≥ L^n(|D|/2^n − o(1)).

Upper bound. Fix η > 0 and δ < η/(4C), with C from Lemma 5.8. By Lemma 5.4, every episode of I(x, y_L) lies in a fixed bounded rescaled region, with probability tending to 1. Lay a grid over that region, of spacing δ/√d in each direction of space and δ in time, so that every point of space is within δ/2 of a grid point. For each grid point g = (q_g, t_g), let e_g be the first episode after time L·t_g of k(L·q_g). There are finitely many, so with probability tending to 1, all at once: every |I(x, e_g)| and |I(e_g, y_L)| is within o(L^n) of its diamond's volume, by item 1; every k(L·q_g) lies within δL/2 of L·q_g; and every e_g comes within δL of time L·t_g. Take any episode z of I(x, y_L) at rescaled position (q, t). Choose grid points g^+ and g^- within δ/2 of q in space, so that their keepers lie within δL of z's keeper, with t_(g^+) between t + η and t + η + δ, and t_(g^-) between t − η − 2δ and t − η − δ. By Lemma 5.9, outside an event of probability tending to 0, every such z reaches k(L·q_(g^+)) within time C(2 + 4δL), which is less than ηL for large L, so z precedes e_(g^+), which comes after time L(t + η). Likewise, by the past form of Lemma 5.9, z is preceded by an episode of k(L·q_(g^-)) after time L(t − η), and e_(g^-), which comes before time L(t − η), is carried to it, so e_(g^-) precedes z. Every episode before z is before e_(g^+), so |I(x, z)| ≤ |I(x, e_(g^+))|, and likewise |I(z, y_L)| ≤ |I(e_(g^-), y_L)|. Diamond volume is a constant times the squared proper time to the power n/2 ≥ 1, taken as 0 outside the cone, which is a Lipschitz function of the endpoints on a bounded region, so the diamonds of those grid episodes differ from the diamonds of z by O(η) in volume. So the smaller of |I(x, z)| and |I(z, y_L)| is at most L^n(|D|/2^n + O(η) + o(1)) for every z at once. Letting η → 0 gives h = L^n(|D|/2^n + o(1)), and m/h → 2^n, so the midpoint dimension tends to n. ▮

Scope. Theorem 5.10 is about large scales only: keepers at rest still carry a frame at the scale of one cell. Its limit is flat spacetime; curvature, gravity and matter are not addressed here. And d is an input: the theorem gives one dimension more than space has, whatever space has. It covers worlds of keepers scattered at random, and its proof uses the scatter in four ways: the world is the same everywhere and in every direction, and its translations are mixing (Lemma 5.5; Theorem 5.10, step 1); cells are never huge and never crowded, except with rapidly vanishing probability (Lemmas 5.3, 5.4 and 5.8); far regions are independent (Lemmas 5.4 and 5.8); and a keeper can be added without disturbing the rest (Lemmas 5.6, 5.7 and 5.9). Theorem 5.11 shows that the fourth is not needed, and that the first three can be weakened to three conditions on a world of cells.

Flat worlds of held cells with W1 to W3. A grown foam is not a Poisson scatter, and a keeper cannot be added to it without changing what grows next. Theorem 5.11 proves the conclusions of Theorem 5.10 for every world of cells with three properties. A fixed grid of points and the triangle inequality do the work of the added keeper of Lemma 5.6, and the only Poisson processes the proof still uses are the keepers' own clocks, which are Poisson by construction. The full proof is in Appendix A.

Definition 6 (Cell world). A *cell world* in d-dimensional space is a random collection of compact convex polytopes, the *cells*, each with nonempty interior, whose interiors are disjoint, which cover space, and of which only finitely many meet any bounded set. Each cell C_k is a keeper k, with a position p_k in the interior of C_k; the world is the cells together with their positions, so a shift or rotation of space moves both. Two keepers are *neighbours* when their cells meet in a piece of dimension d − 1; this is the graph G. Cells need not meet face to face. Clocks, pushes, the order $\lessdot$, the waits w_k and the passage ρ are as in the Setting, with k(u) the keeper whose cell holds u, and lengths are scaled so that there is one cell per unit volume on average. Lemmas 5.1 and 5.2 use nothing about positions and hold as stated.

Notation. B(u, R) is the closed ball of radius R about u, and N(u, R) the number of cells meeting it. For $s > 0$ the grid cubes of side s are the translates $s(z + [0, 1)^d)$, $z \in \mathbb{Z}^d$, of a fixed placement of the grid. A world of cells may have three properties:

- **W1 (settled, the same in every direction, forgetting under shifts).** The law of the world is unchanged by every translation and every rotation, and translations are mixing: $P(E \cap \theta_x F) \to P(E)P(F)$ as $|x| \to \infty$, for all events E and F, where θ_x shifts the world by x.
- **W2 (no runaway clumps).** N(0, 1), the number of cells meeting a unit ball, has finite moments of every order.
- **W3 (crowded or oversized patches forget each other).** Call a closed grid cube of side s *bad* when some cell meeting it has diameter above s, or more than b cells meet it. There is a function p(s, b), as small as we like for suitable s and b, such that for every placement of the grid and every finite set J of grid cubes, $P(\text{every cube of J is bad}) \le p(s, b)^{|J|}$.

What each property is for. W1 gives pushes a single speed, the same in every direction, and episodes a single density. W2 bounds how many cells a short stretch of space can hold. W3 is the precise form of *far regions forget each other*: a long run of bad cubes is as unlikely as a long run of independent rare events. It gives pushes a top speed, cells an exponential tail on their size, and pushes no long delays.

Theorem 5.11 (The Minkowski limit for cell worlds). *Let a cell world in $d \ge 2$ dimensions have W1, W2 and W3. Then items 1 to 4 of Theorem 5.10 hold, with c the constant of Lemma A.3: at large scale the order of held boundaries is (d + 1)-dimensional Minkowski spacetime, in every frame, on the Myrheim–Meyer and the midpoint rulers.*

Corollary 5.12 (The Poisson world is a case). *The Poisson–Voronoi world has W1, W2 and W3 (Proposition A.9), so for $d \ge 2$ Theorem 5.10 is the case of Theorem 5.11 with keepers at the Poisson points.* For d = 1 the proof of Theorem 5.10 stands.

The proof in outline.

- Lemma A.0: a generic segment passes from cell to cell only through the interiors of facets, so the cells it meets form a walk.
- Lemma A.1: cell counts and passage times have moments of every order (W1, W2).
- Lemma A.2: pushes have a top speed, and cells an exponential tail on their size (W3).
- Lemma A.3: along each ray the front settles to one speed c, the same in every direction (W1).
- Lemma A.4: a fixed grid of points stands in for the keepers, by the triangle inequality.
- Lemmas A.5 to A.7: few misjudged episodes and rare slow pushes, as Lemmas 5.7 to 5.9, with the Mecke equation applied to the clocks alone.
- A.8: the proof of Theorem 5.10 then applies as written.

Scope of Theorem 5.11. W1 asks for exact isotropy. If the large-scale shape of pushes were an ellipsoid, a linear change of coordinates would make it round and the limit would still be Minkowski spacetime; if it were not an ellipsoid, the limit would be a spacetime whose light cone has that shape. The theorem is stated for $d \geq 2$, since a walk in one dimension cannot go around a crowded stretch; the proof of Theorem 5.10 covers the Poisson line. Theorem 5.11 reduces the question to one statement about the world the cells make: that it has W1, W2 and W3. Section 6 proves it for the steady state of the tension rule written as a balance of births and deaths.

---

6. The Tension World: Held Cells under Living Tension

What this section proves, in plain words. A growing world's history is hard to prove anything about: as it grows, lumps wider than its cells can read may freeze in. This section sets the history aside and proves what the settled world is. It writes the tension rule, cells divide when stretched and leave when crowded, as a balance between births and deaths, whose steady state is a Gibbs world: the object statistical physics uses for gases and liquids. For gentle tension that world exists, is unique, and has W1, W2 and W3, so Theorem 5.11 applies: in that steady state, held cells under gentle tension make Minkowski spacetime. W3 asks only that far patches forget each other, not that the world be more even than random, so gentle tension is enough.

Definition 7 (Foam and reach). Keepers are points of space with Voronoi cells. Two keepers are neighbours when their cells share a face that some empty ball of radius at most R_cut touches at both keepers: the Voronoi foam with a boundary cut wherever no such ball reaches it. In every world in which every point of every cell lies within R_cut of its keeper, this is the full foam: a point z of the face two neighbours share is then the centre of a ball of radius at most R_cut that touches both and holds no keeper. The *reach* $B_k(r)$ is the number of keepers within r crossings of k in this foam. The cut foam is used only to read reach; pushes run across every shared face of the full foam, as in Definition 6, and Theorem 6.1(iii) is about the full foam.

Definition 8 (Extent). No two keepers are closer than $\delta > 0$. A held boundary has extent (FR15, H2), so a cell cannot shrink to a point.

Definition 9 (Tension energy). For a finite set of depths $\mathscr{D}$ and fixed rest values $\beta_r > 0$, the tension of keeper k at depth r is $\theta_k(r) = B_k(r)/\beta_r - 1$, and its cost is

$$\boldsymbol{\Phi_k = \Sigma_(r \in \mathscr{D})\, \psi(\theta_k(r)),\quad \psi(u) = \min(u^2, M)}$$

a bounded cost of being away from rest, zero at rest. A keeper whose reach at a depth is above its rest there is stretched at that depth, and one below it is crowded. The tension energy of a finite set of keepers is $H = \lambda \cdot \Sigma_k\, \Phi_k$, where $\lambda \geq 0$ is the strength of the tension.

Definition 10 (The tension world). The tension world of strength λ is a Gibbs world for H relative to a Poisson process of intensity z: a random set of keepers whose law, inside any bounded region Λ and given everything outside it, has density proportional to exp(−H_Λ) against a Poisson process of intensity z, restricted to configurations with no two keepers closer than δ. Here H_Λ sums the costs λΦ_k of every keeper, inside Λ or outside it, whose cost depends on the keepers in Λ; by finite range these lie within R_int of Λ (B.1), and adding a keeper at x changes H_Λ by ΔH(x, ω), the full change of B.1. For λ = 0 it is the hard-sphere gas of diameter δ, which for δ small is close to the Poisson world of Section 5.

Why this is the tension rule. Run the world as births and deaths: a new keeper is born at a point x at rate z·exp(−(H(ω ∪ {x}) − H(ω))/2), and a keeper k dies at rate exp(−(H(ω \ {k}) − H(ω))/2). The energy is zero when every keeper is at rest and grows with every keeper's stretch or crowding, so the births and deaths favour every change that brings keepers toward rest, counting every keeper whose reach the change moves. In that sense the balance acts on stretched and crowded keepers alike, in the balance of the whole rather than keeper by keeper, and that is why it is the tension rule. How a keeper's stretch or crowding relates to the density of keepers in space is not a result here: reach does not change when a world is rescaled, and a single birth can move it either way, as at depth 1, where a keeper added among three raises their reach. No theorem here uses that relation. Rates of this balanced form make the Gibbs world the steady state of the births and deaths in every bounded region (Preston 1975).

Theorem 6.1 (Held cells under living tension). *Fix $d \geq 2$, depths $\mathscr{D}$, rest values β_r, the cap M and the cut R_{cut}. For every δ with $0 < \delta < 2R_{cut}$ and z with $z \cdot v_d \delta^d \leq 1/4$, where v_d is the volume of the unit ball, there is $\lambda_0 > 0$ such that for $0 \leq \lambda < \lambda_0$: (i) the tension world exists and is unique; (ii) it is unchanged by translations and rotations and is mixing; (iii) its Voronoi foam, with lengths scaled so that there is one keeper per unit volume, has W1, W2 and W3. Hence, by Theorem 5.11, the order of held boundaries in the tension world converges at large scale to $(d+1)$-dimensional Minkowski spacetime, in every frame.*

The proof in outline. The full proof is in Appendix B.

- **Finite range and local stability** (B.1). The energy a keeper adds or removes is at most $C_\lambda = \lambda \cdot |\mathscr{D}| \cdot M \cdot (K_\delta + 1)$, which tends to 0 with λ, and is decided by the keepers within $2R_{int}$ of it, where $R_{int} = 2R_{cut}(r_{max} + 1)$ and $K_\delta = ((2R_{int} + \delta)/\delta)^d$.
- **Existence** (B.2), by the theory of Gibbs point processes with finite-range interactions (Ruelle 1970; Preston 1976; Dereudre 2019), including Voronoi and Delaunay interactions (Dereudre 2008; Dereudre, Drouilhet and Georgii 2012).
- **Uniqueness** (B.3). Built backwards from an independent process of candidate births, each candidate's fate depends on a clan of earlier candidates (Fernández, Ferrari and Garcia 2001, 2002). A birth looks across its full range only when its fate is in doubt, which happens with probability at most $1 - \exp(-2C_\lambda)$; for $\lambda < \lambda_0$ the clans are subcritical branching processes, finite, with exponential tails. The tension world is then unique, and is built from independent randomness by a rule that commutes with shifts and rotations.
- **W1, W2 and W3** (B.4). Mixing from the clan tails; moments from domination by a Poisson process and a bound on empty balls; and W3 by the route of Proposition A.9, with each cube's clan exploration stopped at its first look outside a neighbourhood of fixed size, so that cubes far enough apart are exactly independent.

Scope. Theorem 6.1 holds for every fixed choice of the rest values β_r; choosing them as the world's own typical reach is a further condition the theorem does not need. The extent δ is FR15's H2, that a held boundary has size, and it is smaller than the cut, $\delta < 2R_{cut}$, so that cells can meet in the cut foam; and the cut R_{cut} never acts in a world in which every point of every cell lies within R_{cut} of its keeper (Definition 7), while a point farther than R_{cut} from its keeper is the centre of an empty ball of radius above R_{cut}, which the bound on empty balls makes exponentially rare; the energy itself reads

only the foam. The births and deaths are one route to the steady state, and the theorem is about the state, not the route: pressing, settling and division are other dynamics, and that they reach the same steady state is not claimed. The interaction reaches across many cells, and λ_0 shrinks as the depth grows. Theorem 6.1 is stated for gentle tension, and strong tension lies outside it.

What it means. For gentle tension, the steady state of held cells under living tension, each reading only its own reach, is a world that Theorem 5.11 covers, and so makes Minkowski spacetime in every frame. Whether worlds that grow by pressing and division reach that state is open. The tension is part of the theorem: it is the energy whose balance defines the world.

7. Two Tiers and Falsification

The governance tier and the ontological tier. As in FR14 Section 8 and FR15 Section 6, this record speaks at two depths, and they are kept apart.

- **Governance independence.** No theorem, definition or gate condition in FR01 to FR13 depends on this record. That perception does not close action (FR01), that computation cannot generate authority (FR02), that the governance gate fails closed without declared authority δ from a natural person at the boundary (FR06, FR07, WP04), that a governance claim must survive subject resistance under the Gildenston Functional and the ILMM (FR10, FR11, FR13), and that AI hardware cannot ground its own floor or scope (ER1 to ER4, WP03) all hold whatever spacetime turns out to be.
- **Where the two tiers meet.** Within the full corpus, a natural person who declares authority δ at a governance boundary (FR06) is an individuated subject holding her own boundary in time and space (FR14, FR15), and the time and space in which she stands are produced by the holding, carrying and mutual testing of boundaries set out here.

Falsification conditions. Within the ontological tier of FR12 Section 7, FR14 and FR15, this record is falsified by any of the following:

1. **Order-equivalence failure.** A benchmark system in FR14 Section 7 whose individuation verdict under the clock-free Definitions 4 and 5 differs from its verdict under FR14's definitions.
2. **Bad-order survival.** A two-layer or three-layer order that admits a keeper satisfying Definition 4 with relaxation depth $d_k \geq 2$.
3. **Limit failure.** A cell world with W1, W2 and W3, or a tension world in the regime of Theorem 6.1, whose order of held boundaries does not converge at large scale to (d + 1)-dimensional Minkowski spacetime on the rulers of Theorem 5.10, or whose readings there tell an interval moving through the keepers from one at rest at equal proper time.

Core Finding

Under the ontology this record works in, consciousness is fundamental and has no boundaries. Spacetime is not a container handed to it from outside; it is what held boundaries make from within. Time is an edge carried forward along a living line (Classical). Space is edges held at once and tested against one another (Quantum). Reality is the ground inside which both are made (Consciousness). At large scale, what they make is Minkowski spacetime, and it stays so in the steady state of living tension, the balance of cells coming and going. Carried, not given.

Appendix A. Proof of Theorem 5.11

Notation. As in Section 5, B(u, R) is the closed ball of radius R about u and N(u, R) the number of cells meeting it. Q(y) is the grid cube holding y, and two cubes are *-neighbours when their indices differ by at most 1 in every coordinate. C, κ and C_q denote constants that depend on d, on the world's law and on q, and may change from line to line. ‖X‖_q = E[|X|^q]^(1/q).

Lemma A.0 (Generic segments). *Let σ be a segment fixed in advance, or a segment from a fixed point g to a point y, where y is chosen in an open set, or on a fixed segment τ off a finite set that depends on the world. Almost surely σ meets no face of dimension d − 2 or less of any cell and shares no piece of positive length with any facet; for the variable segments, every y in the open set outside a set of volume zero, and every y on τ outside finitely many points, gives such a σ. For such σ, the cells it meets, in the order it meets them, form a walk in G.*

Proof. Only finitely many cells meet a bounded set, so the union S of all faces of dimension at most d − 2 is, in every bounded region, a finite union of convex sets of dimension at most d − 2, and the facets are finitely many there. For a fixed segment σ, the shifts x with σ + x meeting S form S − σ, of dimension at most d − 1 and so of volume zero; the shifts making σ + x share a piece of positive length with a facet lie in finitely many sets of dimension d − 1 per bounded region, so also of volume zero. By W1 the probability that σ + x is exceptional is the same for every x; averaging x over a unit cube and exchanging expectation and integral (Fubini) shows it is zero. For a segment [g, y], the exceptional y lie in the cones from g over S and in the hyperplanes through g that contain a facet, all of dimension at most d − 1: volume zero in an open set, and finitely many points on τ unless τ lies in one of these cones or hyperplanes, which by the same shift argument has probability zero. Now let σ be generic. Where it passes from one cell to the next, it does so at a point z on no face of dimension d − 2 or less, so z lies inside a facet of every cell holding it, and near z each such cell is a half-ball cut by a hyperplane through z. Closed half-balls with disjoint interiors that cover a ball are exactly two, cut by the same hyperplane, so exactly two cells hold z and they share a (d − 1)-dimensional piece about it: they are neighbours. σ is not inside that piece, so it passes from one to the other at z. Each cell meets σ in a segment and only finitely many cells meet σ, so the cells met form a walk. ∎

Lemma A.1 (Moments). *For every q ≥ 1 and R ≥ 0, ‖N(u, R)‖_q ≤ C_q(1 + R)^d, and E[ρ(0, x)^q] ≤ C_q(1 + |x|)^q for every x.*

Proof. Cover B(u, R) by at most C(1 + R)^d unit balls. N(u, R) is at most the sum of their counts, each with the law of N(0, 1) by W1, so Minkowski's inequality and W2 give the first bound. By A.0 the cells meeting [0, x] form a walk from k(0) to k(x), so ρ(0, x) is at most the total of their waits. The segment lies in at most |x| + 2 unit balls, so the number N_x of these cells has ‖N_x‖_q ≤ C_q(1 + |x|). Given the world, a total of n ≥ 1 independent Exp(1) waits has q-th moment at most C_(q)n^q, so E[ρ(0, x)^q] ≤ C_(q)E[N_(x)^q] ≤ C_q(1 + |x|)^q. ∎

Lemma A.2 (Pushes have a finite top speed; cells are small). *There are constants a > 0, C and κ > 0 such that, for every r: (a) the probability that ρ(0, p_k) < ar for some keeper k with |p_k| ≥ r is at most Ce^(−κr); (b) the probability that the cell holding a given point has diameter above r is at most Ce^(−κr).*

Proof. Fix s and b, chosen at the end. *Wide cells make bad paths.* A cell is connected, and closed cubes in different *-connected families are disjoint, so the closed cubes meeting one cell form a single *-connected family; if the cell has diameter above s, every one of them is bad. For (b): a cell of diameter above r holding 0 meets a *-connected family of at least r/(s√d) bad cubes containing Q(0); there are at most C_1^j such families of j cubes, so by W3 the probability is at most Σ_(j ≥ r/(s√d)) $C_1^{(j)}p^j$ ≤ Ce^(−κr) once $C_1 p \le e^{-1}$.

For (a), suppose some keeper k with |p_k| ≥ r has ρ(0, p_k) < ar. Then a walk $c_0 = k(0), \ldots, c_n = k$ in G, which we may take self-avoiding, has total wait below ar. Put $y_0 = 0$, choose y_i in the intersection of the cells of c_(i−1) and c_i for 1 ≤ i ≤ n, and put y_(n+1) = p_k; then y_i and y_(i+1) both lie in the

cell of c_i. *A chain of cubes.* If that cell has diameter at most s, Q(y_i) and Q(y_(i+1)) are equal or *-neighbours, since their points differ by at most s in every coordinate. If it is wider, join them by a *-path of closed cubes meeting the cell, all bad. Erasing loops leaves a *-path Γ of m distinct cubes from Q(0) to Q(p_k); each step moves a centre at most s√d, so m ≥ r/(s√d). *Good cubes carry distinct cells.* The cubes added for wide cells are bad, so every good cube of Γ is some Q(y_i) and holds y_i, a point of the walk cell c_i (c_n when i = n + 1). That cell meets a good cube, so its diameter is at most s, and a set of diameter at most s meets at most 2^d cubes. So g good cubes on Γ carry at least g/2^d distinct walk cells, each one of the at most b cells meeting its cube.

Count 1: few bad cubes. There are at most 3^(dm) *-paths of m cubes from Q(0), and at most 2^m ways to choose which of their cubes are bad. By W3, the probability that some such path has at least m/2 bad cubes is at most $\Sigma_{m \geq r/(s\sqrt{d})}\ 3^{(dm)}2^{(m)}p^{(m/2)} \leq 2e^{(-r/(s\sqrt{d}))}$, once $2\cdot3^{(d)}p^{(1/2)} \leq e^{-1}$. *Count 2: no fast walk.* Otherwise Γ has at least m/2 good cubes, so at least $\nu = \lceil m/K \rceil$ distinct walk cells, where K = 2^(d+1). Each cube of Γ is either bad or carries one of the at most b cells meeting it, so the choices number at most (b + 1)^m. These choices depend on the world alone, and the waits are independent of it, so for each choice the probability that ν distinct waits total less than ar is at most $(ar)^\nu/\nu! \leq (e\cdot ar/\nu)^\nu \leq \beta^\nu$, where $\beta = K\cdot e\cdot a\cdot s\sqrt{d}$, using $ar \leq a\cdot s\sqrt{d}\cdot m$. Summing over paths and choices gives at most $\Sigma_m\ (3^{d(b+1)})^{(m)}\beta^{(m/K)} \leq 2e^{(-r/(s\sqrt{d}))}$, once $\beta \leq 1$ and $3^{d(b+1)}\cdot\beta^{(1/K)} \leq e^{-1}$. The constants are fixed once, in this order: s and b by W3, so that $2\cdot3^{(d)}p^{(1/2)} \leq e^{-1}$ and $C_1p \leq e^{-1}$; then a. ∎

Remark (no runaway passage). By (a) and Borel–Cantelli over whole r, almost surely only finitely many keepers are reached from k(0) by any finite time, so the first-passage recursion of Lemma 5.2 is well founded; by W1 the same holds from every fixed point.

Lemma A.3 (The front along each ray). *There is a constant c with 0 < c < ∞ such that, for every q, ρ(0, Lq)/L → |q|/c in probability and in mean as L → ∞ through the reals. The law of ρ(0, x) depends only on |x|, and for every p ≥ 1, E|ρ(0, x) − |x|/c|^p = o(|x|^p).*

Proof. As for Lemma 5.5, with three replacements. The shift by q is mixing by W1, hence ergodic, for the world with its independent waits. E[ρ(0, q)] < ∞ by A.1. For the lower bound, ρ(0, nq) = ρ(0, p_(k(nq))), and |p_(k(nq)) − nq| is at most the diameter of the cell holding nq, whose law does not depend on n and has an exponential tail by A.2(b); so A.2(a) gives γ(q) ≥ a|q| and c ≤ 1/a. By W1 the law of ρ(0, x) is unchanged by rotations about 0, so it depends only on |x| and γ(q) = |q|/c. Finally ρ(0, x)/|x| → 1/c in probability and is bounded in every L^p by A.1, so |ρ(0, x)/|x| − 1/c|^p is uniformly integrable and its mean tends to 0. ∎

Lemma A.4 (A fixed grid stands in for the keepers). *Let 𝔾 be the grid of spacing 1/√d, so that every point is within 1/2 of a grid point, with 0 in 𝔾. For g in 𝔾, let W(g) be the total wait of the cells meeting B(g, 2), and N(g) = N(g, 2). Then: (i) every keeper k with |p_k − g| ≤ 1/2 has ρ(g, p_k) ≤ W(g) and w_k ≤ W(g); (ii) W(g) and N(g) have moments of every order, with laws that do not depend on g; (iii) if keepers j and k have nearest grid points g and h, then |ρ(p_j, p_k) − ρ(g, h)| ≤ W(g) + W(h).*

Proof. (i) p_k lies inside C_k. By A.0, some y inside C_k with |y − p_k| < 1/2 gives a generic segment [g, y], so the cells it meets form a walk from k(g) to k, all meeting B(g, 1); its total wait is at most W(g). C_k holds p_k, so it meets B(g, 2). (ii) A.1, W1, and the moments of a total of N(g) independent waits. (iii) The triangle inequality, twice, with (i). ∎

Lemma A.5 (Few misjudged episodes). *The statement of Lemma 5.7, parts (1) to (3), holds for every cell world with W1, W2 and W3.*

Proof. (1) As in the proof of Lemma 5.7, the expected number of episodes of keeper k misjudged from x is at most E[(1 + Y_k)·1{Y_k > εL}], where Y_k = |τ_x(k) − s − |p_k − p₀|/c|, and by Lemma 5.2, given the world, the times τ_x(k) − s have the joint law of ρ(p₀, p_k) − w_(k(0)), all with the same waits. With g the grid point nearest p_k, A.4 gives Y_k ≤ Z_g, where Z_g = |ρ(0, g) − |g|/c| + W(g) + 2W(0) + |p₀|/c + 1/(2c). At most N(g) keepers have g as their nearest grid point. So the expected

number misjudged is at most $\Sigma_g\, E[N(g)(1 + Z_g)\cdot 1\{Z_g > \varepsilon L\}] \le \Sigma_g\, \|N(g)\|_2 \cdot \|(1 + Z_g)\cdot 1\{Z_g > \varepsilon L\}\|_2$, over the at most CL^d grid points of the spatial part of $L\cdot A$. Where $|g| < \delta L$, the second factor is at most $C(1 + \delta L)$ by A.1, A.2(b) and A.4, and these g contribute at most $C\delta^{(d+1)}L^{(d+1)}$. Where $\delta L \le |g| \le CL$: by A.3, the supremum over this range of $E|\rho(0, g) - |g|/c|$ is at most $CL\cdot \sup_{r \ge \delta L} o(r)/r = o(L)$, so $Z_g/L \to 0$ in probability uniformly in g; and Z_g/L is bounded in L^3 uniformly, by A.1, A.2(b) and A.4. Since $E[X^2\cdot 1\{X > \varepsilon L\}] \le \|X\|_3^2 P(X > \varepsilon L)^{(1/3)}$, the second factor is o(L) uniformly, and these g contribute $o(L^{(d+1)})$. Let $L \to \infty$, then $\delta \to 0$.

(2) As in Lemma 5.7: reversing time about $L\cdot t_0$ turns the past of y into a future of the same law, and by W1 the shift taking $L\cdot u$ to 0 leaves the world's law unchanged, so (1) applies.

(3) Given the world, each clock is a Poisson process, so the Mecke equation applies to the clocks alone: the expected number of pairs is the integral, over keepers k with p_k in the spatial part of $L\cdot A$ and times s' in its time range, of the expected number of episodes of $L\cdot A$ misjudged from an episode added to k's clock at s'. Lemma 5.2 holds from the added episode. With g and h the grid points nearest a source k and a target j, A.4 gives $Y \le Z_{(g,h)} = |\rho(g, h) - |g - h|/c| + 2W(g) + W(h) + 1/c$. Grouping sources and targets by grid points, the count is at most $CL\cdot \Sigma_{(g,h)} \|N(g)\|_4 \|N(h)\|_4 \|(1 + Z_{(g,h)})\cdot 1\{Z_{(g,h)} > \varepsilon L\}\|_2$. By W1 the law of $Z_{(g,h)}$ depends only on $|g - h|$, so the argument of (1) makes the last factor o(L) uniformly over pairs with $|g - h| \ge \delta L$, and at most $C(1 + \delta L)$ for the at most $CL^d(\delta L)^d$ pairs closer than that. The total is $o(L^{(2d+2)}) + C\delta^{(d+1)}L^{(2d+2)}$; let $L \to \infty$, then $\delta \to 0$. ∎

Lemma A.6 (No slow pushes). *Let $d \ge 2$. There is a constant C such that, for every q, $P(\rho(u, v) > C(1 + |u - v|)) \le C_q(1 + |u - v|)^{(-q)}$ for all points u and v.*

Proof. Lemma 5.8's proof for $d \ge 2$, with *fine* read as *good* (W3), s and b fixed so that p is small, and the grid placed so that [u, v] crosses cubes only through the interiors of their facets (W3 holds for every placement, and almost every placement does this). *Few cells per good cube:* at most b, by definition. *Short walks between good cubes sharing a facet:* the segment between their centres is fixed in advance and lies in the two cubes, so it meets at most 2b cells, which form a walk by A.0. *Hulls*, facts (a) and (b), *the walk* and *the splice* are statements about cubes alone and hold as written, with Timár's theorem as cited there. *Count 1* holds with W3 in place of independence beyond range $2L_0$: a set of j cubes, at least $j - M$ of them not good, all fail with probability at most $p^{(j-M)}$. *Count 2:* as there, $P(S_u > t) \le Ce^{(-\kappa t)}$. The walk from k(u) runs along [u, v] to a point z of [u, v] inside e_u, chosen off the finitely many exceptional points of A.0, then along [z, centre of e_u]; every cell it meets before e_u's centre meets $B(u, s\sqrt{d}(S_u + 2))$, so $N_u \le N(u, s\sqrt{d}(S_u + 2))$ and, for $t \ge 1$, with $\tau = t^{(1/(2d))}$, $P(N_u > t) \le P(S_u > \tau) + P(N(u, s\sqrt{d}(\tau + 2)) > t) \le C\cdot\exp(-\kappa\tau) + C_{(q)}t^{(-q/2)}$, by A.1 and Markov's inequality. The same holds at v, and when one hull holds both ends. *Count 3* is unchanged. Adding the failure probabilities gives the bound. ∎

Remark. Lemma 5.8's bound was a stretched exponential; this one beats every power of the distance, which is all Lemma A.7 needs. If W2 is strengthened to a stretched-exponential tail for N(0, 1), the stretched exponential returns.

Lemma A.7 (Rare slow pushes, from every episode). *Let $d \ge 2$, fix $\delta > 0$ and a bounded region A of space-time, and let $C' = C + 1$ with C from A.6. Let F_L be the number of episodes $z = (k, t)$ in $L\cdot A$ for which some keeper j within δL of p_k has $\tau_z(j) - t > C'(6 + 3\delta L)$. Then $E[F_L] \to 0$. The same holds for the past of every episode.*

Proof. By the Mecke equation applied to the clocks, as in A.5(3), $E[F_L] \le CL\cdot E[\Sigma_k \varphi_k]$, over keepers k with p_k in the spatial part A' of $L\cdot A$, where φ_k is the probability, given the world, that from an episode added to k's clock some keeper j within δL of p_k is reached more than $C'(6 + 3\delta L)$ later. By Lemma 5.2 this is the probability, over fresh waits, that $\rho(p_k, p_j) - w_k$ exceeds $C'(6 + 3\delta L)$ for some such j. A.6 bounds a passage only between points far apart, so near pairs go by a detour, as in Lemma 5.9: for grid points g and h with $|g - h| \le \delta L + 1$, let m be the grid point nearest the point δL from their midpoint in a direction perpendicular to $h - g$, which exists because $d \ge 2$; then $|g - m|$ and $|m - h|$ lie between $\delta L - 1$ and $1.2\delta L + 2$. Let E_L be the event that $W(g) > \delta L$ for some grid point

g within CL of the origin, or ρ(a, b) > C(1 + |a − b|) for some such grid points a and b with δL − 1 ≤ |a − b| ≤ 1.2δL + 2. Off E_L, ρ(g, h) ≤ ρ(g, m) + ρ(m, h) ≤ C(6 + 3δL) for every such g and h, and A.4(iii) gives ρ(p_k, p_j) ≤ C(6 + 3δL) + 2δL ≤ C′(6 + 3δL) for every k and j at once, so φ_k ≤ P(E_L | world) for every k. Hence E[F_L] ≤ CL·E[N_L·P(E_L | world)] ≤ CL·‖N_L‖₂P(E_L)^(1/2), where N_L, the number of keepers in A′, has ‖N_L‖₂ ≤ CL^d by A.1. By A.6, and A.4(ii) with Markov's inequality, P(E_L) ≤ CL^d(δL)^d·C_q(δL)^(−q) + CL^d·C_q(δL)^(−q) for every q. Taking q large, E[F_L] → 0. The past follows by the time reversal of Lemma 5.1. ▮

A.8 Proof of Theorem 5.11. The proof of Theorem 5.10, with its proof of item 4, applies as written, with Lemmas A.1 to A.7 for Lemmas 5.3 to 5.9 and two further changes. *Step 1 (density).* The keepers form a stationary point process of intensity 1, each with an independent rate-1 clock, so the episodes form a stationary point process in space and time with one episode per unit volume. Space shifts are mixing for the world by W1, and stay mixing when every keeper carries an independent clock, a standard property of independent marking; so the episode process is ergodic under shifts of space and time, and the ergodic theorem for stationary point processes applies as in Theorem 5.10. *Tips.* The offset of x from the origin is at most the diameter of the cell holding 0, with an exponential tail by A.2(b), plus an Exp(1) time; likewise for y_L, and for the grid episodes e_g of item 4, each of whose keepers lies within δL/2 of L·q_g with probability tending to 1, by A.2(b). In item 4, δ is taken below η/(3C′), so that C′(6 + 3δL) < ηL for large L. ▮

Proposition A.9 (The Poisson world). *For every d ≥ 1, the Poisson–Voronoi world, with keepers at the points of a Poisson process of intensity 1 and cells their Voronoi cells, has W1, W2 and W3.*

Proof. W1. The Poisson process is stationary, isotropic and mixing, and the Voronoi world is built from it by a rule that commutes with shifts and rotations. *W2.* Let D be the largest distance from a point of B(0, 1) to its nearest Poisson point. A cell meeting B(0, 1) at y has its point within D of y, so N(0, 1) ≤ Π(B(0, 1 + D)). Lemma 5.3's argument gives P(D > t) ≤ C·exp(−κt^d), so E[N(0, 1)^q] ≤ Σ_m E[Π(B(0, 1 + m))^(2q)]^(1/2)P(D > m − 1)^(1/2) < ∞. *W3.* Let Q be a closed grid cube of side s, and let E_Q be the event that some cube of side s/(4√d), in the grid of such subcubes covering the points within s of Q, holds no Poisson point, or that more than b Poisson points lie within s/4 of Q. Off E_Q, every point within s of Q has a Poisson point within s/4. Take a cell C_k meeting Q at y_0; then $|y_0$ − p_k| ≤ s/4. If some z in C_k had |z − p_k| > s/2, the segment [y_0, z], which lies in the convex cell, would hold a point z′ with |z′ − p_k| = s/2 and |z′ − y_0| ≤ 3s/4, so within s of Q; z′ would then have a Poisson point within s/4, yet its nearest Poisson point is p_k, at s/2. So every cell meeting Q has diameter at most s and its point within s/4 of Q, and at most b cells meet Q. Hence *bad* implies E_Q, and E_Q depends only on the Poisson points within 2s of Q, so cubes whose indices differ by 6 or more in some coordinate have independent E_Q. Any finite set J of cubes contains at least |J|/6^d cubes pairwise that far apart (one residue class mod 6), so P(all of J bad) ≤ P(E_Q) to the power |J|/6^d, and P(E_Q) ≤ (12√d + 4)^d·exp(−(s/(4√d))^d) + P(Poisson((3s/2)^d) > b), which is as small as we like once s and then b are large. Set p(s, b) = P(E_Q) to the power 1/6^d. ▮

Appendix B. Proof of Theorem 6.1

B.1 Finite range and local stability. With the cut at R_cut, two keepers are neighbours only if they lie on an empty ball of radius at most R_cut, so within 2R_cut of each other, and whether they are is decided by the keepers within 2R_cut of them. By induction on the number of crossings, B_k(r), and with it Φ_k, is decided by the keepers within R_int = 2R_cut(r_max + 1) of k, where r_max is the deepest depth. Every cost is at least 0, so H ≥ 0: the energy is stable. With no two keepers closer than δ, a ball of radius R_int + δ/2 holds at most K_δ = ((2R_int + δ)/δ)^d disjoint balls of radius δ/2, so at most K_δ keepers lie within R_int of any point. Adding or removing one keeper changes at most K_δ + 1 costs, each by at most λ|𝒟|M, so the energy change ΔH(x, ω) = H(ω ∪ {x}) − H(ω) obeys

$$|\Delta H(x, \omega)| \le C_\lambda = \lambda \cdot |\mathcal{D}| \cdot M \cdot (K_\delta + 1)$$

The costs that change are those of keepers within R_int of x, and each is decided by the keepers within R_int of its own keeper, so ΔH(x, ω) is decided by the keepers within 2R_int of x.

B.2 Existence. A stable, finite-range interaction with a hereditary hardcore (removing a keeper never breaks the rule that keepers stay δ apart) has at least one infinite-volume Gibbs world, the limit along growing boxes of its finite-volume Gibbs worlds. The boxes' worlds are tight because each is dominated by a Poisson process of intensity z·exp(C_λ), by local stability, and the limit inherits the local specification because the interaction has finite range (Ruelle 1970; Preston 1976; Dereudre 2019). Results for Voronoi and Delaunay interactions specifically are in Dereudre (2008) and Dereudre, Drouilhet and Georgii (2012).

B.3 Uniqueness, and a factor of independent randomness. Theorem 6.1 concerns the steady state, so any birth-and-death process in detailed balance with it may be used. Take the one in which every keeper dies at rate 1 and a keeper is born at x at rate

$$\mathbf{b(x, \omega) = z \cdot I_\delta(x, \omega) \cdot \exp(-\Delta H(x, \omega))}$$

where I_δ(x, ω) is 1 if no keeper of ω lies within δ of x and 0 otherwise. Its ratio of birth to death rates, z·exp(−ΔH), is that of the split form of Section 6, so it has the same steady state (Preston 1975); and its birth rate never exceeds $\bar{z} = z \cdot \exp(C_\lambda)$.

The free process. Let Ξ be a stationary Poisson process of candidate births ζ = (x, t, s, u) on $\mathbb{R}^d \times \mathbb{R} \times \mathbb{R}_+ \times [0, 1]$, at rate $\bar{z}$ per unit volume and time, each with an exponential lifetime s of mean 1 (alive on [t, t + s)) and a uniform mark u. The candidates alive at any moment form a Poisson process in space of intensity $\bar{z}$. A candidate is kept if and only if

$$\mathbf{u \le b(x, \eta_(t-))/\bar{z} = I_\delta(x, \eta_(t-)) \cdot \exp(-\Delta H(x, \eta_(t-)) - C_\lambda)}$$

where η_(t−) is the set of kept keepers alive just before t. By B.1 the right side is at least exp(−2C_λ) whenever I_δ = 1. So a candidate with u ≤ exp(−2C_λ) is kept exactly when no kept keeper lies within δ of x, and need look only at candidates alive at t− inside B(x, δ); a candidate with u > exp(−2C_λ), a share at most 1 − exp(−2C_λ), must look at those inside B(x, 2R_int) to compute ΔH. Its look radius is δ in the first case and 2R_int in the second. The candidates it looks at are its ancestors; following looks back in time gives its *clan of ancestors* (Fernández, Ferrari and Garcia 2001, 2002).

The branching bound. A candidate's ancestors are the candidates alive at its birth inside its look ball, and its look radius, δ or 2R_int, is fixed by its own mark alone. Explore a clan backwards, one candidate at a time, and replace each part of Ξ already revealed by an independent copy, which can only add to the clan: a candidate's ancestors in a part already revealed were counted when it was revealed, and the rest lie in the unrevealed part of its look ball at its birth, so drawing that whole region afresh gives at least as many new ancestors. The candidates alive at any moment form a Poisson process of intensity $\bar{z}$ with uniform marks independent of everything else, so each candidate found has a Poisson number of ancestors, with mean $m_s = \bar{z} \cdot v_(d)\delta^d$ when its look radius is δ and $m_l = \bar{z} \cdot v_d(2R_int)^d$ when it is 2R_int, and each ancestor, like each candidate at the start, has the larger radius with probability q = 1 − exp(−2C_λ), independently of the candidate that found it. The clan is therefore dominated by a two-type Galton–Watson process whose mean matrix has every row proportional to (1 − q, q), so the expected size of its generation n is the n-th power of m_s(1 − q) + m_(l)q, which is at most

$$\mathbf{\alpha = \bar{z} \cdot v_(d)\delta^d + (1 - \exp(-2C_\lambda)) \cdot \bar{z} \cdot v_d(2R_int)^d}$$

With $z \cdot v_(d)\delta^d \le 1/4$, as Theorem 6.1 assumes, take

$$\mathbf{\lambda_0 = \min(\ln 2 / (|\mathscr{D}|M(K_\delta + 1)),\ 1 / (16 \cdot 2^d \cdot |\mathscr{D}|M(K_\delta + 1) \cdot z \cdot v_(d)R_(int)^d))}$$

For $0 \le \lambda < \lambda_0$, exp(C_λ) ≤ 2 and 1 − exp(−2C_λ) ≤ 2C_λ, so the first term of α is at most 2·(1/4) = 1/2 and the second at most $2C_\lambda \cdot 2z \cdot v_d(2R_int)^d = 4 \cdot 2^(d)\lambda|\mathscr{D}|M(K_\delta + 1)z \cdot v_(d)R_(int)^d < 1/4$. Hence **α ≤ 3/4 < 1**: the process is subcritical.

Tails. Clans are finite almost surely and their size has an exponential tail (total progeny of a subcritical Galton–Watson process). Each ancestor step moves at most 2R_int, since δ < 2R_cut < 2R_int, so a clan reaching distance D has a line of at least n = ⌈D/(2R_int)⌉ generations, which has probability at most the expected size of generation n, α^n ≤ e^(−γD) with γ = |ln α|/(2R_int) > 0. For a bounded region A, the clans of the Poisson($\bar{z}$|A|) candidates alive in A at time 0 reach a distance R(A) with

$$\mathbf{P(\,R(A) > D\,) \le \bar{z}\cdot|A|\cdot e^{(-\gamma D)}}$$

Since every clan is finite, the backward construction decides which candidates in A are kept from Ξ inside the clan, with no condition at infinity. The world so built is the steady state of the process; every Gibbs world for the interaction is reversible for it, and it is the unique one (Preston 1975; Fernández, Ferrari and Garcia 2002). **The tension world is unique, and a factor of the independent process Ξ that commutes with translations and rotations.**

B.4 W1, W2 and W3. *W1.* Ξ is unchanged by translations and rotations, and the cut, the extent and the reach use only distances and the foam; by uniqueness the tension world is stationary and isotropic. Off the event {R(A) > D}, every ancestor of a candidate alive in A lies within D of A, and its look ball within D + 2R_int, so the world in A is decided by Ξ within D + 2R_int of A. Observables on regions A and B at distance more than 2(D + 2R_int) are therefore decided by disjoint parts of Ξ outside the event {R(A) > D} ∪ {R(B) > D}, of probability at most $\bar{z}$(|A| + |B|)e^(−γD) by B.3, so their correlation vanishes as the distance grows: the world is mixing.

W2. Kept keepers are among the free candidates alive at time 0, so the world is dominated by a Poisson process of intensity $\bar{z}$. For empty balls, let t ≥ 3δ and place m ≥ c_1(t/δ)^d test balls B(y_j, δ/4) inside B(0, t − 5δ/4), with centres at least 2δ apart. Given the world outside B(0, t), the probability that B(0, t) is empty is 1/Z, where Z is the conditional partition function inside B(0, t) relative to the empty configuration. Configurations with at most one keeper in each test ball and none elsewhere in B(0, t) obey the hardcore, since test balls are more than δ apart from each other and from everything outside B(0, t), and each added keeper costs at most C_λ by B.1, so

$$\mathbf{P(\,B(0,t)\ empty\,) \le (\,1 + z\cdot\exp(-C_\lambda)\cdot v_d(\delta/4)^d\,)^{(-m)} \le \exp(\,-c_void\cdot z\cdot\exp(-C_\lambda)\cdot t^d\,)}$$

for a constant c_void > 0 depending on d and zδ^d. A Voronoi cell meets the unit ball B(0, 1) with its keeper outside B(0, R) only if some point of B(0, 1) has no keeper within R − 1; covering B(0, 1) by a finite net, this has probability at most C·exp(−c_void·z·exp(−C_λ)·(R − 2)^d). Otherwise every cell meeting B(0, 1) has its keeper in B(0, R), and their number is dominated by Poisson($\bar{z}$·v_(d)R^d). The two tails together give the number of cells meeting a unit ball moments of every order, as in Proposition A.9.

W3. Take grid cubes Q_i of side s > 2R_int, and let E_i be Proposition A.9's local event: some cube of side s/(4√d), in the grid of such subcubes covering the points within s of Q_i, holds no keeper, or more than b keepers lie within s/4 of Q_i. The proof of Proposition A.9 shows, for any set of keepers, that a bad cube has E_i; and E_i is decided by the keepers in $\tilde{Q}$_i = {x : dist(x, Q_i) ≤ 2s}. By the bound on empty balls and Poisson domination, P(E_i) ≤ ε for any ε > 0 once s and then b are large.

Let U_i = {y : dist(y, $\tilde{Q}$_i) ≤ s}. Explore the clans of the candidates alive in $\tilde{Q}$_i at time 0, and let F_i be the event that some ancestor found has a look ball not contained in U_i. Until that happens every ancestor found lies in U_i, so F_i is decided by Ξ restricted to U_i; and F_i requires a clan to reach beyond s − 2R_int, so by B.3

$$\mathbf{P(F_i) \le \bar{z}\cdot(5s)^d\cdot\exp(-\gamma(s - 2R_int))}$$

which is at most ε for s large. Off F_i the explored clans are the true clans, so the world in $\tilde{Q}$_i, and with it E_i, is decided by Ξ inside U_i. Hence $\tilde{E}$_i = E_i ∪ F_i is decided by Ξ inside U_i, with P($\tilde{E}$_i) ≤ 2ε. U_i lies within 3s of Q_i, so cubes whose indices differ by 8 or more in some coordinate have disjoint U_i, and their $\tilde{E}$_i are independent. Any finite set J of cubes contains a set J′ of at least |J|/8^d cubes pairwise that far apart (one residue class mod 8), so

P(every cube of J bad) ≤ P(∩_(i ∈ J′) Ẽ_i) = ∏_(i ∈ J′) P(Ẽ_i) ≤ (2ε)^(|J′|) ≤ p^(|J|)

with p = (2ε) to the power 1/8^d, as small as we like once s and b are large. This is W3, for every placement of the grid by stationarity.

Conclusion. The tension world's keeper density is finite, since the world is dominated by a Poisson process of intensity $\bar{z}$, and positive, by the bound on empty balls. Its Voronoi foam, with keepers at their own points and lengths scaled by that density so that there is one keeper per unit volume, which leaves W1 to W3 unchanged, has W1, W2 and W3, so Theorem 5.11 applies: its order of held boundaries converges at large scale to Minkowski spacetime, in every frame. ∎

References

Formal and empirical records at https://3primitives.io/formal_records/

Passell, P. R., & Gildenston, S. (2025/2026). Three Primitives: Canonical Logic Sequence. Formal Record 1, v2.3.

Gildenston, S., & Passell, P. R. (2026). AI Cannot Govern AI. Formal Record 2, v2.4.

Gildenston, S., & Passell, P. R. (2026). The Law of Declared Authority. Formal Record 6, v2.3.

Gildenston, S., & Passell, P. R. (2026). Bell Non-Closure and the Law of Declared Authority. Formal Record 7, v2.2.

Gildenston, S., & Passell, P. R. (2026). Primitive Stability Theorem. Formal Record 10, v2.2.

Gildenston, S. (2026). The GBSH Correspondence. Formal Record 11, v2.0.

Gildenston, S. (2026). The Forced Bijection: On the Structural Necessity of the Governance Primitives in Conscious Reality. Formal Record 12, v1.7.

Gildenston, S., & Passell, P. R. (2026). The ILMM Coupling Theorem. Formal Record 13, v2.4.

Gildenston, S. (2026). The Individuation Theorem: Life as the Boundary of the Subject. Formal Record 14, v5.8.

Gildenston, S. (2026). The Boundary Theorem: On the Structural Necessity of Held Boundaries in Authority, Personhood and Life. Formal Record 15, v2.4.

Gildenston, S., & Passell, P. R. (2026). The Spontaneous Legibility Assumption: The Chi Recovery Experiment. Empirical Record 1, v1.6.1.

Gildenston, S., & Passell, P. R. (2026). The Structural Floor Hypothesis. Empirical Record 2, v1.3.

Gildenston, S., & Passell, P. R. (2026). The Floor Is Learned. Empirical Record 3, v1.1.

Gildenston, S., & Passell, P. R. (2026). Cross-Architecture Scaling: The Structural Forcing Distinction Is Itself a Training Artifact. Empirical Record 4, v1.1.

Other works

Almada Lopez, D. (2026, 12 October). Anthropic warns small AI models may soon self-replicate on IoT devices. *Crypto Briefing*.

Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3), 195–200.

Boivin, D. (1990). First passage percolation: the stationary case. *Probability Theory and Related Fields*, 86, 491–499.

Bombelli, L., Lee, J., Meyer, D., & Sorkin, R. D. (1987). Space-time as a causal set. *Physical Review Letters*, 59(5), 521–524.

Bowyer, A. (1981). Computing Dirichlet tessellations. *The Computer Journal*, 24(2), 162–166.

Daley, D. J., & Vere-Jones, D. (2008). *An Introduction to the Theory of Point Processes*, Volume II. Springer.

Dereudre, D. (2008). Gibbs Delaunay tessellations with geometric hardcore conditions. *Journal of Statistical Physics*, 131(1), 127–151.

Dereudre, D. (2019). Introduction to the theory of Gibbs point processes. In D. Coupier (Ed.), *Stochastic Geometry*, Lecture Notes in Mathematics 2237 (pp. 181–229). Springer.

Dereudre, D., Drouilhet, R., & Georgii, H.-O. (2012). Existence of Gibbsian point processes with geometry-dependent interactions. *Probability Theory and Related Fields*, 153(3–4), 643–670.

Errera, L. (1886). Sur une condition fondamentale d'équilibre des cellules vivantes. *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences*, 103, 822–824.

Fernández, R., Ferrari, P. A., & Garcia, N. L. (2001). Loss network representation of Peierls contours. *Annals of Probability*, 29(2), 902–937.

Fernández, R., Ferrari, P. A., & Garcia, N. L. (2002). Perfect simulation for interacting point processes, loss networks and Ising models. *Stochastic Processes and their Applications*, 102(1), 63–88.

Hoffman, D. D., Singh, M., & Prakash, C. (2015). The interface theory of perception. *Psychonomic Bulletin & Review*, 22(6), 1480–1506.

Kingman, J. F. C. (1968). The ergodic theory of subadditive stochastic processes. *Journal of the Royal Statistical Society B*, 30, 499–510.

Kleitman, D. J., & Rothschild, B. L. (1975). Asymptotic enumeration of partial orders on a finite set. *Transactions of the American Mathematical Society*, 205, 205–220.

Last, G., & Penrose, M. (2017). *Lectures on the Poisson Process*. Cambridge University Press.

Liggett, T. M. (1985). An improved subadditive ergodic theorem. *Annals of Probability*, 13, 1279–1285.

Loomis, S. P., & Carlip, S. (2018). Suppression of non-manifold-like sets in the causal set path integral. *Classical and Quantum Gravity*, 35(2), 024002.

Malament, D. B. (1977). The class of continuous timelike curves determines the topology of spacetime. *Journal of Mathematical Physics*, 18(7), 1399–1404.

Meyer, D. A. (1988). *The Dimension of Causal Sets*. PhD thesis, Massachusetts Institute of Technology.

Milinkovic, B., & Aru, J. (2025). On biological and artificial consciousness: A case for biological computationalism. *Neuroscience and Biobehavioral Reviews*.

Myrheim, J. (1978). Statistical geometry. CERN preprint TH-2538.

Preston, C. (1975). Spatial birth-and-death processes. *Bulletin of the International Statistical Institute*, 46(2), 371–391.

Preston, C. (1976). *Random Fields*. Lecture Notes in Mathematics 534. Springer.

Rideout, D. P., & Sorkin, R. D. (1999). A classical sequential growth dynamics for causal sets. *Physical Review D*, 61(2), 024002.

Ruelle, D. (1970). Superstable interactions in classical statistical mechanics. *Communications in Mathematical Physics*, 18(2), 127–159.

Sapolsky, R. M. (2023). *Determined: A Science of Life Without Free Will*. Penguin Press.

Timár, Á. (2013). Boundary-connectivity via graph theory. *Proceedings of the American Mathematical Society*, 141, 475–480.

Watson, D. F. (1981). Computing the n-dimensional Delaunay tessellation with application to Voronoi polytopes. *The Computer Journal*, 24(2), 167–172.

Ziesche, S. (2016). First passage percolation in Euclidean space and on random tessellations. arXiv:1611.02005.

Attribution and Licence

This document is the intellectual property of Stacy Gildenston, held under 3primitives.io. Published under Creative Commons Attribution 4.0 International (CC BY 4.0). You are free to share, adapt and build upon this work for any purpose, including commercially, provided you give appropriate credit to 3primitives.io and Stacy Gildenston as originator.

Suggested citation:

Gildenston, S. (2026). Classical, Quantum, Consciousness: On the Production of Causal Order and Spacetime by Held Boundaries. Three Primitives Framework, Formal Record 16, v3.6. Melbourne, Australia. First published 20 October 2026. 3primitives.io. CC BY 4.0.