ChronaQ Foundations II: Compatible Entanglement Geometries and Finite Causal Structure

Working paper

Stephen J. McIntyre

McIntyre Lab, Auckland, New Zealand

stephen.mcintyre@mcintyrelab.org

ORCID: 0009-0008-5202-2005

3 August 2026

Manuscript revision 0.4

Abstract

ChronaQ proposes that observable relational structure constrains geometry without assuming that a unique spacetime metric is primitive. This paper develops that proposal for finite systems. The input is a rendered observable patch: a finite collection of operational objects together with positive distributions, correlations, graph relations, or response data licensed by the ChronaQ axioms. Geometry is introduced through declared constructors and explicit admissibility gates. The resulting CQ2 object is generally a family of compatible geometries, not a single distance matrix. A direct counterexample shows why this distinction is necessary. The literal reciprocal-mutual-information score fails the zero-diagonal metric axiom, and an explicitly diagonal-corrected version still violates the triangle inequality by an arbitrarily large amount and becomes infinite at independence. On one heterogeneous four-object patch, graph shortest-path distance, Hellinger distance between positive observable distributions, and an L1 polyhedral/Finsler chart all satisfy the finite metric axioms while remaining inequivalent. A prospective same-data benchmark then begins with one frozen finite quantum-walk event-count table. Graph shortest-path, Hellinger early-response, and spectral-resistance constructors all produce valid metrics from its unopened early layers under one common 256-replicate uncertainty law. Candidate-specific scales are fitted on common training rows and frozen. On untouched longer-range rows, the graph renderer has a held-out mean absolute error of 0.474 record layer, compared with 2.340 and 2.164 layers for the Hellinger and spectral alternatives. Its uncertainty-bounded advantage clears the preregistered materiality gate against both. The benchmark therefore shows that later operational data can narrow a compatible geometry family without refitting. A separate nearest-neighbour influence assumption gives an exact polyhedral support bound and an earliest-permitted diamond front. Equality with actual first arrival requires an additional nonzero propagation witness and cancellation control. Continuous spectral propagation instead has nonzero tails and supplies an effective causal band rather than a sharp cone. These results support a disciplined form of entanglement-geometry compatibility: observable data constrain and can discriminate admissible representations, while unique physical geometry, Lorentzian signature, physical scale, an entropy-area law, and gravitational dynamics require additional gates.

Keywords: entanglement geometry; mutual information; Hellinger distance; graph metric; Finsler geometry; causal structure; quantum information; emergent spacetime; ChronaQ

1. Introduction

1.1 From relational order to relational geometry

The first ChronaQ Foundations paper constructs an internal order from finite two-boundary record refinement. It separates the number of surviving record lineages, their Hartley entropy, the Shannon entropy of a positive record law, and a chosen normalized Shannon coordinate. None of those quantities is a spatial distance. They order rendered records without specifying how distinct records, regions, or observable contexts are separated from one another (McIntyre, 2026). The next question is therefore not whether information is geometry by definition, but what geometric statements are justified by relational quantum data.

Entanglement has a well-established connection to geometry in several restricted settings. In holographic theories, the Ryu-Takayanagi relation associates boundary entanglement entropy with the area of a bulk extremal surface (Ryu & Takayanagi, 2006). Changes in entanglement can encode linearized gravitational constraints when a holographic dictionary, a suitable state class, modular relations, and a semiclassical bulk are already available (Faulkner et al., 2014). Tensor-network constructions show how scale-dependent entanglement patterns can admit useful geometric representations (Swingle, 2012), and broader programmes investigate how spatial relations might be recovered from Hilbert-space structure (Cao et al., 2017; Van Raamsdonk, 2010). These results motivate a relational geometry programme, but they do not imply that an arbitrary information statistic is automatically a metric or that every finite entanglement pattern identifies one spacetime.

ChronaQ approaches the problem axiomatically. Its fourth axiom states that relational quantum structure constrains a family of compatible geometric descriptions. The word *family* is load-bearing. Observable data may rule out many representations and leave several. A constructor may turn an observable graph, a family of probability distributions, or an exported coordinate chart into a metric. A separate gate must show that the constructed object satisfies the claimed axioms. Further gates are needed before that object can be called causal, Lorentzian, continuum, or physical. Treating these stages separately converts a suggestive slogan into a testable mathematical interface.

This paper develops that interface as CQ2. The finite core has five parts. First, a compatible-geometry construction returns every candidate that is consistent with the declared data, constructor, and metric gates. Second, a counterexample excludes reciprocal mutual information as a general distance, both in its literal form and after an explicit diagonal repair. Third, a heterogeneous four-object model demonstrates typed underidentification: three inequivalent metrics pass their respective gates on one declared patch while reading different supplied descriptors. Fourth, a frozen same-data quantum-walk benchmark compares three independently motivated constructors under common uncertainty and uses untouched response rows to narrow the family without refitting. Fifth, a finite influence model supplies an exact polyhedral support bound and earliest-permitted causal lane, while distinguishing that bound from an attained first-arrival front, latent overlap, and continuous propagation with nonzero tails.

The result is constructive rather than merely cautionary. It shows how a ChronaQ geometry can be represented, audited, perturbed, and refined. It also identifies what additional evidence would select a unique physical geometry: an operational discriminator, a uniqueness theorem, a stable continuum limit, or a common dynamical principle that eliminates the alternatives without being fitted after their outputs are known.

1.2 The scientific programme

The intended Foundations progression is preserved. A1 removes primitive external time; A2 imposes two-boundary admissibility; A3 supplies a positive observable rendering; and A4 makes relational structure geometrically consequential. P1 uses A1-A3 to derive a record order and a dimensionless chart. P2 activates A4. The aim is to move from correlations among rendered regions to admissible geometry, then to specify the conditions under which geometry could support causal and dynamical language. Later papers address latent completion, gluing, action, sources, and experiments.

The mathematical strength of each link must be explicit. The progression is not

> entanglement statistic → unique metric → entropy stationarity → Einstein equation.

It is instead

> rendered relational data → declared geometry constructors → metric and stability gates → compatible geometry family → additional physical selection → continuum and action interfaces.

ChronaQ retains the ambition of deriving familiar spacetime and gravitational structure from its axioms rather than proposing a new gravity law. The programme does not place the conclusion inside the definition of distance. That discipline matters because the strongest later result should follow from the axioms and independent gates, not from an early choice that has been relabelled as a theorem.

1.3 Contributions

The paper makes eight contributions.

1. It states A1-A4 and gives an explicit premise map showing that A4 supplies the P2 geometry interface while P1's CQ1 output supplies order only.
2. It defines a finite rendered patch, a compatible state family, geometry constructors, metric gates, equivalence relations, and the family-valued CQ2 output.

3. It gives an exact classical probability counterexample showing that the literal reciprocal-mutual-information score fails the zero-diagonal axiom and that an explicitly diagonal-corrected version still fails the triangle inequality.
4. It constructs three distinct metrics from three declared descriptor types on one four-object patch and records their finite metric, perturbation, and refinement audits.
5. It proves the corresponding finite non-identification statement: when two inequivalent candidates survive all declared inputs and gates, those inputs do not entail a unique geometry.
6. It holds one synthetic quantum-walk event dataset fixed, propagates one multinomial uncertainty law through three independently justified metric constructors, and obtains distinct no-refit held-out operational predictions.
7. It distinguishes latent completion adjacency, rendered distance, exact controllable-influence support bounds, separately witnessed first-arrival fronts, effective propagation bands, and physical light cones.
8. It states explicit interfaces for entropy-area relations, Lorentzian recovery, continuum limits, action principles, conserved sources, the existing Einstein or Einstein-Cartan law, and physical scale.

1.4 Organization and status language

Section 2 places CQ2 relative to entanglement geometry, information metrics, graph geometry, locality bounds, and Finsler structure. Section 3 states the ChronaQ axioms and their P2 roles. Section 4 defines the finite construction. Section 5 excludes reciprocal mutual information as a general metric and states CQ2. Section 6 gives the heterogeneous finite model. Section 7 gives the frozen same-data discriminator. Section 8 develops the causal-geometry interface. Section 9 addresses entropy, area, continuum geometry, and gravity. Sections 10 and 11 give controls, falsifiers, interpretation, and limitations. Appendices provide proofs and complete numerical tables.

Four status terms keep premise and conclusion distinct. *Chosen* denotes an axiom, definition, constructor, coordinate chart, tolerance, or modelling convention. *Forced* denotes a conclusion whose contrary is excluded by proof or counterexample under the stated assumptions. *Conditional* denotes a theorem or calculation valid on a declared domain. *Conjectural* denotes a proposed physical extension awaiting an independent discriminator. A4 is chosen as a ChronaQ axiom; failure of reciprocal mutual information as a universal metric is forced; the finite compatible-family and polyhedral results are conditional; and unique Lorentzian recovery remains conjectural.

2. Background and relation to established approaches

2.1 Entanglement and geometry

Entanglement entropy is not ordinarily a distance between arbitrary subsystems. It quantifies a property of a state relative to a bipartition. Mutual information measures total statistical dependence between two systems and is nonnegative, symmetric, and zero for a product state, but it does not generally obey a distance triangle inequality. Those distinctions coexist with powerful geometric results in special theories.

The Ryu-Takayanagi formula relates the entropy of a boundary region in an appropriate holographic state to the area of a codimension-two bulk surface (Ryu & Takayanagi, 2006). Van Raamsdonk (2010) argued that reducing entanglement can be understood as pulling spacetime regions apart. Swingle (2012) connected tensor-network renormalization structure to an emergent holographic direction. Faulkner et al. (2014) showed how the entanglement first law, combined with holographic ingredients, yields linearized gravitational equations. Jacobson (2016) obtained the Einstein equation from a local entanglement-equilibrium hypothesis together with assumptions about small causal diamonds and ultraviolet entanglement. These are important precedents for ChronaQ's goal, but each conclusion is conditional on a substantial mathematical setting.

The lesson adopted here is twofold. Entanglement can constrain geometry, and the constraint must be typed. One must state the regions or algebras, the class of states, the information quantity, the map to geometry, the equivalence relation, and the regime in which the map is expected to hold. A result established in anti-de Sitter/conformal field theory cannot be imported unchanged into a finite two-boundary construction. Conversely, a finite relational construction need not imitate holography to be meaningful; it must pass the gates appropriate to the geometry it claims.

Mutual information remains useful even when its reciprocal is not a metric. It can identify dependence, bound connected correlations, provide graph weights, and help infer neighbourhood structure under additional assumptions (Cover & Thomas, 2006; Wolf et al., 2008). It can therefore enter a geometry constructor as data. The error would be to assume that the transform from dependence to distance is unique or automatically metric.

2.2 Metrics on finite observable objects

A metric on a finite set X is a function $d: X \times X \to [0,\infty)$ satisfying finiteness, nonnegativity, identity of indiscernibles, symmetry, and the triangle inequality. A pseudometric relaxes identity of indiscernibles, allowing distinct objects at zero distance. An extended metric may permit infinity. These are different claims and should not be interchanged because a candidate behaves badly at an edge case.

Finite data admit many conventional metric constructions. Positive edge weights on a connected undirected graph induce a shortest-path metric. Normalized probability vectors admit Hellinger distance, which is Euclidean distance after the square-root embedding up to a conventional normalization; the square root of the Jensen-Shannon divergence supplies another genuine probability metric (Endres & Schindelin, 2003). Partitions of one underlying set admit the variation-of-information metric, which uses conditional entropies rather than the reciprocal of mutual information (Meilă, 2007). A coordinate chart with an L1 norm gives a polyhedral metric whose unit ball is a diamond in two dimensions. Each construction has a different input ontology. Graph distance reads connectivity and edge costs; Hellinger and Jensen-Shannon metrics compare probability distributions; variation of information reads a pair of partitions of a common domain; L1 distance reads a declared coordinate chart and norm. Agreement among them is an empirical or mathematical result, not a prerequisite for calling each one a metric on its own domain.

Information geometry often places Riemannian structure on smooth statistical models, most notably through the Fisher information metric (Amari, 2016). Quantum statistical distance likewise has a developed state-space geometry tied to measurement distinguishability and the quantum Fisher information (Braunstein & Caves, 1994). Those settings assume a specified state or parameter manifold and sufficiently regular statistical structure. A finite classical Hellinger geometry can be a useful precursor or control, but the existence of a finite statistical or quantum-state metric does not by itself supply a spacetime manifold, Lorentzian signature, or physical units.

2.3 Graph, normed, and Finsler geometry

Graph metrics and normed vector spaces are exact geometries, not merely numerical approximations to Riemannian geometry. A square lattice with equal coordinate-step costs has L1 shortest-path distance. Its unit ball is polyhedral. Refining the lattice while preserving those costs converges, in the relevant chart, to an L1 norm rather than to the Euclidean norm. Smoothness and rotational isotropy therefore cannot be inferred from refinement alone.

Finsler geometry generalizes Riemannian geometry by allowing a positively homogeneous norm $F(x,v)$ on tangent directions rather than requiring the square root of a quadratic form. Under suitable regularity and convexity conditions, path length is the integral of F along a curve. Polyhedral norms sit naturally at the finite or nonsmooth boundary of this broader framework. The terminology must remain precise: an L1 gauge is a legitimate norm and length structure, but it is not strongly convex or smooth on every nonzero direction in the sense used by some definitions of a classical Finsler metric. Here *polyhedral/Finsler* denotes a finite normed causal renderer and its length functional; any smooth Finsler or pseudo-Finsler completion requires additional work (Javaloyes & Sánchez, 2020).

Lorentzian geometry differs again. Its quadratic form is indefinite, causal classification depends on the sign of that form, and the associated spacetime distance is not a positive symmetric metric on events. A finite positive metric can describe separation on a rendered patch without being a Lorentzian spacetime metric. A causal influence order can coexist with that positive metric, but their compatibility and continuum combination must be shown. In quantum lattice systems, Lieb and Robinson (1972) provide the conventional model for making locality quantitative: a finite-range interaction gives a velocity-dependent bound on distant commutators rather than an automatic relativistic light cone. The discrete support theorem below is sharper only because it assumes a strictly finite one-edge-per-update rule; the continuous spectral lane returns to a bounded-tail setting.

2.4 Identification versus representation

A representation problem asks whether data can be encoded by at least one geometric object. An identification problem asks whether the data select one object, possibly up to a declared equivalence such as relabelling, isometry, gauge, or coordinate change. The distinction is central to inverse problems. A solver returning one admissible matrix establishes existence only if it has not also proved that all admissible solutions are equivalent.

ChronaQ uses the same distinction at two levels. Observable data can admit several latent state completions, and one rendered representative can admit several geometry constructors. Even when the latent representative is unique, the geometry need not be. Conversely, a carefully designed operational test might identify a geometry even when some latent details remain underdetermined. CQ2 therefore returns a family and records the quotient under the declared equivalence relation. Singleton language is permitted only when that quotient contains exactly one class.

3. ChronaQ axioms and the P2 premise map

3.1 A1: timeless totality

A1. The complete description is timeless: the total quantum object has no distinguished external time parameter. P2 inherits from P1 an internally ordered sequence of rendered records. The index attached to that sequence is not assumed to be a background temporal coordinate. A geometry constructor may use the order to compare successive patches or to orient influence, but it may not convert the index into physical duration without an independent clock calibration.

A1 matters because a geometry assembled from correlations is relational from the outset. There is no pre-existing spacetime in which the observables are placed and then rediscovered. At the same time, A1 does not forbid convenient charts. Coordinates are permitted as representations if active changes of the rendered geometry are distinguished from passive coordinate relabellings.

3.2 A2: two-boundary admissibility

A2. Complete histories are admissible only when they satisfy both boundary conditions. P2 does not repeat the finite survivor-lineage proof. It takes a nonempty rendered patch whose records have already passed the P1 admissibility and continuation gates. Geometry is constructed on those rendered objects or on a declared compatible family associated with them.

The boundary conditions can affect which histories and correlations survive and therefore can affect the geometry-compatible family. This dependence is not retrocausal signalling. It is a dependence of the globally conditioned model. If a boundary choice makes the rendered patch empty, nonpositive, or nondecoherent, P2's positive geometry constructors do not apply.

3.3 A3: positive observable rendering

A3. Observable alternatives are rendered through positive, normalized probabilities when their interference is absent or controlled and stable records distinguish them. A3 licenses the probability vectors and mutual informations used below. It does not require every latent ChronaQ object to be positive. Signed amplitudes, phases, or other latent structures must first be mapped to an observable positive chart before Hellinger distance, Shannon entropy, or classical mutual information is evaluated.

This separation prevents an authority inversion. A useful positive metric on rendered probabilities does not erase the latent quantum description, and a latent overlap relation is not automatically rendered spatial proximity. They are different mathematical types connected only by a declared export map.

3.4 A4: entanglement-geometry compatibility

A4. Relational quantum structure constrains a family of compatible geometric descriptions. Entanglement or other relational data may provide inputs to a geometry interface, but a metric, signature, scale, and continuum limit require their own admissibility conditions. A4 is not the equation $d=1/I$, nor does it say that every allowed constructor is physically equivalent.

A4 has three consequences for method. First, geometry candidates must be functions of named relational inputs rather than arbitrary matrices introduced after inspection. Second, every candidate must pass the axioms of the geometric type claimed. Third, failure of uniqueness is a result to be retained. The family may narrow as later evidence is added.

3.5 Premise and output map

Table 1 records the role of each axiom and prevents later conclusions from being read back into P2.

Table 1

ChronaQ axioms and their role in Foundations II

Premise	P2 input	P2 use	Not supplied by the premise

The explanatory flow is summarized in Figure 1.

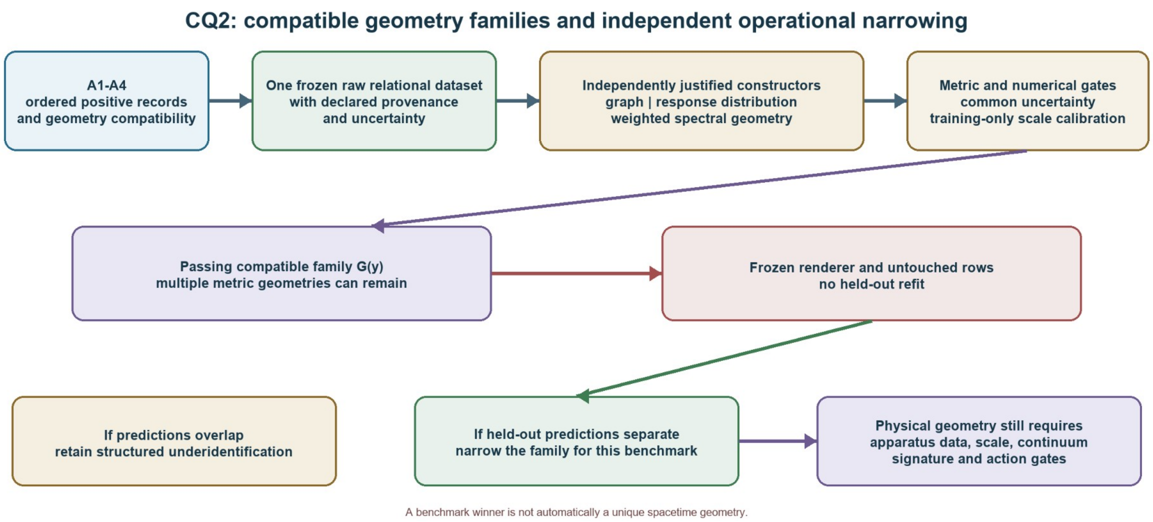

Figure 1. CQ2 compatible-family and operational-narrowing pipeline. One frozen relational dataset feeds independently justified constructors. Metric gates, shared uncertainty, and training-only calibration return a compatible family; a frozen renderer and untouched rows can then narrow that family for the benchmark. Physical geometry still requires independent apparatus, scale, continuum, signature, and action gates.

4. Finite mathematical setting

4.1 Rendered patches and observable descriptors

Let $X=\{x_1,\ldots,x_n\}$ be a finite set of rendered observable objects. An object may be a stable record, an observable context, a region, or another operationally specified unit. The rendered data y may contain several typed components: positive outcome distributions p_i, pairwise correlation summaries, an undirected adjacency graph, calibrated chart coordinates q_i, or response data obtained from allowed interventions. A component is absent unless it is actually supplied.

The term *patch* denotes the pair P=(X,y) together with its provenance and admissibility conditions. It does not imply a continuum neighbourhood. Two objects may share one observable descriptor and differ in another. A metric constructor must say which descriptor it reads.

For positive joint variables X_i and X_j, the classical mutual information in nats is

$$I(X_i:X_j)=\sum_{a,b} p_{ij}(a,b)\ln\frac{p_{ij}(a,b)}{p_i(a)\,p_j(b)} \tag{1}$$

where the sum is over the positive support and terms with zero joint probability are interpreted by continuity. Mutual information is an observable dependence measure in this classical rendered chart. Quantum mutual information could be used in another declared chart, but the metric question would still require separate proof.

4.2 Compatible state families

Finite observations may not identify a unique underlying state or completion. Let C(y) denote the set of state, model, or latent representatives that reproduce y within declared tolerances and satisfy the relevant positivity, normalization, boundary, and operational constraints. In a fully observed classical fixture C(y) may contain one representative. In a partial tomography problem it may be a convex or more general feasible set.

P2 does not select an element of C(y) silently. A constructor that depends on a representative rho must either return a geometry for every admissible rho, prove representative independence, or state a selector. This convention is inherited by the later finite gluing paper.

4.3 Geometry constructors and gates

Let K be a preregistered collection of geometry constructors. A constructor κ takes the permitted portion of (y,ρ), together with declared hyperparameters, to a candidate geometric object. For a finite positive metric candidate d, the acceptance gate is

$$d_{ij} \in [0,\infty), d_{ij} = 0 \Leftrightarrow i = j, d_{ij} = d_{ji}, d_{ik} \leq d_{ij} + d_{jk} \tag{2}$$

for all i,j,k, with a declared numerical tolerance when calculations are approximate. Additional gates may test invariance under relabelling, stability to observation noise, refinement behaviour, and independence from forbidden targets. Passing the metric gate says exactly that d is a metric on X. It does not identify its physical meaning.

Let approximately-equal denote the declared equivalence relation on candidates. For finite distance matrices this may be equality after an allowed relabelling or an isometry-preserving coordinate transformation. The compatible geometry family is the quotient

$$G(y) = \{d_{\kappa,\rho} : \rho \in C(y), \kappa \in K, A_\kappa(y,\rho) = \text{pass}\} / \sim \tag{3}$$

where Aκ records the domain and all gates for constructor κ. Returning equivalence classes prevents the same geometry written in different passive coordinates from being counted several times.

4.4 The finite non-identification proposition

Proposition 1 (finite compatible-family logic). Suppose P=(X,y) is fixed, K and all constructor parameters are declared before target comparison, and each retained candidate passes its stated domain, metric, covariance, and numerical gates. Then G(y) is the complete compatible output relative to those inputs. If G(y) contains two inequivalent elements, the premises do not entail a unique geometry. If G(y) contains one equivalence class and the constructor family is demonstrably exhaustive for the physical hypothesis being tested, the geometry is identified relative to that hypothesis.

The proof is logical. A conclusion is entailed only if it holds in every model satisfying the premises. Two inequivalent passing candidates are two countermodels to uniqueness. A singleton among a narrow hand-picked constructor list is not enough unless the list is exhaustive for the claim. This last condition matters because an optimizer can produce a unique minimizer inside one parameterization while an untested geometry exists outside it.

4.5 Forced, chosen, conditional, and conjectural components

The compatible-family formalism exposes the status of each step.

- Choosing X, the observable descriptors, a constructor family, and tolerances is a modelling act.
- Excluding a candidate that fails a claimed axiom is forced under the finite audit.
- Retaining a candidate that passes its gates is conditional evidence for that representation.
- Declaring one retained geometry to be physical is conjectural until a selection gate or uniqueness theorem closes.

This classification does not weaken A4. It states what the axiom does: it turns relational data into constraints on a family. A stronger output is allowed when stronger inputs are proved.

5. CQ2 and the reciprocal-mutual-information exclusion

5.1 Why inverse dependence is attractive

If strong entanglement suggests proximity, it is natural to seek a decreasing transform of an information quantity. Define the literal reciprocal score on positive mutual information by

$$r_I(U,V) = \frac{1}{I(U:V)}, I(U:V) > 0 \tag{4}$$

Large dependence then gives a smaller score and weak dependence gives a larger one. The score is symmetric, so it can be a useful heuristic weight in a restricted constructor.

A metric, however, requires more. For any nondegenerate variable U, I(U:U)=H(U), and therefore $r_I(U,U)=1/H(U)>0$. The literal rule already fails the zero-diagonal axiom. Independence gives I=0 and a singular off-diagonal reciprocal. Distinct but perfectly correlated variables can create other identification problems depending on the object definition. Most importantly, even manually setting the diagonal to zero does not make a decreasing transform of a symmetric dependence measure inherit the triangle inequality.

5.2 Exact counterexample

Let A and C be unbiased binary variables with joint law

$$P(A,C)=\begin{bmatrix}0.275 & 0.225\\ 0.225 & 0.275\end{bmatrix} \tag{5}$$

and define B=(A,C). Both marginals are uniform. Because B contains A and C, I(A:B)=H(A)=ln 2 and I(B:C)=H(C)=ln 2. Direct evaluation gives I(A:C)=0.0050083668463568876 nats. The literal reciprocal score has diagonal values $r_I(A,A)=r_I(C,C)=1/\ln 2=1.4426950408889634$ and $r_I(B,B)=1/H(B)=0.7239630345531265$, so it is not a metric. Its off-diagonal values are

$$d_I(A,B)=d_I(B,C)=1.4426950408889634,\, d_I(A,C)=199.66588524309182 \tag{6}$$

To isolate the triangle question from the diagonal defect, define an explicitly repaired extended dissimilarity

$$\tilde{d}_I(U,V)=\begin{cases}0 & \text{if } U=V\\ \dfrac{1}{I(U:V)} & \text{if } U\neq V, I>0\\ +\infty & \text{if } U\neq V, I=0\end{cases} \tag{7}$$

The triangle excess along A-B-C is 196.7804951613139: the repaired A-C value exceeds the sum of the repaired A-B and B-C values. By moving the A-C law closer to independence while keeping a positive correlation, I(A:C) can be made arbitrarily small and the violation arbitrarily large. At exact independence the off-diagonal value is infinity, so the corrected object is extended rather than a finite metric.

Proposition 2 (reciprocal-MI exclusion). The literal reciprocal rule is not a general finite metric because it fails the zero-diagonal axiom. The explicitly diagonal-corrected extended dissimilarity in Equation 7 is also not a general metric because a valid positive law violates its triangle inequality.

The example proves both parts of the proposition. It is not an objection to mutual information itself. Mutual information can still weight graph edges, enter a likelihood, constrain a compatible family, or support a restricted transform whose metric properties are proved on a specified domain. Variation of information, for example, is a genuine metric on partitions of a common underlying set (Meilă, 2007). What is excluded is universal metric status for the reciprocal transform on arbitrary random variables or rendered ChronaQ objects.

5.3 CQ2: compatible entanglement geometry

CQ2 is the following layered statement.

CQ2 structural layer. A finite A3-rendered patch and its relational descriptors constrain the family G(y) of declared geometry candidates that pass their type-specific gates.

CQ2 identification layer. Geometry is unique only when the surviving family contains one equivalence class relative to an exhaustive declared hypothesis. Multiple surviving classes are a positive underidentification result.

CQ2 causal export. A retained geometry may support a causal renderer only after an intervention or propagation rule is supplied and its influence gates pass. An exact finite renderer need not be Lorentzian.

CQ2 physical export. Spacetime signature, dimension, continuum topology, physical length, limiting speed, entropy-area normalization, and gravitational dynamics are downstream properties. They are not inherited from the word *entanglement* or from satisfaction of the positive metric axioms.

CQ2 therefore formalizes A4 without overselecting. It is compatible with a future theorem that makes G(y) a singleton. It is also compatible with the possibility that different physical regimes select different effective geometries.

6. A four-object compatible-geometry model

6.1 Heterogeneous descriptor fixture

Let X={A,B,C,D}. Three descriptor types are declared on the same labels. The coordinate chart, probability rows, and graph are separate inputs to one heterogeneous rendered patch; none is derived from either of the others in this fixture.

The coordinate chart is A=(0,0), B=(1,0), C=(1,1), and D=(0,1). The positive three-outcome distributions are

$$\begin{pmatrix} p_A \\ p_B \\ p_C \\ p_D \end{pmatrix} = \begin{bmatrix} 0.70 & 0.20 & 0.10 \\ 0.20 & 0.70 & 0.10 \\ 0.10 & 0.20 & 0.70 \\ 0.20 & 0.10 & 0.70 \end{bmatrix} \quad (8)$$

The undirected graph has edge lengths AB=0.8, BC=1.1, CD=0.9, and DA=1.3. No diagonal edge is supplied. These inputs are deliberately modest: they are sufficient to construct and compare finite metrics without importing a target spacetime. The example therefore tests typed compatibility and non-identification for a supplied descriptor package. It does not show that one fixed entanglement dataset, processed by three independently justified rules, yields three physically competitive geometries.

6.2 Graph shortest-path geometry

For positive edge costs w(e), define

$$d_G(x_i, x_j) = \min_{P: i \to j} \sum_{e \in P} w(e) \quad (9)$$

where the minimum is over graph paths from x_i to x_j. Connectivity and positive edge weights give finiteness and separation; path reversal gives symmetry; concatenation gives the triangle inequality. For the fixture, the distance matrix in A,B,C,D order is

	A	**B**	**C**	**D**
A	0.0	0.8	1.9	1.3
B	0.8	0.0	1.1	2.0
C	1.9	1.1	0.0	0.9
D	1.3	2.0	0.9	0.0

This geometry represents declared connectivity and traversal cost. It does not read the probability vectors.

6.3 Hellinger observable geometry

For positive normalized distributions p and q, define the Hellinger distance (Hellinger, 1909)

$$d_H(p, q) = \frac{1}{\sqrt{2}} \left\| \sqrt{p} - \sqrt{q} \right\|_2 \quad (10)$$

The square-root map embeds the probability simplex in the positive orthant of the Euclidean unit sphere. The Euclidean triangle inequality therefore proves that d_H is a metric. The conventional factor 1/sqrt(2) bounds the distance by one.

The fixture gives

	A	**B**	**C**	**D**
A	0.000000	0.389446	0.520432	0.468869
B	0.389446	0.000000	0.468869	0.520432
C	0.520432	0.468869	0.000000	0.130986

	A	B	C	D
D	0.468869	0.520432	0.130986	0.000000

This geometry represents distinguishability between rendered probability laws. It is insensitive to the supplied graph unless the graph was itself derived from those laws by a declared rule.

6.4 L1 polyhedral chart geometry

In the coordinate chart, define

$$F(v)=\|v\|_1=\sum_r |v_r|, d_1(q_i,q_j)=F(q_i-q_j) \tag{11}$$

The associated distance matrix is

	A	B	C	D
A	0.0	1.0	2.0	1.0
B	1.0	0.0	1.0	2.0
C	2.0	1.0	0.0	1.0
D	1.0	2.0	1.0	0.0

The metric axioms follow from the norm axioms. In two dimensions, the radius-one set for the norm in Equation 11 is a diamond. This geometry represents a chosen chart and cost norm. The chart may ultimately be derived from tomography, but the present metric remains conditional on that export.

6.5 Metric audit and non-identification

All three matrices are finite, nonnegative, symmetric, zero only on the diagonal, and satisfy the triangle inequality to the calculation tolerance. They are not the same geometry. The maximum entrywise differences between the three candidate pairs are 1.4795677394827624, 0.30000000000000004, and 1.4795677394827624. No relabelling can remove their different distance spectra and pair orderings.

Proposition 1 therefore applies. The heterogeneous patch proves existence of several compatible finite geometries and excludes uniqueness under the declared descriptor package and constructors. It does not say that all imaginable geometries are equally good, nor does it establish entanglement-only underidentification. Each surviving candidate has passed a typed gate and retains a different operational interpretation. Section 7 supplies the stronger complement: it holds one raw relational dataset fixed, justifies several competing constructors independently, propagates common uncertainty, and asks whether later operational data favour one without refitting.

Table 2 summarizes the comparison.

Table 2

Three passing geometry candidates on the heterogeneous finite patch

Candidate	Reads	Why it is a metric	Physical interpretation at this stage
Graph shortest path	Positive graph edge costs	Minimum path construction	Conditional traversal/connectivity geometry
Hellinger	Positive normalized observable distributions	Euclidean square-root embedding	Observable distinguishability geometry
L1 chart	Rendered coordinates and polyhedral norm	Norm-induced distance	Conditional finite spatial or causal renderer

6.6 Noise control

The Hellinger inputs are perturbed by adding 1.0×10^{-4} to the first outcome probability and subtracting the same amount from the second outcome probability in every row. Positivity and normalization are retained. The largest distance change is 4.92×10^{-5}, and the perturbed matrix still passes every metric gate. This is a finite continuity check. It does not establish a uniform noise theorem near the boundary of the probability simplex, where square-root derivatives require greater care. Appendix B gives every perturbed probability row and the complete distance matrix.

Noise stability and uniqueness are different. A candidate can be stable under perturbation while remaining one of several stable candidates. Conversely, a unique fit at one precision can split into a family when uncertainty is represented honestly.

6.7 Refinement control

The L1 path from (0,0) to (1,1) has length 2 along a coordinate-ordered route. Splitting every straight segment into two collinear subsegments leaves the sum exactly 2; the recorded residual is zero. More generally, additivity on collinear same-direction segments preserves length under this refinement.

This control shows that the declared discrete path is not an artefact of one edge partition. It does not show convergence to a smooth Riemannian manifold. Repeated refinement of coordinate steps retains the L1 unit ball unless the stencil or cost law changes. If rotationally isotropic quadratic geometry is desired, isotropy must be supplied or derived through a richer refinement process and tested prospectively.

7. Same-data operational discrimination

7.1 Frozen dataset and prospective split

The compatible-family theorem says what the declared inputs permit; it does not imply that every surviving geometry makes the same operational predictions. To test this distinction, one synthetic coined quantum-walk event dataset was generated and frozen before the candidate comparison (Aharonov et al., 1993). The walk uses a periodic 9×9 lattice with four coin states and twelve separately prepared ensemble layers. Each of the 81 source preparations supplies a destination-count distribution at every layer, with 32,768 shots per source–layer distribution. The simulator is not built from any of the three candidate geometry constructors.

Geometry construction receives node identifiers and counts from layers one and two, but not coordinate metadata. The operational target begins at layer three: the first separately prepared ensemble layer at which a destination probability reaches 0.004. This threshold response-arrival layer is neither a repeated-measurement first-detection time nor a physical speed. Coordinate metadata are used only to freeze and audit the displacement-class split: 648 training pairs, 486 validation pairs, and 972 untouched longer-range off-axis holdout pairs. Every holdout target resolves within the twelve-layer window.

Figure 1 summarizes the prospective pipeline. One raw table feeds all constructors; common training rows set one scalar scale per candidate; shared uncertainty is propagated before untouched targets are opened.

7.2 Independently motivated constructors

All candidates read the same early count table but represent different operational hypotheses.

- **Graph shortest path:** a symmetrized layer-one transition support creates an undirected empirical graph; unit edge costs and path composition yield the graph metric.
- **Hellinger early response:** each source is represented by the equally weighted concatenation of its layer-one and layer-two destination distributions; Hellinger distance measures statistical distinguishability.
- **Spectral resistance:** symmetrized weighted layer-one response supplies a graph Laplacian; square-root effective resistance gives a spectral energy/diffusion metric (Klein & Randić, 1993).

Every candidate passes nonnegativity, separation, symmetry, and triangle gates on the frozen data. For each geometry, one scalar renderer is fitted by least squares through the origin on the common training classes and then frozen. No intercept, candidate-specific feature, validation refit, or holdout refit is permitted.

7.3 Common uncertainty and decision rule

Finite shot noise is propagated through geometry construction, scale fitting, target construction, and scoring by 256 rowwise multinomial bootstrap replicates with seed 314159. Within each replicate the same resampled count table is passed to all candidates, so pairwise intervals share the same fluctuation history. The intervals are conditional on the frozen simulator and do not include apparatus or model-form uncertainty.

The preregistered pass requires all metric and split gates, resolved holdout targets, a lower 95% bound of at least 0.25 layer for at least two pairwise prediction differences, and a lower 95% bound of at least 0.25 mean-absolute-error layer for the best candidate against each alternative. Metric passage alone cannot determine the winner.

7.4 Held-out results

Table 3 reports the no-refit point results. The graph renderer predicts the held-out mean response-arrival layer near the observed value and has substantially smaller mean absolute error than the other two constructors.

Table 3

Same-data training, validation, and held-out geometry results

Constructor	Fitted scale	Training MAE	Validation MAE	Held-out prediction	Held-out observation	Held-out MAE
Graph shortest path	1.031869	0.206854	0.526561	5.503302	5.740741	0.473965
Hellinger early response	3.400969	0.566828	1.160947	3.400969	5.740741	2.339771
Spectral resistance	1.831112	0.439395	0.972977	3.576728	5.740741	2.164013

The 95% bootstrap intervals for the held-out mean prediction are [5.457364, 5.510767] for graph shortest path, [3.369253, 3.406130] for Hellinger, and [3.541944, 3.582404] for spectral resistance. The graph–Hellinger prediction difference is [2.088120, 2.104651] layers and the graph–spectral difference is [1.915445, 1.928372] layers. Hellinger and spectral resistance differ by only [0.172707, 0.176277] layer and therefore do not clear the 0.25-layer materiality gate against each other.

The graph renderer's held-out MAE improvement is [1.861012, 1.917262] layers over Hellinger and [1.684822, 1.744569] layers over spectral resistance. Both lower bounds clear the preregistered gate. Figure 2 shows the point predictions, uncertainty intervals, and held-out error comparison.

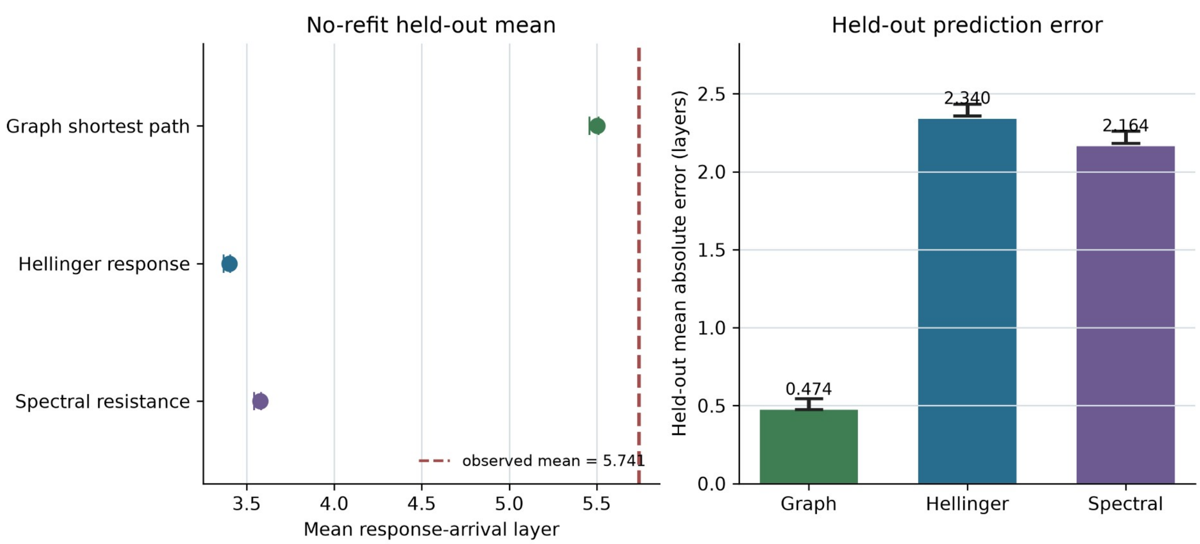

Figure 2. No-refit same-data holdout. Left: candidate mean response-arrival predictions with 95% shared-bootstrap intervals and the held-out observation. Right: held-out mean absolute errors with 95% intervals. Graph shortest path is benchmark-best for the declared renderer; this ranking is not a unique physical-geometry claim.

7.5 Consequence for CQ2

This benchmark closes the same-data obligation in a finite prospective setting. One raw count table admits three mathematically valid metric constructions, but their frozen renderers do not make equivalent untouched operational predictions. Later response data therefore narrow the compatible family: graph shortest path is the benchmark-best candidate under the declared renderer and uncertainty law.

The conclusion is not that graph distance is the unique ChronaQ geometry or physical spacetime metric. The dataset is synthetic, the target is dimensionless, and the comparison conditions on one simulator and constructor family. The result supplies an existence proof for operational discrimination within CQ2, not Lorentzian recovery, a physical length or speed calibration, or empirical confirmation of ChronaQ.

8. Finite causal structure

8.1 Five different relations

The word *causal* is especially vulnerable to premature promotion. Five relations must be distinguished.

- **Latent overlap** states that two rendered objects share a compatible completion, latent cell, or constraint relation.
- **Rendered distance** is calculated after an observable export and positive metric are supplied.
- **Controllable influence** asks whether an allowed intervention at one context can change accessible statistics at another after a stated number of record updates.
- **Propagation amplitude or response** may be nonzero outside an earliest-permitted support layer while being bounded or suppressed.
- **Physical light-cone structure** combines an operational causal order with physical units, continuum localization, and a Lorentzian or other justified spacetime geometry.

Entanglement correlation alone is not controllable influence. A local nonselective quantum operation cannot change a remote reduced state, even though conditioning on a selected outcome can change the conditional description. ChronaQ retains this standard no-signalling boundary. Latent co-location can account for shared consistency without becoming a superluminal message.

8.2 Exact finite support bound and earliest-permitted front

Consider a rectangular cover of finite observable contexts. Two contexts are adjacent when their overlap contains a declared nontrivial observable generator; the identity alone does not create an edge. Let one elementary record update transmit controllable influence across at most one adjacency. Induction then places the possible influence support $S_k(s)$ after k updates inside the graph ball $B_k(s)$ of radius k about a source s.

Let $W_k(s,x)$ be a separately defined nonnegative witness of controllable influence from s to x after k updates, and define $t_{arr}(s,x)$ as the first k for which $W_k(s,x)>0$, with $t_{arr}=\infty$ if no such k exists. The support theorem implies

$$S_k(s)\subseteq B_k(s)\Rightarrow t_{\text{arr}}(s,x)\geq d_G(s,x) \tag{12}$$

On a square nearest-neighbour chart this becomes

$$t_{\text{arr}}(\Delta x,\Delta y)\geq|\Delta x|+|\Delta y| \tag{13}$$

where Δx and Δy are coordinate differences in lattice units. The boundary $|\Delta x|+|\Delta y|=k$ is therefore a diamond-shaped earliest-permitted arrival layer. The theorem excludes arrival outside $B_k(s)$; it does not prove that a nonzero witness reaches every point on the boundary. Equality requires a specified update channel, a nonzero shortest-path witness, and control of destructive cancellation. Those ingredients have not been supplied here.

A disconnected control remains outside every reachable component; an all-to-all update violates the one-edge locality assumption; and a fully commuting channel may remove every nontrivial witness even inside the permitted ball. These controls distinguish a locality bound from an attained propagation law. The induction is exact for the declared discrete support rule, but it relies on observable-local composition in addition to restriction consistency. Boundary closure alone would not forbid a dynamical map from coupling disjoint contexts. Observable locality is therefore a declared physical assumption of this renderer, not a theorem from A1-A4 alone.

8.3 Norm, dual norm, and length functional

The continuum notation associated with the finite diamond-shaped support ball is

$$F(v)=|v_1|+|v_2|,F^*(p)=max(|p_1|,|p_2|) \tag{14}$$

The dual norm determines the support function and response to covectors. A dimensionless Finsler length functional can be written

$$L_F[\gamma]=\int_{\lambda_0}^{\lambda_1} F(\gamma(\lambda),\dot{\gamma}(\lambda))d\lambda \tag{15}$$

for absolutely continuous paths in the declared chart. Minimizers reproduce L1 geodesic length. Because the norm is not strictly convex, several coordinate-monotone paths can share the same length. The response can consequently be set-valued at nonsmooth directions, with Karush-Kuhn-Tucker conditions selecting an admissible face rather than a unique differentiable velocity.

This is a legitimate finite length result. It becomes an action only if a later model deliberately chooses the same functional as a variational action and supplies its degrees of freedom and endpoint conditions. It is dimensionless because no physical length or time unit has been introduced. It is polyhedral rather than quadratic. Calling it Finsler records that length depends on a norm of directions; calling it Lorentzian would be incorrect without an indefinite spacetime construction and a recovery theorem.

8.4 Continuous spectral flow and effective causal bands

A continuous generator built from graph adjacency behaves differently from a strictly local discrete update. Matrix exponentials generally contain higher powers of the adjacency operator at every positive parameter value. Consequently, influence develops nonzero tails at arbitrarily large graph distance, even when those tails satisfy a Lieb-Robinson-style suppression bound (Lieb & Robinson, 1972). The correct object is then an effective causal band, not an exact compact-support cone. P2 does not derive a Lieb-Robinson velocity; it uses that conventional result to distinguish bounded tails from the stricter support assumption in Section 8.2.

This distinction does not make the continuous model unusable. Many physical systems have finite-velocity bounds without exact zeros. It does prevent the continuous spectral flow from being used as proof of the discrete sharp boundary. ChronaQ may retain both lanes if it says which one supplies an exact support exclusion and which one supplies approximate propagation. Neither lane establishes actual first-arrival equality without an explicit nonzero witness.

8.5 Rendered distance and latent overlap are independent types

A finite countermodel can hold the latent completion graph and its pair response fixed while changing the rendered coordinate export, thereby changing rendered distance. Conversely, two pairs can be assigned the same rendered distance while retaining different latent adjacency and response. Passive coordinate transformations accompanied by the correct metric pullback leave all distances invariant, so the effect is not a coordinate-labelling artefact.

The conclusion is structural: A1-A4 do not force latent overlap to equal rendered spatial distance. A physical export law must connect them if a relation is intended. This leaves room for ChronaQ's intuition that objects separated in the rendered chart may share latent completion structure, while preserving ordinary no-signalling and avoiding an unsupported force-distance claim.

8.6 What would select a smooth or Lorentzian lane

A smooth geometry claim requires a refinement family, a topology or convergence notion, uniform controls, and a limit object. A Lorentzian claim additionally requires an indefinite signature, an operational causal order compatible with that signature, local dimension, and a mapping from the finite record coordinate to a temporal variable. A prospective test should freeze the refinement stencil and success criteria before comparing with a quadratic cone.

Several outcomes are scientifically useful. Convergence to L1 confirms a polyhedral limit. Convergence to a smooth strictly convex Finsler norm selects a nonquadratic continuum lane. Convergence to a quadratic norm may support an isotropic spatial sector. Recovery of a Lorentzian cone requires time as well as space and cannot be inferred from a positive spatial metric alone. Failure to identify among several limits should remain a family-valued result.

9. Entanglement, area, continuum geometry, and gravity

9.1 Conditional entropy-area relations

The holographic entropy-area relation is an important model for what a mature ChronaQ bridge could accomplish. It is not a generic identity for every entangled system. A ChronaQ analogue would need at least a class of rendered regions R, a boundary or extremal object Σ_R defined by the selected geometry, a measure A_g, a state-dependent entropy S_R, and a controlled correction term. A schematic conditional relation is

$$S_R = \eta\, A_g(\Sigma_R) + S_{\text{bulk}}(R) + \varepsilon_R \tag{16}$$

where η carries the required units and normalization. Every symbol depends on prior construction. If g is family-valued, so are Σ_R and A_g unless the area observable is invariant across the family. If the patch is finite and dimensionless, η cannot be identified with $1/(4G\hbar)$ without a physical calibration and a gravity derivation.

ChronaQ can test such relations without assuming them. A target-free protocol would define regions and candidate surfaces, compute both sides independently, include conventional and null controls, and state the consequence of non-agreement. Evidence that one geometry makes the relation stable under refinement could become a selector. Agreement obtained by fitting eta separately to each opened case would not identify the law.

9.2 Hessian and potential constructions

Another possible route assigns a scalar potential Φ to relational data and defines a local tensor through a Hessian,

$$g_{ab}^{\Phi} = \partial_a \partial_b \Phi \tag{17}$$

This is familiar in information geometry and convex analysis. It remains a chosen construction until the domain, differentiable structure, coordinate transformation law, convexity or signature, and observable meaning of Φ are established.

A positive Hessian can define a Riemannian information metric on a statistical manifold. It cannot by itself produce a Lorentzian signature. An indefinite Hessian requires a different stability analysis and may fail to be a distance metric. Nor may a potential-first construction override the finite non-identification result by definition. Its output joins G(y) and must be compared with the graph, Hellinger, and polyhedral lanes under the same selection discipline.

9.3 Continuum and dimension

A finite distance matrix does not uniquely determine a continuum manifold. One must specify a sequence of patches, embeddings or correspondences, convergence of distances, dimensional estimators, topology, and regularity. Different continuum spaces can approximate the same small finite metric within uncertainty. Dimension should therefore be treated as an inferred property with controls, not as the number of coordinates used for a convenient plot.

The CQ1 record coordinate creates an ordered family of patches. It may index a refinement or slice sequence, but it does not supply physical lapse or proper time. To build spacetime, P2 needs a rule that relates changes between patches to the spatial or causal geometry, and that rule must be invariant under permissible reparameterizations of the dimensionless CQ1 chart.

9.4 Boundary to the existing gravitational law

The target of the ChronaQ programme is to derive the existing Einstein or Einstein-Cartan law from its axioms and action structure, not to introduce a replacement law. P2 provides geometric candidates and the prerequisites that such a derivation must respect. It does not complete the derivation.

At minimum, the downstream bridge requires:

- a selected or physically invariant geometric structure with a justified signature and continuum interpretation;
- a common action or equivalent variational object containing all retained geometric, connection, coframe, matter, and source variables;
- the complete transformation laws and off-shell Ward identities;
- the corresponding Bianchi or integrability structure;
- a conserved observable stress-energy or Einstein-Cartan source with a proved relation to the matter sector;
- boundary terms and a well-posed variational principle;
- a physical normalization map and independent determination of the coupling identified with G; and
- a controlled limit in which the full, not merely linearized, existing field equation follows.

Entanglement-based gravity results in the literature illustrate how strong premises can close parts of this chain (Faulkner et al., 2014; Jacobson, 2016). ChronaQ must supply its own premises and derivation. A distance constructor cannot substitute for a common action, and an entropy stationarity statement cannot substitute for the missing variations and conservation identities.

9.5 Downstream paper interfaces

P3 may use the CQ2 distinction between a compatible family and a selected representative when reconstructing positive observable states from finite local data. P5.5 may ask whether compatible local state and geometry families glue globally and whether a selector is continuous, natural, and cocycle-consistent. P4 must independently build the action, variation, currents, sources, and gravitational bridge. P5 may design experimental discriminators whose renderers are sensitive to differences among the surviving geometries.

The dependency is one-way. A later action may select among P2 candidates if the selection is derived independently. P2 may then be updated with that result. It may not anticipate the answer by presenting a preferred metric as already forced.

10. Verification, controls, and falsifiers

10.1 Metric gates

Every proposed finite metric should report the full distance matrix and the maximum residual or violation for each axiom. Singular entries, zero distances between operationally distinct objects, asymmetry, and triangle excess are separate failure modes. Replacing *metric* with *similarity*, *pseudometric*, or *extended metric* is acceptable only when that type is scientifically intended and used consistently downstream.

The literal reciprocal-mutual-information score fails the zero-diagonal gate, and the explicitly diagonal-corrected dissimilarity fails the triangle gate on the same valid positive law. The graph, Hellinger, and L1 candidates pass on the heterogeneous fixture. These outcomes are not balanced by averaging: one failed universal claim is excluded, while three typed constructions survive on their declared descriptor domains.

10.2 Prospective comparison

A geometry comparison should not use the desired spacetime answer to choose constructors or parameters and then cite recovery of that answer as evidence. Inputs, candidate lanes, tolerances, and equivalences should be frozen before the discriminating values are opened. Conventional baselines and deliberately wrong controls help establish that the test can reject something.

When several candidates pass, the result is underidentification relative to the declared data bundle and constructor set. Section 7 demonstrates the corresponding prospective narrowing procedure: several constructors read one frozen raw event table, one common uncertainty law is propagated, and an operational renderer is frozen before holdout opening. That benchmark favours graph shortest path over the two alternatives for the declared target, but the general discipline remains the same. New observables should be designed to separate surviving candidates without choosing the desired spacetime answer in advance. Independent first-arrival data, refinement anisotropy, and coordinate-covariance probes remain useful examples.

10.3 Noise and missing data

Finite observational uncertainty should be propagated to a set or distribution of distance matrices. A positive metric at the central estimate may cease to separate points within uncertainty. Hellinger geometry is continuous on the probability simplex, but inference near zero-probability faces requires explicit treatment. Graph connectivity may change discontinuously when edges are thresholded. A robust claim should therefore report sensitivity to thresholds and explain whether topology changes are physical or procedural.

Missing correlations cannot be silently set to zero. Zero dependence is an empirical statement; missing dependence is absence of information. The compatible-family construction naturally retains alternatives until data or assumptions remove them.

10.4 Refinement and convergence falsifiers

Refinement is not automatically improvement. A sequence can preserve selected path lengths while changing curvature surrogates, dimension estimates, or causal anisotropy. A continuum claim should declare which quantities converge and which remain bounded. If different refinement routes yield inequivalent limits, unique continuum language fails.

For the present L1 lane, exact preservation of a split coordinate path is a passing local control. A claim of Euclidean or Lorentzian recovery would fail unless a richer stencil causes the relevant norms or cones to converge prospectively within the declared tolerance. A continuous generator with nonzero tails fails an exact-support claim even if it passes an exponential-bound claim.

10.5 Physical falsifiers

The finite CQ2 theorem would be misapplied if any of the following were claimed without added evidence:

- reciprocal mutual information is a universal distance;
- the compatible family is a unique physical spacetime;
- a positive spatial metric is a Lorentzian metric;
- latent overlap is observable spatial coincidence;
- entanglement correlation transmits controllable influence;
- a dimensionless graph velocity is the measured speed of light;
- an entropy-area fit determines G; or
- an information metric alone yields the Einstein equation.

These are not permanent prohibitions. Each is a gate. A later theorem or preregistered experiment may close one, at which point the claim can be promoted with its assumptions visible.

11. Discussion

11.1 What has been derived

The first result is structured underidentification. Finite rendered descriptors can support real geometric conclusions without selecting one geometry. The reciprocal-MI counterexample removes a tempting but invalid shortcut. The heterogeneous fixture then demonstrates that the positive result is not emptiness: three explicit metrics pass on their respective descriptor domains, have different operational meanings, and remain distinct. The same-data benchmark supplies the next step on one fixed raw quantum-walk count dataset: all three candidate metrics pass, but their frozen operational predictions separate on untouched rows, and graph shortest path has the lowest held-out error by the preregistered margin. The finite causal calculation adds a separate exact L1 support exclusion and earliest-permitted arrival layer for one declared local update rule; it does not prove arrival equality.

This is a stronger foundation for later work than an unsupported unique metric. It says exactly what a new observation must accomplish. The Section 7 renderer is one successful finite discriminator, not a universal selector. A physical selector must additionally remain stable under apparatus uncertainty and refinement, respect passive covariance, and avoid using the target answer as input. Failure can still have a constructive interpretation: ChronaQ may remain family-valued at a layer if observables are invariant across the surviving family.

11.2 Geometry as a constrained export

The distinction between latent and rendered structure is central to ChronaQ. A latent completion can encode relations that do not look local in a selected observable chart. That possibility is often described informally as co-location beneath spacetime. The finite result supports only the logical independence: latent adjacency and rendered distance need not be functions of one another. It does not establish a hidden spatial dimension, a nonlocal force, or a measurable anomaly.

Rendered geometry is active when the export map changes while the metric convention is held fixed. It is passive when coordinates change and the metric is transformed so that distances remain invariant. Confusing these two cases can manufacture apparent geometry change. Any future tomography-to-spacetime map must state its covariance law.

11.3 Relation to holographic and entanglement-equilibrium results

ChronaQ shares with holographic and entanglement-equilibrium approaches the idea that quantum relational structure can constrain geometry. Its present finite result is more modest in physical scope. It assumes neither an anti-de Sitter boundary nor a conformal field theory, and it has not derived an extremal-area law or local Lorentzian vacuum. Therefore it should not claim the conclusions obtained in those settings.

The comparison is nevertheless useful. Holography shows that a carefully specified entanglement-geometry dictionary can be powerful. Entanglement-first-law derivations show that gravitational equations can emerge when the geometric, modular, state, and variational ingredients align. Those examples define standards for the downstream ChronaQ bridge: the dictionary, state class, locality structure, variation, and normalization must all be explicit.

11.4 Limitations

The construction is finite and dimensionless. The four-object model remains a heterogeneous-descriptor consistency and non-identification example rather than a reconstruction of observed space. The same-data comparison uses synthetic coined-quantum-walk counts, not apparatus data or a generic entanglement dataset; its bootstrap represents finite shot noise conditional on the frozen simulator and excludes model-form uncertainty. Its candidate ranking is tied to the declared threshold-response renderer. The earlier noise control is one directional perturbation, and the refinement control addresses one class of path splits. The exact causal result is a support bound under a discrete observable-local update assumption; attainability remains open, while the continuous spectral lane has nonzero tails. No unique smooth limit is selected.

The paper does not derive spatial dimension, topology, curvature, Lorentzian signature, physical distance, proper time, the speed of light, an entropy-area coefficient, stress-energy, Newton's constant, or the Einstein or Einstein-Cartan equation. These limitations are concentrated here because they are interfaces for the next papers, not because the finite CQ2 result is tentative within its domain.

11.5 Reproducibility and accessibility

The counterexample, all three heterogeneous distance matrices, metric audits, perturbation control, and path-refinement control are deterministic finite calculations. The same-data supplement additionally contains the frozen event-count table, constructor and renderer code, training/validation/holdout split, common-bootstrap seed and replicates, machine-readable results, figures, and SHA-256 manifest. The complete core values appear in Appendix B, and the accompanying reproducibility archive permits direct replay. Reproducibility materials are available from the author on reasonable request.

The figures use shape, labels, and colour redundantly and have text alternatives in the editable document. Tables repeat their header rows and do not encode meaning by colour. Display equations are editable mathematical objects rather than images.

12. Conclusion

ChronaQ Foundations II turns the claim that entanglement constrains geometry into a finite, auditable construction. A1-A3 provide the timeless, two-boundary, positive-rendered setting inherited from P1. A4 supplies the new premise: relational quantum structure constrains a family of compatible geometries. CQ2 returns that family after declared constructors and geometric gates are applied.

Reciprocal mutual information is excluded as a general metric twice over: the literal score fails the zero-diagonal axiom, and an explicit diagonal correction still fails the triangle inequality. Graph shortest-path, Hellinger, and L1 chart geometries all pass on one heterogeneous finite patch while remaining inequivalent. On a separate frozen quantum-walk count dataset, graph shortest-path, Hellinger early-response, and spectral-resistance metrics also pass together, but their no-refit held-out response predictions do not. The graph renderer's held-out error advantage survives common count uncertainty and clears the prospective selection gate against both alternatives. One separate local influence assumption selects an exact polyhedral support bound and earliest-permitted arrival layer together with a dimensionless Finsler length functional; a continuous spectral flow instead supplies an effective causal band.

The same-data result shows how operational evidence can narrow a compatible family, but its benchmark winner is not automatically a physical spacetime geometry. Apparatus evidence, alternative renderers, continuum behaviour, Lorentzian signature, entropy-area, action, identity, source, normalization, and controlled-limit gates must still close before ChronaQ can claim a derivation of the existing gravitational law. P2 supplies the finite geometry foundation for that attempt without assuming its conclusion.

Appendix A. Proofs and finite propositions

A.1 Graph shortest-path metric

Let G=(X,E,w) be a connected undirected finite graph with w(e)>0. Define $d_G(x_i,x_j)$ as the minimum total weight among paths joining x_i and x_j. Finiteness follows from connectivity and finiteness of the graph. Every path weight is nonnegative, and a zero path occurs only when i=j because every nontrivial path contains a positive edge. Reversal preserves path weight, giving symmetry. If P_ij and P_jk are minimizing paths, their concatenation is a path from i to k with weight $d_G(i,j)+d_G(j,k)$. The minimum over all i-k paths can be no larger, so $d_G(i,k)<=d_G(i,j)+d_G(j,k)$.

A.2 Hellinger metric

For normalized probability vectors p and q, map p to √p componentwise. These images lie on the unit sphere in Euclidean space. The Hellinger distance is 1/√2 times the Euclidean chord distance between the images. Nonnegativity and symmetry follow from the Euclidean norm. Equality holds only when every square root, and therefore every probability, is equal. The Euclidean triangle inequality is preserved by multiplication by the positive constant 1/√2.

Strict positivity is convenient for the finite perturbation used here but is not necessary for the basic metric on the closed probability simplex. When statistical inference rather than exact distributions supplies the vectors, boundary uncertainty must be propagated separately.

A.3 L1 metric and refinement

For q_i in R^m, $d_1(q_i,q_j)=\|q_i-q_j\|_1$. The norm axioms prove the metric axioms. For a coordinate-monotone segment split at $q_m$id, every component displacement keeps the same sign, so

$$\|q_j - q_i\|_1 = \|q_{\text{mid}} - q_i\|_1 + \|q_j - q_{\text{mid}}\|_1 \tag{18}$$

Repeated collinear or coordinate-monotone subdivision therefore preserves total length. A detour or reversal can increase the length. The result is exact for the declared refinement and does not assert path independence for arbitrary curves.

A.4 Reciprocal-mutual-information counterexample

For the joint law in Equation 4, $p_A=p_C=(1/2,1/2)$. The mutual information is

$$I(A:C) = 2(0.275)\ln(1.1) + 2(0.225)\ln(0.9) = 0.0050083668463568876 \tag{19}$$

Because B=(A,C), H(A|B)=H(C|B)=0, and hence I(A:B)=I(B:C)=ln 2. The literal diagonal values are 1/H(A), 1/H(B), and 1/H(C), all strictly positive. It therefore fails the zero-diagonal axiom. Setting those diagonal values to zero by definition gives the off-diagonal values in Equation 5. Their triangle excess is

$$199.66588524309182 - 2(1.4426950408889634) = 196.7804951613139 > 0 \tag{20}$$

One valid counterexample is sufficient to exclude a universal metric claim. The two failures are logically distinct: the literal reciprocal score is not diagonal-zero, and the diagonal-corrected extended dissimilarity is not triangular.

A.5 Compatible-family non-identification

Let P denote all declared premises and gates. If d_a and d_b both satisfy P but are inequivalent, suppose for contradiction that P entails a unique geometry class [d]. Then every model satisfying P must have class [d]. Since d_a and d_b each satisfy P, $[d_a]=[d]=[d_b]$, contradicting their inequivalence. Therefore P does not entail uniqueness. The proposition is relative to the constructor and model domain. Enlarging the premise set may remove a candidate; enlarging the constructor domain may add one.

A.6 Finite influence induction

Let the influence-support set at update zero contain only the source. Suppose the support after k updates is contained in the radius-k graph ball. An elementary observable-local update acts only within a context or across one declared overlap edge, so the next support lies in the previous support and its neighbours, hence in the radius-(k+1) ball. Induction gives the support inclusion stated in Equation 12 for every nonnegative update count. If a positive witness implies membership in that support, first arrival cannot precede graph distance. On a nearest-neighbour square chart, this is the L1 coordinate bound in Equation 13. Attainability at the boundary requires a separate nonzero witness and cancellation control; it does not follow from upper-support induction alone.

Appendix B. Complete finite calculation record

B.1 Reciprocal-mutual-information audit

Pair	Mutual information (nats)	Literal reciprocal score
A-A	0.6931471805599453	1.4426950408889634
B-B	1.3812859942735338	0.7239630345531265
C-C	0.6931471805599453	1.4426950408889634
A-B	0.6931471805599453	1.4426950408889634
B-C	0.6931471805599453	1.4426950408889634
A-C	0.0050083668463568876	199.66588524309182

The literal score is finite on this correlated example, nonnegative, symmetric, and separating off the diagonal, but it is positive rather than zero on the diagonal. After the diagonal is manually set to zero, the resulting extended dissimilarity still fails the triangle inequality with maximum excess 196.7804951613139. For a distinct independent pair its off-diagonal value is infinite.

B.2 Candidate metric gates

Candidate	Finite	Nonnegative	Zero diagonal	Symmetric	Separates points	Triangle inequality
Graph shortest path	Pass	Pass	Pass	Pass	Pass	Pass
Hellinger	Pass	Pass	Pass	Pass	Pass	Pass
L1 chart	Pass	Pass	Pass	Pass	Pass	Pass

B.3 Candidate difference and stability record

Check	Result
Maximum graph-Hellinger matrix difference	1.4795677394827624
Maximum graph-L1 matrix difference	0.30000000000000004
Maximum Hellinger-L1 matrix difference	1.4795677394827624
Maximum probability perturbation	0.00010000000000000286
Maximum Hellinger distance drift	0.000049197258056299376
Coarse L1 path length	2.0
Refined L1 path length	2.0
Refinement residual	0.0

The numerical outputs are reported beyond their physical significance so the deterministic calculation can be checked exactly. They are not measurements and do not carry physical units.

B.4 Complete Hellinger perturbation

The perturbation adds (0.0001,-0.0001,0) to every probability row. The complete inputs are

Object	Base probabilities	Perturbed probabilities
A	(0.7000, 0.2000, 0.1000)	(0.7001, 0.1999, 0.1000)

Object	Base probabilities	Perturbed probabilities
B	(0.2000, 0.7000, 0.1000)	(0.2001, 0.6999, 0.1000)
C	(0.1000, 0.2000, 0.7000)	(0.1001, 0.1999, 0.7000)
D	(0.2000, 0.1000, 0.7000)	(0.2001, 0.0999, 0.7000)

Every row remains positive and normalized. The resulting Hellinger matrix is

	A	B	C	D
A	0.0000000000000000	0.3894464463524073	0.5203831053000461	0.4688534769238946
B	0.3894464463524073	0.0000000000000000	0.4688837658782783	0.5204814577752939
C	0.5203831053000461	0.4688837658782783	0.0000000000000000	0.1309858632227954
D	0.4688534769238946	0.5204814577752939	0.1309858632227954	0.0000000000000000

The maximum change from the base distance matrix is $4.9197258056299376 \times 10^{-5}$ for pair B-D. This is one declared directional perturbation. It is not a worst-case perturbation ball, a sampling model, a covariance analysis, or a uniform stability theorem.

Appendix C. Claim-permission summary

Table C1

Permitted and unpermitted readings of the finite P2 results

Result	Permitted wording	Additional gate required for stronger wording
Reciprocal-MI counterexample	Literal 1/I fails zero diagonal; diagonal repair still fails triangle	Restricted-domain metric proof
Three heterogeneous candidates	Supplied typed descriptors support a compatible geometry family	A common raw-data reconstruction or exhaustive singleton/selector
Same-data quantum-walk benchmark	Three valid metrics read one frozen count table and make distinct no-refit held-out predictions; graph shortest path is best for this renderer	Apparatus data, alternate-renderer stability, physical calibration, continuum and signature gates
Hellinger perturbation	Stable on the declared finite perturbation	Uniform statistical error theorem
L1 split path	Exact refinement stability for the declared path	Convergence theorem for a patch sequence
Exact graph support bound	No influence before graph distance under the declared local update assumption	Nonzero propagation witness, cancellation control, physical export, and continuum compatibility
Continuous spectral propagation	Effective bounded causal band	Exact locality mechanism for a sharp cone
Dimensionless Finsler length	Valid length functional in the declared chart	Separate variational-action choice, physical units, matter coupling, and common-action identity
Entanglement-geometry compatibility	Relational data constrain admissible geometries	Area law, Lorentzian limit, source, G, and field-equation derivation

References

Amari, S. (2016). *Information geometry and its applications*. Springer. https://doi.org/10.1007/978-4-431-55978-8

Aharonov, Y., Davidovich, L., & Zagury, N. (1993). Quantum random walks. *Physical Review A, 48*(2), 1687–1690. https://doi.org/10.1103/PhysRevA.48.1687

Braunstein, S. L., & Caves, C. M. (1994). Statistical distance and the geometry of quantum states. *Physical Review Letters, 72*(22), 3439–3443. https://doi.org/10.1103/PhysRevLett.72.3439

Cao, C., Carroll, S. M., & Michalakis, S. (2017). Space from Hilbert space: Recovering geometry from bulk entanglement. *Physical Review D, 95*(2), 024031. https://doi.org/10.1103/PhysRevD.95.024031

Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory* (2nd ed.). Wiley. https://doi.org/10.1002/047174882X

Endres, D. M., & Schindelin, J. E. (2003). A new metric for probability distributions. *IEEE Transactions on Information Theory, 49*(7), 1858–1860. https://doi.org/10.1109/TIT.2003.813506

Faulkner, T., Guica, M., Hartman, T., Myers, R. C., & Van Raamsdonk, M. (2014). Gravitation from entanglement in holographic CFTs. *Journal of High Energy Physics, 2014*, Article 51. https://doi.org/10.1007/JHEP03(2014)051

Hellinger, E. (1909). Neue Begründung der Theorie quadratischer Formen von unendlichvielen Veränderlichen. *Journal für die reine und angewandte Mathematik, 136*, 210–271. https://doi.org/10.1515/crll.1909.136.210

Jacobson, T. (2016). Entanglement equilibrium and the Einstein equation. *Physical Review Letters, 116*(20), 201101. https://doi.org/10.1103/PhysRevLett.116.201101

Javaloyes, M. A., & Sánchez, M. (2020). On the definition and examples of Finsler metrics. *Revista de la Real Academia de Ciencias Exactas, Físicas y Naturales. Serie A. Matemáticas, 114*, Article 30. https://doi.org/10.1007/s13398-019-00736-y

Klein, D. J., & Randić, M. (1993). Resistance distance. *Journal of Mathematical Chemistry, 12*, 81–95. https://doi.org/10.1007/BF01164627

Lieb, E. H., & Robinson, D. W. (1972). The finite group velocity of quantum spin systems. *Communications in Mathematical Physics, 28*(3), 251–257. https://doi.org/10.1007/BF01645779

McIntyre, S. J. (2026). *ChronaQ Foundations I: Finite survivor-lineage entropy and relational order in two-boundary quantum histories* [Preprint]. McIntyre Lab.

Meilă, M. (2007). Comparing clusterings—An information based distance. *Journal of Multivariate Analysis, 98*(5), 873–895. https://doi.org/10.1016/j.jmva.2006.11.013

Ryu, S., & Takayanagi, T. (2006). Holographic derivation of entanglement entropy from AdS/CFT. *Physical Review Letters, 96*(18), 181602. https://doi.org/10.1103/PhysRevLett.96.181602

Swingle, B. (2012). Entanglement renormalization and holography. *Physical Review D, 86*(6), 065007. https://doi.org/10.1103/PhysRevD.86.065007

Van Raamsdonk, M. (2010). Building up spacetime with quantum entanglement. *General Relativity and Gravitation, 42*(10), 2323–2329. https://doi.org/10.1007/s10714-010-1034-0

Wolf, M. M., Verstraete, F., Hastings, M. B., & Cirac, J. I. (2008). Area laws in quantum systems: Mutual information and correlations. *Physical Review Letters, 100*(7), 070502. https://doi.org/10.1103/PhysRevLett.100.070502