

CAUSAL CONVERGENCE CONDITIONS THROUGH VARIABLE TIMELIKE RICCI CURVATURE BOUNDS

MATHIAS BRAUN AND ROBERT J. MCCANN

ABSTRACT. We describe a nonsmooth notion of globally hyperbolic, regular length metric spacetimes (M, l) . It is based on ideas of Kunzinger–Sämman, but does not require Lipschitz continuity of causal curves. We study geodesics on M and the space of probability measures over M in detail.

Furthermore, for such a spacetime endowed with a reference measure \mathfrak{m} , a lower semicontinuous function $k: M \rightarrow \mathbf{R}$, and constants $0 < p < 1$ and $N \geq 1$, we introduce and study the entropic timelike curvature dimension condition $\text{TCD}_p^e(k, N)$ with variable Ricci curvature bound k . This provides a unified synthetic approach to general relativistic energy conditions, including

- the Hawking–Penrose strong energy condition $\text{Ric} \geq 0$, or more generally $\text{Ric} \geq K$ for constant $K \in \mathbf{R}$, in all timelike directions,
- the weak energy condition $\text{Ric} \geq \text{scal} - \Lambda$ in all timelike directions, and
- the null energy condition $\text{Ric} \geq 0$ in all null directions.

Our approach also allows for the synthetic quantification of asymptotic conditions or integral controls on the timelike Ricci curvature. For example, we give a nonsmooth generalization of a timelike diameter estimate of Frankel–Galloway (and Schneider), and of a Hawking-type singularity theorem which requires only that the negative Ricci curvature have small enough integral relative to the maximal mean curvature of an achronal slice.

As further applications, we discuss the stability of our notion and provide timelike geometric inequalities. To obtain sharp constants in the latter, we develop the localization paradigm in the variable k framework.

CONTENTS

1.	Introduction	2
2.	Metric spacetimes	9
2.1.	Basic notions	9
2.2.	Length metric spacetimes	11
2.3.	Causality theory for length metric spacetimes	13
2.4.	Geodesy	19
2.5.	Regularity	21
2.6.	Lorentzian optimal transport	25
2.7.	Geodesics of probability measures	29
2.8.	Taking the reference measure \mathfrak{m} into account	36
3.	Timelike curvature-dimension condition	37
3.1.	Distortion coefficients along l -geodesics	37
3.2.	Distortion coefficients along plans	41

2020 *Mathematics Subject Classification*. Primary 51K10; Secondary 28A50, 49Q22, 51F99, 53Z99, 83C75.

MB’s research was supported in part by the Fields Institute for Research in Mathematical Sciences. RM’s research is supported in part by the Canada Research Chairs program CRC-2020-00289, the Simons Foundation, Natural Sciences and Engineering Research Council of Canada Discovery Grant RGPIN-2020-04162. We thank Fabio Cavalletti and Andrea Mondino for communicating their constant k versions of Theorem 6.37 and Theorem 7.2 to us, and for pointing out their related work [46].

©April 4, 2024 by the authors.

3.3.	Definition and basic properties	42
3.4.	Qualitative properties from local(ly uniform) curvature bounds	46
3.5.	Pathwise characterization	50
3.6.	Local-to-global property	51
4.	Examples	53
4.1.	Timelike convergence condition and stability	53
4.2.	Regular solutions to the Einstein equations	56
4.3.	Null convergence condition	57
5.	Geometric inequalities	58
5.1.	Timelike Brunn–Minkowski inequality	58
5.2.	Timelike Bonnet–Myers inequality	60
5.3.	Timelike Schneider inequality	60
5.4.	Timelike Bishop–Gromov inequality	62
6.	One-dimensional localization/needle decomposition	66
6.1.	Framework	66
6.2.	Transport relations and transport sets	67
6.3.	Ray map and quotient map	72
6.4.	Negligibility of bad points	74
6.5.	Disintegration of \mathfrak{m} on the transport set	79
6.6.	Localization of $\text{TCD}_p^e(k, N)$	81
7.	Applications	85
7.1.	Localization by optimal transport	85
7.2.	Sharp geometric inequalities	87
7.3.	Hawking-type singularity theorem and integral curvature bounds	90
Appendix A.	Variable timelike measure contraction property	94
Appendix B.	Limit curves; lifting causality from events to measures	98
Appendix C.	Kantorovich duality formula for ℓ_1	103
References		104

1. INTRODUCTION

Outline. This work contributes to the recent fascinating developments in mathematical general relativity from the angle of metric measure geometry [71, 129]. Such a nonsmooth approach to gravity seems highly desirable, as expressed e.g. in the introductions of [44, 87] and acknowledged in various special issues [45, 119]. We consider a lower semicontinuous function $k: M \rightarrow \mathbf{R}$ on an abstract space M to be specified and a “dimensional” parameter $N \geq 1$. By using tools from optimal transport in Lorentzian signature [44, 51, 78, 91, 99, 125], we introduce and study the (*entropic*) *timelike curvature-dimension condition* $\text{TCD}_p^e(k, N)$ for M , where $0 < p < 1$. This synthetically quantifies the “smooth” conditions

- $\text{Ric} \geq k$ in all timelike directions, and
- $\dim M \leq N$.

We also survey its dimensionless counterpart $\text{TCD}_p(k, \infty)$. To reduce the level of technicality, we leave the precise definition of these conditions as a black box for now, and refer to the “Motivation of Definition 3.13” paragraph below for more details. The reader unfamiliar with optimal transport or metric measure geometry is invited to think of the special case of M being a smooth, globally hyperbolic Lorentzian spacetime, endowed with its volume form (or a weighted version thereof); some of our results are new even in this situation.

Our main motivation for this work is the crucial role played by

$$\text{Ric} \geq k \quad \text{in all timelike directions} \quad (1.1)$$

in classical general relativity as an energy condition. We provide a unified synthetic treatment of the following three special cases; cf. [36, 74, 130] for background.

- **Timelike convergence condition.** This refers to (1.1) for constant k ; for vanishing k and assuming general relativity, this property is also termed *strong energy condition*. Independently, McCann [91] and Mondino–Suhr [99] pioneered its optimal transport characterization. This led to the timelike curvature-dimension condition of Cavalletti–Mondino [44]. An alternative notion was later proposed by Braun [21].
- **Weak energy condition.** Assuming general relativity, this special property asserts lower boundedness of Ric by the function $\text{scal}/2 - \Lambda$ in every timelike direction. Here, $\Lambda \in \mathbf{R}$ is the cosmological constant. It is believed to hold in most physically reasonable circumstances.
- **Null convergence condition.** A far less evident scenario fitting into (1.1) is the nonnegativity of the Ricci curvature in all *null* directions. Its equivalence to locally uniform lower bounds on the tensor Ric in all *timelike* directions — thus to (1.1) — has recently been shown by McCann [92], thereby addressing an open question of [45]. An alternative approach is suggested by work of Ketterer [83]. The null energy condition describes the stress-energy experienced by light rays and is expected to hold throughout the entire non-quantum universe. Besides this physical relevance, it is the basis for Penrose’s singularity theorem [107], Hawking’s area monotonicity [73], Galloway’s null splitting theorem [64], topological censorship [52], etc.

Moreover, our abstract formulation of (1.1) accommodates two further important situations on a synthetic level.

- **Asymptotic conditions.** Several spacetimes that describe cosmological phenomena fall into the category of Friedmann–Lemaître–Robertson–Walker geometries [13, 94, 105, 130], or related models such as those of [127]. Such warped products are the mathematically best understood models of the universe. In general, their timelike Ricci curvature is not uniformly bounded from below, but instead behaves as a function of the time coordinate; asymptotic conditions on it yield singularity theorems [58, 62]. Moreover, (1.1) allows one to phrase averaged energy conditions [57, 126] (in terms of integral assumptions on k along geodesics, so-called “worldline bounds”), which lead to smooth singularity theorems as well [20, 112].
- **Integral bounds.** Integrated controls on the negative part of the Ricci tensor play a fundamental role in Riemannian geometry [48]. Analogously to semiclassical phase space estimates for Schrödinger operators, L^p -bounds for $2p > \dim M$ are particularly frequent, e.g. [10, 108, 109]; on metric measure spaces, their stability has been explored in [82]. In general relativity, these are called “worldvolume bounds”. Several relevant models exhibit curvature behavior which is tame in an integrated way despite being badly behaved pointwise [29, 28, 68]. Activities in this direction have also triggered interest in the parallel null case [56, 59].

We will justify the synthetic treatment of asymptotic conditions and integral bounds by establishing one singularity theorem in each framework, see Theorem 1.2 and Theorem 1.3 below for details.

Results.

Metric measure spacetimes. We work in the abstract setting of a *metric measure spacetime*. Roughly speaking, this is a triple $\mathcal{M} := (\mathbb{M}, l, \mathfrak{m})$ consisting of

- a space \mathbb{M} endowed with a first-countable topology,
- a *signed time separation function* $l: \mathbb{M}^2 \rightarrow [0, \infty] \cup \{-\infty\}$ satisfying the reverse triangle inequality, and
- a Radon measure \mathfrak{m} on \mathbb{M} (assumed with full support in this introduction to simplify the presentation).

Its definition is tailored towards three further key properties, namely

- **global hyperbolicity**, i.e. the causality relation induced by l is a closed order and causal diamonds are compact, following the analogous definition for topological preordered spaces [101] suggested by Minguzzi [97],
- a **length** property, i.e. for every $x, y \in \mathbb{M}$ the quantity $l(x, y)$ is arbitrarily well approximated by the (l -)length of a causal curve from x to y , and
- **regularity**, i.e. if $l(x, y) > 0$, every nowhere constant causal curve with length $l(x, y)$ is already timelike.

Such a nonsmooth approach to “spacetime structures” is not new. First proposals by Busemann [34] and Kronheimer–Penrose [85] already date back to the sixties. Causality theory on continuous spacetimes has been studied order-theoretically by Sorkin–Woolgar [115] (see also [49, 65]). A recent systematic theory capturing various causality phenomena are Kunzinger–Sämman’s *Lorentzian length spaces* [87], motivated by the well-developed theory of length metric spaces in positive signature. Their article also introduces timelike sectional curvature bounds à la Alexandrov [4, 9, 30]; moreover, the timelike curvature-dimension conditions of [21, 44] have been set up in this framework. The theory of Kunzinger–Sämman has been streamlined in [92] in the globally hyperbolic, regular case. Related to [19, 87, 102], and tailor-made for Lorentz–Gromov–Hausdorff convergence, are Minguzzi–Suhr’s bounded Lorentzian metric spaces [98]. (For another approach to Lorentz–Gromov–Hausdorff convergence see [100].) In place of a time separation function, Sormani–Vega encode the “metric” properties of a continuous spacetime by their null distance [116]. This has been adapted to Lorentzian length spaces in [88].

Our first contribution further relaxes the approach of [92]. In fact, many of our results here follow known proofs distributed over various works, e.g. [3, 32, 95, 96, 97, 98]; we provide a self-contained and streamlined exposition on these. Our innovation lies in the generality of our setting. We completely waive rectifiability assumptions on causal curves and any other property which depends on the choice of a metrizable topology, cf. Remark 2.13. In particular, our compactness results solely rely on causality arguments. This is useful since in the strict sense of [87], causality of curves is *not* preserved under affine or proper-time reparametrizations. Moreover, Lipschitz continuity of causal curves per se imposes an additional requirement to deal with in the construction of limits.

Under quite weak assumptions, cf. Hypothesis 2.8 (length property and compact diamonds), we establish key properties like global hyperbolicity (Theorem 2.12), finiteness of l_+ (Corollary 2.16), topological facts pertaining to Hausdorffness, local compactness, and Polishness (Corollary 2.17 and Theorem 2.18), the existence of generalized time functions (Theorem 2.21), a limit curve theorem (Theorem 2.32), and more. The analog of Avez–Seifert’s theorem stated in Theorem 2.36 holds as well: any two points $x, y \in \mathbb{M}$ with $l(x, y) > 0$ are connected by an l -geodesic. The latter means a map $\gamma: [0, 1] \rightarrow \mathbb{M}$ such that for every $0 \leq s < t \leq 1$,

$$l(\gamma_s, \gamma_t) = (t - s)l(\gamma_0, \gamma_1) > 0.$$

The additional assumption of regularity, cf. Hypothesis 2.41, will ensure every such curve is continuous. This fact observed in [92] is recalled in Lemma 2.46.

Second, we study the “Lorentzian” geometry of the space of Borel probability measures $\mathcal{P}(\mathcal{M})$. Following [51], a natural causality relation \preceq on $\mathcal{P}(\mathcal{M})$ is defined by setting $\mu \preceq \nu$ if there is a coupling $\pi \in \Pi(\mu, \nu)$ with support in $\{l \geq 0\}$.

Theorem 1.1 (Global hyperbolicity, cf. Theorem B.5). *Hypothesis 2.8 implies the space $\mathcal{P}(\mathcal{M})$ is globally hyperbolic in the following sense.*

- (i) *The relation \preceq is a narrowly closed order.*
- (ii) *If $\mathcal{C} \subset \mathcal{P}(\mathcal{M})$ is narrowly compact, so is its causal emerald $\mathcal{J}(\mathcal{C}, \mathcal{C})$.*

This result is new even in the smooth setting of [51]. In other words, global hyperbolicity “lifts” from \mathcal{M} to $\mathcal{P}(\mathcal{M})$. This adds an entire class of new topological ordered spaces to nonsmooth general relativity. It motivates a systematic study of causality (and chronology) on $\mathcal{P}(\mathcal{M})$. An interesting question — beyond the scope of our work — is to which extent other phenomena lift from \mathcal{M} to $\mathcal{P}(\mathcal{M})$ or “project down” in the reverse direction. This would parallel general known metric correspondences between Wasserstein spaces and their bases [129]. In addition, we thoroughly study the analog of l -geodesics for probability measures [91], set up in Definition 2.61 in terms of the total transport cost ℓ_p defined in (2.12); here $0 < p < 1$. Notably, we establish criteria for their narrow continuity using regularity (Corollary 2.64), a lifting theorem (Proposition 2.66), and a limit curve theorem (Theorem 2.68). This relates the narrow topology to properties merely defined in terms of ℓ_p .

Timelike curvature-dimension conditions. In the second part of this work, we define the conditions $\text{TCD}_p^e(k, N)$ and $\text{TCD}_p(k, \infty)$ for \mathcal{M} and derive several properties. Our approach is inspired by the theory of metric measure spaces with variable Ricci curvature lower bounds without [122] and with [79, 80] dimensional restrictions (see also [25, 123]). In rough terms, these properties are phrased by convexity properties of the Boltzmann entropy with respect to \mathfrak{m} along ℓ_p -geodesics interpolating any two appropriate elements of $\mathcal{P}(\mathcal{M})$. We discuss some of the motivating examples indicated in the “Outline” paragraph in Chapter 4. We also state a stability result in Theorem 4.5¹. We prove geometric inequalities à la

- Brunn–Minkowski (Theorem 5.1),
- Bonnet–Myers (Theorem 5.5), and
- Bishop–Gromov (Theorem 5.9).

Moreover, we extend a singularity theorem by Frankel–Galloway from the smooth [58, 62] to the synthetic setting; see Theorem 1.2 below for its sharp form.

Finer results are established under the hypothesis of timelike p -essential non-branching, cf. Definition 2.70. Roughly speaking, the latter means that ℓ_p -optimal transports never charge sets of forward or backward branching l -geodesics. This notion introduced in [21] following [111] relaxes timelike nonbranching [44]. First, by local lower boundedness of k , all qualitative properties from [21, 22, 44] transfer to our setting: uniqueness of chronological ℓ_p -optimal couplings (Theorem 3.23) and ℓ_p -geodesics (Theorem 3.24), and the existence of good geodesics (Theorem 3.26). Our contribution lies in proving the foregoing results by only assuming $\text{TCD}_p^e(k, N)$ locally, cf. Definition 3.22. This gives a more streamlined proof of the equivalence of $\text{TCD}_{p,\text{loc}}^e(k, N)$ and $\text{TCD}_p^e(k, N)$ (Theorem 3.30), pioneered in [21] for constant k . Lastly, via a pathwise characterization of $\text{TCD}_p^e(k, N)$ we identify the latter with its above mentioned weak version (Theorem 3.27).

¹As explained in the introduction of [21] (and already contained in the axioms of [44]), this necessitates the definition of weak versions of $\text{TCD}_p^e(k, N)$ and $\text{TCD}_p(k, \infty)$.

From a technical point of view, one may ask if $\text{TCD}_p^e(k, N)$ is independent of the transport exponent p under suitable nonbranching hypotheses. This is unknown even for constant k [45], yet expected by analogous facts for CD metric measure spaces [2, 24] based on [39].

Localization. In the third (rather technical) part, we develop the Lorentzian localization paradigm in the variable case. The broad term “localization” from convex geometry refers to the reduction of a high-dimensional problem to an often easier one-dimensional task. It originates in Payne–Weinberger’s treatment of the optimal Euclidean Poincaré inequality [106] and was further developed by Gromov–Milman [72], Lovász–Simonovits [90], and Kannan–Lovász–Simonovits [76]. Related ideas were used to study optimal transport maps in \mathbf{R}^n by Sudakov [124], Ambrosio [6], Caffarelli–Feldman–McCann [35], and Trudinger–Wang [128], and in more general geometries by Feldman–McCann [55], Bianchini–Cavalletti [16], and Cavalletti [37]. As a technique for deriving inequalities on smooth Riemannian manifolds, this method was pioneered by Klartag [84]. This inspired Cavalletti–Mondino’s localization paradigm for CD metric measure spaces [40] and their Lorentzian entropic TCD spaces [44].

We consider a Borel function $u: E \rightarrow \mathbf{R}$ on an l -geodesically convex Borel set $E \subset M$. Assume every $x, y \in E$ satisfies

$$u(y) - u(x) \geq l(x, y). \quad (1.2)$$

Then an appropriate subset $\mathcal{J}_u \subset E$ admits a partition into l -geodesic “rays” M_α — also termed “needles” — passing through $\alpha \in Q$, where $Q \subset E$ is an index set, with the following property. The restriction $\mathbf{m}' := \mathbf{m} \llcorner \mathcal{J}_u$ admits a disintegration $d\mathbf{m}'(x) = d\mathbf{m}_\alpha(x) d\mathbf{q}(\alpha)$ into “conditional measures” \mathbf{m}_α concentrated on M_α , where \mathbf{q} is a Borel probability measure on Q (Theorem 6.31). In fact, for \mathbf{q} -a.e. $\alpha \in Q$ we show \mathbf{m}_α is absolutely continuous with respect to the one-dimensional Hausdorff measure on M_α (Proposition 6.33) and the induced metric measure space obeys a variable curvature-dimension condition in the sense of Ketterer [79] (Theorem 6.37). In addition, equality holds in (1.2) along \mathbf{q} -a.e. needle. Viewing these observations from the perspective of the null convergence condition, it is interesting to note that nontrivial curvature information can be extracted for \mathbf{q} -a.e. ray from this “codimension one” hypothesis.

Besides the variability of k , our localization complements [44] as follows. First, by starting from a stronger assumption, instead of the measure contraction property we establish a stronger curvature property of \mathbf{q} -a.e. ray. For constant k , such a result has been achieved by Cavalletti–Mondino in connection with their work on Lorentzian isoperimetry [46], as they first communicated to us in 2022. Second, we assume mere timelike p -essential nonbranching. This does not exclude branching l -geodesics, and some effort is needed to show branching points are \mathbf{m} -negligible (Corollary 6.29); our arguments are inspired by [37]. Third, in [44] u was chosen as the signed distance function from an achronal, future timelike complete (FTC) Borel set. The advantage in allowing for more general u comes from ℓ_1 -optimal transport theory. In typical applications, u will be a Kantorovich potential for an ℓ_1 -optimal transport problem. The technical background, Theorem 7.2, seems new in Lorentzian signature, as does the accompanying Kantorovich duality formula for ℓ_1 (Theorem C.3). For constant k , similar results appear in [46], which we did not see until after the present work had been circulated. Related developments are also pursued in [1].

Applications. To simplify the presentation, let $N > 1$. Using Theorem 7.2, inspired by [41] we improve all geometric inequalities from Chapter 5 to their sharp versions; we refer to [44, Rem. 5.11] and Example 7.9 for elaborations on sharpness. This

notably pertains to the timelike Brunn–Minkowski inequality (Theorem 7.6), whose sharp form under $\text{TCD}_p^e(k, N)$ was unknown even for constant k .

A particular application is the subsequent sharp singularity theorem, valid under an *asymptotic* condition on the timelike Ricci curvature. It extends a smooth result by Frankel–Galloway [58, 62].

Theorem 1.2 (Synthetic Frankel–Galloway singularity theorem, cf. Corollary 7.8). *Assume \mathcal{M} is a timelike p -essentially nonbranching $\text{TCD}_p^e(k, N)$ space, and suppose that some $o \in \mathcal{M}$, $\beta > 0$, and $R > 0$ satisfy*

$$k \geq (N - 1) \left[\frac{1}{4} + \beta^2 \right] l(o, \cdot)^{-2} \quad \text{on } \{l(o, \cdot) > R\}.$$

Then the l -distance of any two points in $I^+(o)$ does not exceed $Re^{\pi/\beta}$.

An analogous past version is also available. The importance of Theorem 1.2 is that k is not hypothesized to be everywhere nonnegative. It already predicts timelike geodesic incompleteness provided the timelike Ricci curvature does not decay too quickly in the far future of o . This result has predecessors in smooth [10, 62, 114] and nonsmooth [80] Riemannian geometry. Since Schneider’s work [114] was the first to state the explicit asymptotic condition of Theorem 1.2, in the sequel we call its conclusion the *timelike Schneider inequality*.

Lastly, we show the following Hawking-type singularity theorem under an *integral bound* on the negative part of the timelike Ricci curvature. Its smooth predecessor has recently been shown by Graf–Kontou–Ohanyan–Schinnerl [68] under an additional uniform lower boundedness assumption on the timelike Ricci curvature. For simplicity, we state the synthetic counterpart to their result here; an extension is provided in Theorem 7.14. Two quantities require a quick introduction: \mathfrak{h}_0 is the “surface measure” of V , and $V^{0,T}$ is the radial “future development” of V up to temporal height T ; see Subsections 7.3.1 and 7.3.2.

Theorem 1.3 (Synthetic Hawking-type singularity theorem, cf. Theorem 7.14). *Let \mathcal{M} be a timelike p -essentially nonbranching proper $\text{TCD}_p^e(k, N)$ space with $k \geq K$ for some constant $K < 0$. Let $V \subset \mathcal{M}$ be an achronal FTC Borel set with $\mathfrak{m}[V] = 0$ and $\mathfrak{h}_0[V] < \infty$ whose forward mean curvature is bounded from above by $-\beta$ in the sense of Definition 7.12, where $\beta > 0$. Assume there exist $T > 0$ and $\delta > 0$ with*

$$\mathfrak{h}_0[V]^{-1} \int_{V^{0,T}} k_- \, \mathfrak{d}\mathfrak{m} < \left[\frac{\sinh(\sqrt{-K/(N-1)} \delta)}{\sinh(\sqrt{-K/(N-1)} (T + \delta))} \right]^{N-1} \left[\beta - \frac{N-1}{T} \right].$$

Then there is a radial l -geodesic starting in V whose maximal domain of definition in $I^+(V) \cup V$ is strictly contained in $[0, T + \delta)$.

The smooth result of [68] has been inspired by Cheeger–Colding [48] and Sprouse [117]. A different synthetic version of Hawking’s singularity theorem under uniform timelike Ricci curvature bounds has been shown in [44]. A counterpart in positive signature can be found in [33]. An interesting ingredient of independent interest in our proof is a new nonsmooth form of Cheeger–Colding’s *segment inequality* (Proposition 7.13). It was used in [48] in the context of the almost splitting theorem. Roughly speaking, it estimates the line integral of a function along a single needle by its integral over a normal tube.

We point out that Theorem 1.2 and Theorem 1.3 are new even in spacetimes of low regularity, notably for Lorentzian metrics of class $C^{1,1}$ (or C^1 plus timelike p -essential nonbranching) thanks to Theorem 4.2 and [23]. Under the strong energy condition however, a more traditional analog of the Hawking theorem was obtained for C^1 -metrics by Graf; see [67] and its references for related results.

Motivation of Definition 3.13. The setup of $\text{TCD}_p^e(k, N)$ for variable k is a bit technical. Loosely speaking, it is based on (Jacques Charles François) Sturm’s comparison theorem for certain ODEs involving k . Below, we illustrate its definition by outlining the smooth proof of “ $\text{Ric} \geq k$ in all timelike directions implies $\text{TCD}_p^e(k, N)$ ”; a similar line of thought motivates the definition of $\text{TCD}_p(k, \infty)$, cf. e.g. the introduction of [121]. We temporarily neglect regularity issues and refer to Theorem 4.1 for details.

Let $\mu_0, \mu_1 \in \mathcal{P}(\mathbf{M})$ be appropriate mass distributions with \mathfrak{m} -densities ρ_0 and ρ_1 , respectively. The results of [91] imply the unique ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ from μ_0 to μ_1 is given by $\mu_t = (T_t)_\# \mu_0$. Here $(T_t(x))_{t \in [0,1]}$ is an l -geodesic for μ_0 -a.e. $x \in \mathbf{M}$. Furthermore, the \mathfrak{m} -density ρ_t of μ_t solves the Monge–Ampère-type equation

$$\rho_0 = (\rho_t \circ T_t) |\det dT_t| \quad \mu_0\text{-a.e.}$$

The negative logarithmic Jacobian $-\log j$, where $j := |\det dT_t|$, solves a Riccati-type equation. The hypothesis “ $\text{Ric} \geq k$ in all timelike directions” and several manipulations thus imply the following ODI for every $0 < t < 1$:

$$(j_t^{1/N})'' \leq -\frac{(k \circ T_t) |T_t'|^2}{N} j_t^{1/N} \quad \mu_0\text{-a.e.}$$

Sturm’s comparison theorem relates $j_t^{1/N}$ μ_0 -a.e. to the solution v of the ODE

$$v''(t) + \frac{(k \circ T_t) \theta^2}{N} v(t) = 0$$

satisfying $v(0) = j_0^{1/N}$ and $v(1) = j_1^{1/N}$, where $\theta := |T_t'| = l(\cdot, T_1)$, viz.

$$j_t^{1/N} \geq v(t) \quad \mu_0\text{-a.e.} \tag{1.3}$$

On the other hand, v is explicitly given by

$$v(t) = \sigma_{\kappa^-/N}^{(1-t)}(\theta) j_0^{1/N} + \sigma_{\kappa^+/N}^{(t)}(\theta) j_1^{1/N}.$$

Here, for μ_0 -a.e. $x \in \mathbf{M}$ the functions $\kappa^\pm: [0, \theta] \rightarrow \mathbf{R}$ designate the forward and backward potentials along the given l -geodesic $(T_t(x))_{t \in [0,1]}$ defined by the relations $\kappa^+(t\theta) := k \circ T_t(x)$ and $\kappa^-(t\theta) := k \circ T_{1-t}(x)$. The inherent *distortion coefficients* $\sigma_{\kappa^\pm/N}^{(t)}(\theta)$, discussed in detail in Definition 3.1 and Definition 3.4, arise from the unique solution $u(t) := \sigma_{\kappa^+/N}^{(t)}(\theta)$ of the ODE

$$0 = u''(t) + \frac{\kappa^+(t\theta) \theta^2}{N} u(t) = u''(t) + \frac{(k \circ T_t) \theta^2}{N} u(t)$$

with boundary data $u(0) = 0$ and $u(1) = 1$; analogously for κ^- . An integrated, thus more robust version of (1.3) is precisely the estimate we use to phrase $\text{TCD}_p^e(k, N)$ in Definition 3.13. Compare also with the proof of Theorem 3.27.

In the previous discussion, we have implicitly regarded k as continuous. The lower semicontinuous case will be handled by approximation, recalling that every such function is the locally monotone limit of continuous functions, cf. (3.6).

Organization. In Chapter 2, we develop our notion of globally hyperbolic, regular length metric measure spacetimes. The proofs of some results therein which require not only l -geodesics but their merely causal analogs are outsourced to Appendix B. Chapter 3 introduces our variable timelike curvature-dimension condition, which is preceded by a thorough discussion of the relevant distortion coefficients and succeeded by first elementary properties. In Chapter 4, we provide examples of variable TCD spaces obeying the timelike convergence condition, the weak energy condition, and the null convergence condition, respectively. In Chapter 5, we show geometric inequalities. Chapter 6 develops the localization paradigm with respect to a general 1-step potential. In Chapter 7, these insights are applied to improve all geometric inequalities from Chapter 5 to their sharp form, provided a timelike

essential nonbranching condition holds; the Kantorovich duality formula for ℓ_1 required for the inherent localization by optimal transport is shown in Appendix C. Chapter 7 also proves Theorem 1.2 and Theorem 1.3. Finally, Appendix A outlines the variable timelike measure-contraction property and its consequences.

2. METRIC SPACETIMES

We now describe our setting of *metric spacetimes*. It is a version of the approach to Lorentzian (pre-)length spaces by Kunzinger–Sämman [87], but differs from it in two major aspects.

- It simplifies the axiomatization by taking into account the recent contributions [3, 92, 96, 97, 98]. (We refer to [92] for a similar simplification, and to [14] for a slightly different forthcoming approach.)
- It is tailored towards the requirements of *global hyperbolicity* and a *length* condition, which we assume almost from the outset, deriving all other desired properties from these; the convenient assumption of *regularity* will be added a bit later.

Although an attempt to impose as few assumptions as possible has been made, some machinery needs to be developed, and many results from [3, 32, 95, 96, 97, 98] have to be reproven. If interested in a shortcut, we refer the reader to Theorem 2.42 for an equivalent definition of globally hyperbolic, regular length metric spacetimes as proposed by Definition 2.7 and Theorem 2.12.

2.1. Basic notions. Let M be a space endowed with a first-countable topology τ . Our discussion will be accompanied by several more sophisticated properties of τ implied by our standing Hypothesis 2.8 (length metric spacetime with compact diamonds). Most notably, in Corollary 2.17 it is shown to be locally compact and Hausdorff; in particular, τ is Polish if and only if it is second-countable. This will be assumed from Section 2.6, but our framework can be partly generalized if we only work with probability measures living on compact subsets of M ; cf. Theorem 2.18.

2.1.1. *Signed time separation function.* We adopt the conventions

$$\infty + (-\infty) := -\infty + \infty := -\infty + (-\infty) := -\infty. \quad (2.1)$$

Definition 2.1 (Signed time separation). *A function $l: M^2 \rightarrow [0, \infty] \cup \{-\infty\}$ is a signed time separation function if it has the following properties.*

- a. **Causality** [87, Def. 2.35]. *It is nonnegative on the diagonal of M^2 , and whenever two points $x, y \in M$ satisfy*

$$\min\{l(x, y), l(y, x)\} > -\infty,$$

we already have $x = y$.

- b. **Continuity.** *It is upper semicontinuous, and its positive part l_+ is lower semicontinuous.*

- c. **Reverse triangle inequality.** *For every $x, y, z \in M$,*

$$l(x, z) \geq l(x, y) + l(y, z). \quad (2.2)$$

The second property of causality is naturally called *antisymmetry*. By (2.2) and causality, we have $l(x, x) \in \{0, \infty\}$ for every $x \in M$.

The continuity of l_+ has to be understood with respect to the extended topology of the half-line $[0, \infty]$; we do not exclude the case $l(x, x) = \infty$ for some $x \in M$ a priori. Until Corollary 2.16 — which is only used around Lemma B.1 — upper semicontinuity of l_+ will not play any role.

The following remark can be skipped at first reading.

Remark 2.2 (Possible generalizations). In two aspects, it is interesting to study whether our assumptions made thus far can be weakened. To make the exposition less cumbersome, though, we leave these possible generalizations to future work.

- **Sequential topology.** One could merely assume τ to be *sequential*, i.e. it is completely determined by convergence of sequences; see [53] for details. Every (quotient of a) first-countable space is sequential, and unlike the latter property, first-countability is not stable under quotients in general. That way, one could always guarantee antisymmetry by taking a quotient [92, Lem. 1]. On the other hand, sequential topologies come with various subtleties to be addressed with care. For instance, the product of two sequential topologies is not necessarily sequential; neither does compactness imply sequential compactness in this generality, nor vice versa.
- **Upper semicontinuity near the diagonal.** It seems merely assuming upper semicontinuity of l near the diagonal of M^2 suffices. For smooth spacetimes, this follows from strong causality [12, Cor. 3.5], in turn implied by compact diamonds in all physically relevant cases [75, Thm. 2.8]. In our approach, it appears possible to prove the affected results Lemma B.1, Theorem 2.32, and Lemma 2.35 first in small chronological diamonds, then get upper semicontinuity of l as in [87, Thm. 3.28], and finally redo the proofs of the previous results globally. Corollary 2.16 and Theorem 2.18 only need continuity of l_+ near the diagonal.

Let such a signed time separation function l be fixed in the sequel. For brevity, every property of the tuple (M, l) will be formulated for M only.

2.1.2. *Causality and chronology.* We define a relation \leq on M by

$$x \leq y \quad :\iff \quad l(x, y) \geq 0$$

and call it *causality*. Definition 2.1 ensures the reflexivity of \leq and its antisymmetry, i.e. if $x, y \in M$ satisfy $x \leq y$ and $y \leq x$, then necessarily $x = y$. The reverse triangle inequality (2.2) implies transitivity of \leq . These observations are summarized by saying the triple (M, τ, \leq) is a topological ordered space [97, 101]. If $x \leq y$ we term x to lie in the *causal past* of y , and y to lie in the *causal future* of x ; a future orientation is implicit in this and subsequent terminology. We define

- the *causal future* of x by $J^+(x) := \{l(x, \cdot) \geq 0\}$,
- the *causal past* of y by $J^-(y) := \{l(\cdot, y) \geq 0\}$, and
- the *causal diamond* between x and y by $J(x, y) := J^+(x) \cap J^-(y)$.

Given any set $X \subset M$, we define

$$J^+(X) := \bigcup_{x \in X} J^+(x),$$

and analogously $J^-(Y)$ for a given set $Y \subset M$. We also define

$$J(X, Y) := J^+(X) \cap J^-(Y)$$

and call this a *causal emerald* provided X and Y are compact.

We will write $x < y$ if $x \leq y$ yet $x \neq y$.

Furthermore, we define a relation \ll on M by

$$x \ll y \quad :\iff \quad l(x, y) > 0.$$

The induced relation \ll is called *chronology*. By (2.2), it is transitive. All notions from the preceding paragraph are defined analogously for \ll instead of \leq , where every terminological occurrence of “causal” is replaced by “chronological”. Define

- the *chronological future* of x by $I^+(x) := \{l(x, \cdot) > 0\}$,
- the *chronological past* of y by $I^-(y) := \{l(\cdot, y) > 0\}$, and

- the *chronological diamond* between x and y by $I(x, y) := I^+(x) \cap I^-(y)$.

Given any sets X and Y as above, the open sets $I^+(X)$, $I^-(Y)$, and $I(X, Y)$ are defined analogously.

These properties turn the triple (M, \ll, \leq) into a *causal space* [87, Def. 2.1], as originally studied in a slightly different form by Kronheimer–Penrose [85].

Let us collect the following elementary properties, cf. [87, Lems. 2.10, 2.12, Prop. 2.13]. The push-up property follows from (2.2), while openness is implied by lower semicontinuity of l_+ .

Lemma 2.3 (Push-up and openness of chronology). *The two relations \leq and \ll have the following properties for every $x, y, z \in M$ and every $X \subset M$.*

- Push-up.** *If $x \leq y \ll z$ or $x \ll y \leq z$, then $x \ll z$.*
- Openness.** *The set $I^\pm(X)$ is open. Moreover, \ll is an open relation in the sense that $\{l > 0\}$ is an open subset of M^2 .*

Here, the set X does not need to be open, or to have any other regularity.

2.2. Length metric spacetimes. We now set up a condensed and more flexible notion of length metric spacetimes akin to Lorentzian length spaces [87, Sec. 3.4]. When coupled with compact diamonds, as we explore below it implies almost all desired causal properties of M except regularity, cf. Section 2.5.

In the following two subsections, let $a < b$ and $c < d$ be real numbers.

2.2.1. Causal curves and loops.

Definition 2.4 (Causal character). *A curve² $\gamma: [a, b] \rightarrow M$ is called*

- causal if it is continuous, and $\gamma_s \leq \gamma_t$ for every $a \leq s < t \leq b$,*
- strictly causal if it is causal yet nowhere constant, i.e. $\gamma_s < \gamma_t$ for every $a \leq s < t \leq b$,*
- timelike if it is continuous, and $\gamma_s \ll \gamma_t$ for every $a \leq s < t \leq b$, and*
- null if it is causal yet $\gamma_s \not\ll \gamma_t$ for every $a \leq s < t \leq b$.*

Any of these causality properties may be referred to as the *causal character* of γ . If a curve $\gamma: [a, b] \rightarrow M$ has one of the mentioned causal characters yet is not required to be continuous, we term it *rough causal*, *rough strictly causal*, *rough timelike*, or *rough null*, respectively. Moreover, any nonconstant path $\gamma: [a, b] \rightarrow M$ with $\gamma_a = \gamma_b$ is called a *loop*.

Since we have not excluded the possibility $l(x, x) > 0$ for some $x \in M$ yet, it is not clear at this point that every timelike curve is strictly causal.

Lastly, by reverse triangle inequality (2.2) a curve $\gamma: [a, b] \rightarrow M$ is rough null if and only if it is rough causal and $\gamma_a \not\ll \gamma_b$.

2.2.2. Length.

Definition 2.5 (Length functional). *The length of a curve $\gamma: [a, b] \rightarrow M$ is*

$$\text{Len}_l(\gamma) := \inf \sum_{i=1}^m l(\gamma_{t_{i-1}}, \gamma_{t_i}),$$

where the infimum is taken over all $m \in \mathbf{N}$ and all partitions $a = t_0 < t_1 < \dots < t_{m-1} < t_m = b$ of $[a, b]$ of cardinality $m + 1$.

²We use the terms *curve* or *path* synonymously for a map whose domain is a convex subset of \mathbf{R} . In particular, we do not require its continuity a priori. If relevant, (possible dis)continuity is always clear from our terminology or stressed explicitly.

These notions can be extended to intervals of the form (a, b) , $[a, b)$, and $(a, b]$, including unbounded ones, using limits.

Note that $\text{Len}_l(\gamma) \geq 0$ if and only if $\gamma: [a, b] \rightarrow \mathbf{M}$ is rough causal by (2.1).

The following properties are proven in exactly the same way as [87, Lems. 2.25, 2.27, 2.28]. As usual, by a *reparametrization* of a curve $\gamma: [a, b] \rightarrow \mathbf{M}$ we mean a map $\sigma: [c, d] \rightarrow \mathbf{M}$ such that $\sigma = \gamma \circ \phi$, where $\phi: [c, d] \rightarrow [a, b]$ is *continuous* and strictly increasing. Reparametrization does not change the causal character of a rough curve [87, Lem. 2.27].

Lemma 2.6 (Concatenation and reparameterization invariance). *The following properties hold.*

(i) **Additivity.** *If $\gamma: [a, b] \rightarrow \mathbf{R}$ is a rough causal curve and $a < c < b$, then*

$$\text{Len}_l(\gamma) = \text{Len}_l(\gamma|_{[a,c]}) + \text{Len}_l(\gamma|_{[c,b]}).$$

(ii) **Reparametrization invariance.** *The functional Len_l is invariant under reparametrizations of rough causal curves.*

We will thus mostly restrict our attention to rough curves defined on $[0, 1]$.

2.2.3. *Imposing a length property.* Every rough causal curve $\gamma: [0, 1] \rightarrow \mathbf{M}$ satisfies

$$\text{Len}_l(\gamma) \leq l(\gamma_0, \gamma_1).$$

In general, this inequality is strict, but one can at least ask these quantities to be arbitrarily close to each other once the endpoints are fixed. This leads to the Lorentzian analog of length metric spaces [31], see also [87, 98].

Definition 2.7 (Length metric spacetime). *A set \mathbf{M} equipped with a first countable topology τ and a signed time separation function l as above is termed length metric spacetime if the following statements hold.*

- a. **Nontrivial chronology.** *Every point $x \in \mathbf{M}$ is the midpoint of a timelike curve $\gamma: [0, 1] \rightarrow \mathbf{M}$, i.e. $x = \gamma_{1/2}$.*
- b. **Length property.** *If $x, y \in \mathbf{M}$ satisfy $x \ll y$, then*

$$l(x, y) = \sup \text{Len}_l(\gamma),$$

where the supremum is taken over all strictly causal curves $\gamma: [0, 1] \rightarrow \mathbf{M}$ such that $\gamma_0 = x$ and $\gamma_1 = y$.

The subtle properties of continuity and nowhere constancy of admissible curves in the length property will be crucial for us. Inter alia, chronology will be implied, i.e. the vanishing of l on the diagonal of \mathbf{M}^2 ; cf. Corollary 2.14.

The nontrivial chronology condition implies $I^\pm(x) \neq \emptyset$ for every $x \in \mathbf{M}$. This is one requirement for localizability in [87, Def. 3.16]. Unlike [87, Def. 3.1], we merely assume chronologically related points to be connected by strictly causal (as opposed to timelike) curves, and we do not assume any a priori connectedness between merely causally related points. These will be made up for by our assumption of regularity in Section 2.5, as well as Theorem 2.36, nontrivial chronology, and continuity of l_+ , respectively.

We will exclusively deal with length metric spacetimes satisfying the following.

Hypothesis 2.8 (First standing assumption). *From now on, we assume \mathbf{M} to be a length metric spacetime with compact diamonds.*

Here and in the sequel, we say \mathbf{M} has

- *compact diamonds* if $J(x, y)$ is compact for every $x, y \in \mathbf{M}$, and
- *compact emeralds* if $J(X, Y)$ is compact for every compact $X, Y \subset \mathbf{M}$.

As in the smooth setting [105, Lem. 14.10], Hypothesis 2.8 is not well-behaved under restriction to subsets, especially causal diamonds. However, this will not be an issue for us since we can always locally use all relevant properties that hold globally, cf. Remark 2.15.

2.3. Causality theory for length metric spacetimes. Now we study numerous consequences of Hypothesis 2.8. We address

- causality and chronology properties along with
- global hyperbolicity and
- refined properties of the topology τ ,
- existence of generalized time functions,
- the causal ladder, especially with regards to strong causality, and
- a limit curve theorem.

The existence of maximizers in the length property is postponed to Section 2.4.

2.3.1. Causality, chronology, and global hyperbolicity. We will use the following notations for “monotone” convergence. For a sequence $(x_n)_{n \in \mathbf{N}}$ in M and $x \in M$, we write $x_n \downarrow x$ as $n \rightarrow \infty$ or $x_n \uparrow x$ as $n \rightarrow \infty$, respectively, if $(x_n)_{n \in \mathbf{N}}$ converges to x and, respectively, $x_{n+1} \ll x_n$ or $x_n \ll x_{n+1}$ for every $n \in \mathbf{N}$.

A cornerstone for our discussion is the following result from [3, Lem. 2.18], called the “sequence lemma” therein.

Lemma 2.9 (Timelike monotone approximation). *For every $x \in M$, there exist sequences $(a_i)_{i \in \mathbf{N}}$ and $(b_i)_{i \in \mathbf{N}}$ such that $a_i \uparrow x$ as $i \rightarrow \infty$ and $b_i \downarrow x$ as $i \rightarrow \infty$.*

Proof. Owing to the nontrivial chronology from Definition 2.7, let $\gamma: [0, 1] \rightarrow M$ be a timelike curve with midpoint x . Then setting $a_i := \gamma_{1/2-1/2^i}$ and $b_i := \gamma_{1/2+1/2^i}$ for $i \in \mathbf{N}$ leads to the desired sequences by continuity. \square

Note again that the sequences constructed above might be constant if $l(x, x) > 0$ for the given point x .

The following two consequences of Lemma 2.9 are variants of [96, Thm. 4.12] and [75, Prop. 2.3], whose proofs we adopt with no essential change. The following causal closedness is clear by upper semicontinuity of l , but we prefer to give an independent proof which does not rely on this property.

Proposition 2.10 (Causal closedness [87, Def. 3.4]). *The relation \leq is closed in the product topology τ^2 of M^2 .*

Proof. We first claim closedness of $J^\pm(x)$ for every $x \in M$. We focus on $J^+(x)$, the proof for closedness of $J^-(x)$ is analogous. Given any y in the closure $\bar{J}^+(x)$ of $J^+(x)$, there exists a directed net $(y_\alpha)_{\alpha \in A}$ in $J^+(x)$ converging to y . Let $y' \in I^+(y)$ be as provided by nontrivial chronology. By compactness of diamonds, $J(x, y')$ closed. By the openness asserted by Lemma 2.3, we have $y_\alpha \in J^+(x) \cap I^-(y')$ for sufficiently large $\alpha \in A$. Since the latter is a subset of $J(x, y')$, we obtain $y \in J^+(x)$.

To prove closedness of \leq , let $(x_\alpha, y_\alpha)_{\alpha \in A}$ be a directed net in M^2 converging to $(x, y) \in M^2$ such that $x_\alpha \leq y_\alpha$ for every $\alpha \in A$. Let $(b_i)_{i \in \mathbf{N}}$ be a sequence with $b_i \downarrow y$ as $i \rightarrow \infty$ as given by Lemma 2.9. Given any $i \in \mathbf{N}$, we obtain $y_\alpha \in I^-(b_i)$ for sufficiently large $\alpha \in A$ by Lemma 2.3. The same result implies $x_\alpha \in I^-(b_i)$ for sufficiently large $\alpha \in A$, and therefore $x \in J^-(b_i)$ by the preceding paragraph; in other words, $b_i \in J^+(x)$. Letting $i \rightarrow \infty$ and using closedness of $J^+(x)$ yields $y \in J^+(x)$, or equivalently $x \leq y$. \square

Corollary 2.11 (Compactness of emeralds). *The space M has compact emeralds.*

Proof. By Proposition 2.10 and [97, Thm. 3.3] it suffices to prove compactness of $J(C, C)$ if $C \subset M$ is compact. Closedness of $J(C, C)$ follows from Proposition 2.10 and the theory of closed ordered spaces [101], cf. [96, Thm. 4.12].

The rest follows as for [75, Prop. 2.3]. Given any $x \in C$, by nontrivial chronology there exist $a, b \in M$ with $c \in I(a, b)$. By Lemma 2.3, these diamonds are an open cover of C , hence $C \subset \bigcup_{i=1}^n I(a_i, b_i)$ for some $n \in \mathbf{N}$ and some $a_1, b_1, \dots, a_n, b_n \in M$. Taking causal pasts and causal futures of both sides of the preceding inclusion yields $J(C, C) \subset \bigcup_{i=1}^n J(a_i, b_i)$. Therefore, as a closed subset of a compact subset, the causal diamond $J(C, C)$ is itself compact. \square

As a remarkable consequence of Proposition 2.10 and Corollary 2.11, all length metric spacetimes M having compact diamonds are *globally hyperbolic* (after the modification of the classical definition of this property suggested in [95]).

Theorem 2.12 (Global hyperbolicity). *Hypothesis 2.8 implies M is globally hyperbolic, i.e. it satisfies the following two properties.*

- (i) *The relation \leq is a closed order.*
- (ii) *The causally convex hull $J(C, C)$ of every compact set $C \subset M$ is compact.*

Remark 2.13 (About various notions of global hyperbolicity). There are several slightly different definitions of global hyperbolicity in the literature. We shortly elaborate on the connections to the definition we use.

Traditionally [13, 105], for smooth spacetimes global hyperbolicity means nonexistence of almost closed strictly causal curves and compact diamonds. This property will follow from Lemma 2.27 and Proposition 2.28 below. In particular, M satisfies the Bernal–Sánchez definition of global hyperbolicity from the smooth case [15]. More generally, Theorem 2.12 echoes the results of Hounnonkpe–Minguzzi [75, 97] which inspire it: for smooth noncompact spacetimes of dimension three or higher, endowed with a $C^{1,1}$ -Lorentzian metric, [75, Thm. 2.8] asserts the causality requirements in the indicated traditional definition of global hyperbolicity are superfluous. Our result does not extend [75], though, since we *assume* causality.

For various characterizations of global hyperbolicity for spacetimes with continuous metrics, we refer to [113].

In the more recent approach [87], global hyperbolicity for Lorentzian pre-length spaces is defined by non-total imprisonment [87, Def. 2.35] and compact diamonds. By [97, Thm. 3.7], for Lorentzian length spaces [87, Def. 3.22] this is equivalent to the definition of global hyperbolicity from Theorem 2.12. ([32] shows further equivalent characterizations of global hyperbolicity in this abstract setting.) On the other hand, non-total imprisonment is a metric property, hence becomes less comparable to our setting in which τ is not required to be metrizable a priori. Yet, in fact for Lorentzian length spaces global hyperbolicity depends on the reference metric only through the topology it induces [92, Rem. 14].

The assumption of compact emeralds for Lorentzian length spaces was termed “ \mathcal{K} -global hyperbolicity” in [21, 22, 44].

We go on with further properties of \leq and \ll . Recall that a loop refers to any nonconstant path $\gamma: [0, 1] \rightarrow M$ with $\gamma_0 = \gamma_1$.

Corollary 2.14 (Enhanced properties). *The following properties hold.*

- (i) **No rough causal loops.** *If a rough causal curve $\gamma: [0, 1] \rightarrow M$ satisfies $\gamma_s = \gamma_t$ for some $0 \leq s < t \leq 1$, then $\gamma|_{[s,t]}$ is constant.*
- (ii) **Chronology** [87, Def. 2.35]. *The relation \ll is irreflexive. In other words, every $x \in M$ obeys $x \not\ll x$, and in particular $l(x, x) = 0$.*

Proof. The existence of rough causal loops would contradict causality assumed in Definition 2.1. Indeed, if $\gamma: [0, 1] \rightarrow M$ is a rough causal loop (hence nonconstant), then we get $\gamma_0 \leq \gamma_t \leq \gamma_1$ yet $\gamma_t \neq \gamma_0 = \gamma_1$ for some $0 < t < 1$. This gives (i).

To show (ii), assume $x \ll x$ for some $x \in M$. As $l(x, x) > 0$, the length property would imply the existence of a causal loop $\gamma: [0, 1] \rightarrow M$ starting and ending at x , a contradiction to (i). This shows $l(x, x) \leq 0$, while the inequality $l(x, x) \geq 0$ follows from our assumption on l from Definition 2.1. \square

Remark 2.15 (Restriction caveat). Corollary 2.14 implies that causal diamonds $J(x, y)$, where $x, y \in M$, endowed with the restriction $l|_{J(x, y)^2}$ will never obey Hypothesis 2.8 if M does. Indeed, otherwise x would be the midpoint of a timelike curve through it, and since $l(x, x) = 0$ by Corollary 2.14 we would find $z \in J(x, y) \setminus \{x\}$ with $z \ll x$. On the other hand $x \leq z$, a contradiction to said corollary.

However, this will not affect our “local” discussions, since all relevant properties we use (compactness, limit curve theorems, basic optimal transport theory, etc.) are inherited by the considered subsets from the ambient space.

Lastly, we show everywhere finiteness of l_+ . Note that we only use continuity of l_+ near the diagonal of M^2 .

Corollary 2.16 (Finiteness of time separation). *The function l_+ is real-valued.*

Proof. By Corollary 2.14, it suffices to consider $x, y \in M$ with $l(x, y) > 0$. Let $\gamma: [0, 1] \rightarrow M$ be a strictly causal curve connecting x and y , as provided through the length property. For every $0 \leq t \leq 1$, by the continuity of l_+ near the diagonal of M^2 as well as Corollary 2.14 there exists $\delta > 0$ with the property $l_+(\gamma_s, \gamma_{s'}) \leq 1$ for every $s, s' \in (t - \delta, t + \delta) \cap [0, 1]$. By compactness, finitely many such intervals $(t_i - \delta_i, t_i + \delta_i) \cap [0, 1]$ cover $[0, 1]$, where $i = 0, \dots, m$ and $m \in \mathbf{N}$. Without loss of generality, we assume $0 = t_0 < t_1 < \dots < t_{m-1} < t_m = 1$. Invoking Corollary 2.14, a telescoping sum, and the reverse triangle inequality (2.2), we obtain the desired conclusion:

$$l_+(x, y) = \sum_{i=1}^m [l_+(\gamma_{t_0}, \gamma_{t_i}) - l_+(\gamma_{t_0}, \gamma_{t_{i-1}})] \leq \sum_{i=1}^m l_+(\gamma_{t_{i-1}}, \gamma_{t_i}) \leq m. \quad \square$$

2.3.2. Topological consequences.

Corollary 2.17 (Topological properties and metrizability). *The topology τ is Hausdorff, locally compact yet noncompact, and does not have isolated points.*

In particular, τ is Polish if and only if it is second-countable.

Proof. The global hyperbolicity of Theorem 2.12 combines with [97, Thm. 3.3] to imply the Hausdorff property of τ .

Local compactness follows from nontrivial chronology and compact diamonds. Indeed, given any $x \in M$, taking appropriate points on a timelike segment passing through x and considering the spanned chronological diamond gives an open neighborhood of x with compact closure.

Compactness of τ is contradicted in the standard way, cf. [105, Lem. 14.10].

To show nonexistence of isolated points, assume to the contrary that $x \in M$ is such a point. By continuity, every timelike curve with midpoint x as given by nontrivial chronology is constant around $1/2$, which contradicts Corollary 2.14.

The last statement follows straightforwardly from these considerations. Every Polish topology is second-countable. Conversely, every second-countable, locally compact Hausdorff topology is Polish. \square

More refined topological properties can be deduced for all compact subsets of M . The proof relies on linking the theory of bounded Lorentzian metric spaces from

[98] to our setting. To streamline the presentation, we defer it to Appendix B, but recall that it only requires l_+ to be continuous near the diagonal of M^2 .

Theorem 2.18 (Polish compact sets). *Hypothesis 2.8 implies the relative topology of τ on every compact subset $C \subset M$ is Polish.*

Next, we pass to the distinction property of M shown in Corollary 2.20. To this aim, the closedness of causal pasts and causal futures implicitly proven for Proposition 2.10 leads to the identities in the following Lemma 2.19. They should be read as “non-bubbling” under Hypothesis 2.8 (in the terminology of Chruściel–Grant [49]). The proof is performed as for [3, Prop. 2.19, Lem. 3.11].

Lemma 2.19 (No causal bubbles). *For every $x \in M$,*

$$\bar{I}^\pm(x) = \{y \in M : I^\pm(y) \subset I^\pm(x)\} = J^\pm(x),$$

where $\bar{I}^\pm(x)$ denotes the closure of $I^\pm(x)$.

Proof. Since $I^\pm(x) \subset J^\pm(x)$, taking closures of both sides yields the inclusion “ \subset ” by Proposition 2.10.

On the other hand, Lemma 2.9 (timelike monotone approximation) as well as Lemma 2.3 (push-up lemma) show each $y \in J^+(x)$ to be the decreasing limit of a sequence $(y_i)_{i \in \mathbf{N}}$ in $I^+(x)$ to establish the opposite inclusion “ \supset ”.

The inclusion $\bar{I}^-(x) \supset J^-(x)$ is argued in an analogous way (or by time reversal as in Example 2.45). \square

Corollary 2.20 (Future and past distinction [3, Def. 3.3]). *If two points $x, y \in M$ satisfy $I^+(x) = I^+(y)$ or $I^-(x) = I^-(y)$, then $x = y$. In particular, the space M is distinguishing, i.e. if $I^+(x) = I^+(y)$ and $I^-(x) = I^-(y)$, then $x = y$.*

Proof. In the first case, Lemma 2.19 yields $J^+(x) = J^+(y)$ by taking closures on both sides. In turn, since \leq is reflexive this implies $x \in J^+(y)$ and $y \in J^+(x)$. The claim follows from Corollary 2.14 or the antisymmetry from Definition 2.1.

The second case is argued analogously. \square

2.3.3. *Generalized time functions.* Following [32], this corollary will be used to establish the existence of *generalized time functions* through averaged volumes.

Theorem 2.21 (Existence of generalized time functions). *Every compact subset $C \subset M$ admits a generalized time function, i.e. there exists a function $\mathfrak{t} : C \rightarrow \mathbf{R}$ with the property $\mathfrak{t}(x) < \mathfrak{t}(y)$ whenever $x, y \in C$ satisfy $x < y$.*

Proof. By covering C with finitely many chronological diamonds and taking the closure, without loss of generality we may and will assume C lies in the interior of a larger compact set $C' \subset M$. By Theorem 2.18, C' is Polish — with topology induced by, say, the metric d — and hence there is a Borel probability measure $\mu \in \mathcal{P}(M)$ with $\text{spt } \mu = C'$ [32, Prop. 4.1] (see Subsection 2.6.2 for relevant notation). For $x \in M$, define the *past averaged volume function* [32, Def. 4.2] $\mathfrak{t}^- : M \rightarrow [0, 1]$ by

$$\mathfrak{t}^-(x) := \int_0^1 \mu[d(\cdot, I^-(x)) < r] \, dr.$$

By Lemma 2.3, if $x, y \in M$ satisfy $x \leq y$ then $\mathfrak{t}^-(x) \leq \mathfrak{t}^-(y)$.

Following the proof of [32, Prop. 4.8] and applying Corollary 2.20, in fact we get $x < y$ implies $\mathfrak{t}^-(x) < \mathfrak{t}^-(y)$ for every $x, y \in C$. Here, we used the construction of C' , which ensures all sequences converging to points in C to eventually lie in C' , and hence the applicability of the “past-approximation property” from the proof of [32, Prop. 4.8] yielded by nontrivial chronology. \square

Remark 2.22 (Alternative construction). In the previous proof, one could have also taken the *future averaged volume function* $\mathfrak{t}^+ : M \rightarrow [-1, 0]$ defined by

$$\mathfrak{t}^+(x) := - \int_0^1 \mu[d(\cdot, I^+(x)) < r] dr.$$

Remark 2.23 (Global existence of generalized time functions). If τ is induced by a separable metric, Theorem 2.21 even holds for C replaced by the noncompact space M . The proof is similar, in fact simpler in this case.

2.3.4. *Causal ladder*. Now we “descend” the causal ladder from [3, Thm. 3.16] and outline how global hyperbolicity implies further standard causality properties. The proof follows [3, Props. 3.12, 3.13] by using Lemmas 2.9 and 2.19. We point out that by Proposition 2.10, there is no need to discuss *stable causality* [3, Def. 3.9].

Theorem 2.24 (Causal ladder). *The following hold in decreasing strength if we assume Hypothesis 2.8.*

- (i) **Causal simplicity** [3, Def. 3.7]. *The space M is causal in the sense of Corollary 2.14, and $J^\pm(x)$ is closed for every $x \in M$.*
- (ii) **Causal continuity** [3, Def. 3.7]. *The space M is distinguishing in the sense of Corollary 2.20, and reflective, i.e. every $x, y \in M$ satisfy $I^+(x) \subset I^+(y)$ if and only if $I^-(y) \subset I^-(x)$.*

The property of strong causality [87, Def. 2.35], which is another usual part of the causal ladders [3, Thm. 3.16] and [87, Thm. 3.26] is investigated now.

Definition 2.25 (Strong causality). *Let α be the Alexandrov topology on M , i.e. the topology on M having $\{I(x, y) : x, y \in M\}$ as a subbase.*

We call M strongly causal if the α coincides with the original topology τ .

Remark 2.26 (Chronological topology). Under strong causality, both α and τ agree with the *chronological topology* ι [87, Def. 2.4, Prop. 2.7]. The latter is simply the topology on M having $\{I^\pm(x) : x \in M\}$ as a subbase.

We obtain the following characterization of strong causality by the traditional definition through nonexistence of almost closed strictly causal curves [13, 105]. As always, topological notions such as neighborhoods are understood with respect to τ .

Lemma 2.27 (Characterizing strong causality). *Strong causality is equivalent to the following. For every $x \in M$ and every neighborhood $U \subset M$ of x , there exists a neighborhood $V \subset U$ of x such that every strictly causal curve $\gamma : [0, 1] \rightarrow M$ with endpoints in V does not leave U .*

Proof. The implication from strong causality to the second part is argued as in [87, Lem. 2.38] by using Lemma 2.3.

The converse is similar to [87, Thm. 3.26]; we give a proof since we do not use the notion of localizability. Assume to the contrary that M is not strongly causal. This means the existence of $x \in M$ and a neighborhood $U \subset M$ of x such that no chronological diamond is contained in U . By assumption, there exists an open set $V \subset U$ containing x such that all strictly causal curves with endpoints in V do not leave U . Nontrivial chronology implies existence of a timelike curve $\gamma : [0, 1] \rightarrow M$ passing through x . For sufficiently small $\delta > 0$, by continuity we have $\gamma_{1/2-\delta}, \gamma_{1/2+\delta} \in V$. By Corollary 2.14, γ is strictly causal, and therefore $x \in I(\gamma_{1/2-\delta}, \gamma_{1/2+\delta}) \subset U$. This contradicts our assumption. \square

Proposition 2.28 (Strong causality). *The space M is strongly causal.*

Proof. Assume to the contrary that there exist $x \in M$ and a neighborhood $U \subset M$ of x such that for every neighborhood $V \subset U$ of x , there exists a strictly causal curve $\gamma: [0, 1] \rightarrow M$ with endpoints in V , but which leaves U .

We claim the existence of a nonincreasing sequence $(W_n)_{n \in \mathbf{N}}$ of neighborhoods of x such that

- $\bigcap_{n \in \mathbf{N}} W_n = \{x\}$, and
- every neighborhood of x contains W_n for every sufficiently large $n \in \mathbf{N}$.

Let $(N_n)_{n \in \mathbf{N}}$ be a sequence of neighborhoods of x provided by first-countability of τ , and set $W_n := N_1 \cap \dots \cap N_n$, where $n \in \mathbf{N}$. Then $x \in \bigcap_{n \in \mathbf{N}} W_n$ by construction. Now suppose that the latter contains another point $z \in M \setminus \{x\}$. Let $W \subset M$ be a neighborhood of x which does not contain z , as provided by the Hausdorff property of τ shown in Corollary 2.17. First-countability of τ , however, yields $W_n \subset W$ for some $n \in \mathbf{N}$, a contradiction to $z \in W_n$.

By nontrivial chronology as well as Corollary 2.14, there exist $a, b \in U$ with $a \ll x \ll b$, and these points are mutually distinct. Since $I(a, b)$ is open, again by first-countability we have $W_n \subset I(a, b)$ for sufficiently large $n \in \mathbf{N}$. There is a strictly causal curve $\gamma^n: [0, 1] \rightarrow M$ with $\gamma_0^n, \gamma_1^n \in W_n$ leaving U . Since $\bigcap_{n \in \mathbf{N}} W_n = \{x\}$ and employing first-countability of τ once more, $(\gamma_0^n)_{n \in \mathbf{N}}$ and $(\gamma_1^n)_{n \in \mathbf{N}}$ converge to x . On the other hand, since $\gamma_{[0,1]}^n \subset J(a, b)$ yet γ^n leaves U , we find a sequence $(z_n)_{n \in \mathbf{N}}$ in $J(a, b) \setminus U$ such that $\gamma_0^n \leq z_n \leq \gamma_1^n$. By compactness of $J(a, b) \setminus U$ and Proposition 2.10, there is a point $z \in J(a, b) \setminus U$ with $x \leq z \leq x$. Since $z \notin U$, necessarily $x \neq z$, which contradicts Corollary 2.14. \square

Remark 2.29 (Exhaustion by emeralds). The foregoing discussions yield that if τ is induced by a separable metric or is σ -compact, then it can be exhausted by compact emeralds. In the first situation, simply cover M with countably many precompact chronological diamonds, and successively take their causally convex hulls.

In particular, in this case M can be exhausted by Polish spaces by Theorem 2.18.

2.3.5. Limit curve theorem. Since our limit curve theorems involve uniform convergence of curves, we need to clarify the meaning of this notion for the reference topology τ being merely first-countable. Also, we have to make sense of uniform convergence of curves with varying domains; cf. [98, Def. 5.10].

We employ the notation $\mathbf{N}_\infty := \mathbf{N} \cup \{\infty\}$.

Definition 2.30 (Uniform convergence of curves). *Given any $n \in \mathbf{N}_\infty$, let $\alpha_n \leq \beta_n$ be given real numbers. Set $\alpha := \alpha_\infty$ and $\beta := \beta_\infty$. Let $(\gamma^n)_{n \in \mathbf{N}_\infty}$ be a family of curves $\gamma^n: [\alpha_n, \beta_n] \rightarrow M$. We define the extension $\hat{\gamma}^n: \mathbf{R} \rightarrow M$ of γ^n by*

$$\hat{\gamma}_s^n := \begin{cases} \gamma_{\alpha_n}^n & \text{if } s < \alpha_n, \\ \gamma_s^n & \text{if } \alpha_n \leq s \leq \beta_n, \\ \gamma_{\beta_n}^n & \text{otherwise.} \end{cases} \quad (2.3)$$

We term $(\gamma^n)_{n \in \mathbf{N}}$ to converge uniformly to γ^∞ if

- a. $(\alpha_n)_{n \in \mathbf{N}}$ converges to α ,
- b. $(\beta_n)_{n \in \mathbf{N}}$ converges to β , and
- c. there exists a metric d inducing the relative topology of τ on a neighborhood of $J(\gamma_\alpha^\infty, \gamma_\beta^\infty)$ such that

$$\lim_{n \rightarrow \infty} \sup_{s \in \mathbf{R}} d(\hat{\gamma}_s^n, \hat{\gamma}_s^\infty) = 0.$$

Remark 2.31 (Disambiguation). When dealing with uniform convergence, we always work on compact subsets of M . These are Polish by Theorem 2.18, whence they always admit a uniform structure. In other words, a metric d as in Definition 2.30 will always be given on compact subsets of M .

In fact, since our notion of uniform convergence works with compact metrizable spaces, the uniform structure will always be *unique* [131, Thm. 36.19]. In particular, uniform convergence with respect to the metric from the abstract Theorem 2.18 will hold even if uniform convergence is already known with respect to some other metric (e.g. the one from a Polish τ).

We now turn to our limit curve theorem. For now, we have to allow for uniform convergence up to reparametrization. See, however, Corollary 2.47 below. Anyway this is no issue for the next results since the length functional is invariant under reparametrizations by Lemma 2.6.

Theorem 2.32 (Limit curve theorem). *Assume Hypothesis 2.8. Let $C_0, C_1 \subset \mathbf{M}$ be compact subsets, and let $(\gamma^n)_{n \in \mathbf{N}}$ be a sequence of strictly causal curves on $[0, 1]$ such that $\gamma_0^n \in C_0$ and $\gamma_1^n \in C_1$ for every $n \in \mathbf{N}$. Then $(\gamma^n)_{n \in \mathbf{N}}$ has a subsequence which — after reparametrization — converges uniformly, in the sense of Definition 2.30,*

- (i) *to a strictly causal curve defined on a compact interval $[\alpha, \beta] \subset \mathbf{R}$ of positive length unless the endpoints of the extracted subsequence converge to the same limit, and*
- (ii) *to a constant causal curve defined on $[0, 1]$ otherwise.*

As for Theorem 2.18, the lengthy proof relies on constructions from [98] and is deferred to Appendix B.

2.4. Geodesy. We now turn to *maximizing curves*, especially *geodesics*. See [87, Def. 2.33] or [98, Def. 5.9] for previous Lorentzian analogs. Global hyperbolicity somewhat replaces the local compactness assumption from metric geometry, which yields existence of curves *minimizing* length in this setting [31, Prop. 2.5.19].

Definition 2.33 (Maximizing curves). *A curve $\gamma: [0, 1] \rightarrow \mathbf{M}$ is called*

- a. *maximizing if it is causal — in particular continuous — and*

$$l(\gamma_0, \gamma_1) = \text{Len}_l(\gamma),$$

- b. *strictly causal maximizing if it is maximizing and strictly causal,*
- c. *timelike maximizing if it is maximizing and timelike, and*
- d. *null maximizing if it is maximizing and null.*

Analogously, as after Definition 2.4 we define rough versions of these. Note our convention of rough maximizers γ as above always satisfying $\text{Len}_l(\gamma) \geq 0$.

The proof of the following is standard, cf. [87, Prop. 2.34] and [98, Prop. 6.2].

Lemma 2.34 (Characterizing rough maximizers). *The following hold for every rough causal curve $\gamma: [0, 1] \rightarrow \mathbf{M}$.*

- (i) *The curve γ is a rough maximizer if and only if for every $0 \leq r < s < t \leq 1$,*

$$l(\gamma_r, \gamma_s) + l(\gamma_s, \gamma_t) = l(\gamma_r, \gamma_t).$$

- (ii) *If γ is rough null, it is a rough null maximizer.*
- (iii) *If γ is a rough causal maximizer, then any of its subsegments are rough causal maximizing. The same holds by replacing every occurrence of “rough causal” by “rough strictly causal”, “rough timelike”, or “rough null”.*

To prove existence of maximizers, we need the following lemma, whose proof is analogous to [98, Thm. 6.3].

Lemma 2.35 (Upper semicontinuity of the length functional). *Let $(\gamma^n)_{n \in \mathbf{N}}$ be a sequence of rough causal curves $\gamma^n: [0, 1] \rightarrow \mathbf{M}$ converging pointwise to a rough causal curve $\gamma: [0, 1] \rightarrow \mathbf{M}$. Then*

$$\limsup_{n \rightarrow \infty} \text{Len}_l(\gamma^n) \leq \text{Len}_l(\gamma).$$

Proof. Recall by Corollary 2.16 that l_+ is finite, hence $\text{Len}_l(\gamma) < \infty$. By definition of $\text{Len}_l(\gamma)$, given $\varepsilon > 0$ there exist $m \in \mathbf{N}$ and a partition $0 = t_0 < t_1 < \dots < t_{m-1} < t_m = 1$ of $[0, 1]$ such that

$$\sum_{i=1}^m l(\gamma_{t_{i-1}}, \gamma_{t_i}) \leq \text{Len}_l(\gamma) + \frac{\varepsilon}{2}.$$

By pointwise convergence and causality we have

$$l(\gamma_{t_{i-1}}^n, \gamma_{t_i}^n) \leq l(\gamma_{t_{i-1}}, \gamma_{t_i}) + \frac{\varepsilon}{2m}$$

for every $i = 1, \dots, m$ and every sufficiently large $n \in \mathbf{N}$. Combining these two inequalities with Lemma 2.6 then leads to

$$\text{Len}_l(\gamma^n) = \sum_{i=1}^m \text{Len}_l(\gamma^n|_{[t_{i-1}, t_i]}) \leq \sum_{i=1}^m l(\gamma_{t_{i-1}}^n, \gamma_{t_i}^n) \leq \text{Len}_l(\gamma) + \varepsilon. \quad \square$$

Theorem 2.36 (Nonsmooth Avez–Seifert theorem). *Assume Hypothesis 2.8. For every $x, y \in \mathbf{M}$ satisfying $x < y$, there exists a strictly causal maximizing curve connecting x to y . In particular, the length property from Definition 2.7 holds for x and y merely obeying $x < y$.*

Moreover, every rough causal curve connecting x to itself is a constant, and in particular a continuous maximizer.

Proof. The last statement follows from Corollary 2.14.

To show the remaining claim, first we assume $x \ll y$, and in particular $x \neq y$. Recall that $l(x, y) < \infty$ by Corollary 2.16. Thus, owing to the length property let $(\gamma^n)_{n \in \mathbf{N}}$ be a sequence of strictly causal curves from x to y with $\text{Len}_l(\gamma^n) \rightarrow l(x, y)$ as $n \rightarrow \infty$. Theorem 2.32 thus yields a uniform limit of reparametrizations of $(\gamma^n)_{n \in \mathbf{N}}$ which is a strictly causal curve connecting x to y . Since $x \neq y$, we easily see these curves can be rescaled to curves $\sigma^n: [0, 1] \rightarrow \mathbf{M}$, where $n \in \mathbf{N}$, and $\sigma: [0, 1] \rightarrow \mathbf{M}$ defined on $[0, 1]$, respectively, in such a way that $(\sigma^n)_{n \in \mathbf{N}}$ converges pointwise to σ . Since the length functional is invariant under reparametrizations by Lemma 2.6, from Lemma 2.35 we obtain

$$l(x, y) \geq \text{Len}_l(\sigma) \geq \limsup_{n \rightarrow \infty} \text{Len}_l(\sigma^n) = l(x, y).$$

This forces equality to hold throughout.

If only $x < y$, we first use Lemma 2.9 to find sequences $(x_n)_{n \in \mathbf{N}}$ and $(y_n)_{n \in \mathbf{N}}$ with $x_n \uparrow x$ as $n \rightarrow \infty$ and $y_n \downarrow y$ as $n \rightarrow \infty$, respectively. Then for every $n \in \mathbf{N}$, the points x_n and y_n are connected by a strictly causal maximizer $\gamma^n: [0, 1] \rightarrow \mathbf{M}$ by the preceding part. As above, by also taking continuity of l_+ and the fact $x \neq y$ into account, this leads to the desired strictly causal maximizer from x to y . \square

At this point, we do not know in general whether chronologically related points x and y can be connected by a *timelike* maximizer. It might happen that all such curves have null subsegments. This will be ruled out later by the assumption of regularity in Section 2.5, which will be very useful for our purposes.

A topic we first address is the question of reparametrization of curves or maximizers by l -arclength. Again the procedure is quite standard [31], and we refer the reader to [87, Prop. 3.34, Cor. 3.35] for details; the inherent assumption $l_+(x, x) = 0$ for every $x \in \mathbf{M}$ is guaranteed by Corollary 2.14.

Let us call a rough causal curve $\gamma: [0, 1] \rightarrow \mathbf{M}$ *l -rectifiable* if $\text{Len}_l(\gamma|_{[s, t]}) > 0$ for every $0 \leq s < t \leq 1$ [87, Def. 2.29]. By definition of the length functional, every l -rectifiable curve is timelike [87, Lem. 2.30].

Proposition 2.37 (Reparameterizations of maximizers). *Every given causal curve $\gamma: [0, 1] \rightarrow \mathbb{M}$ which is l -rectifiable has a reparameterization $\bar{\gamma}: [0, \text{Len}_l(\gamma)] \rightarrow \mathbb{M}$ by l -arclength, i.e. for every $0 \leq t \leq \text{Len}_l(\gamma)$,*

$$\text{Len}_l(\bar{\gamma}|_{[0,t]}) = t.$$

In particular, every timelike maximizer $\gamma: [0, 1] \rightarrow \mathbb{M}$ has a proper-time reparameterization $\bar{\gamma}: [0, l(\gamma_0, \gamma_1)] \rightarrow \mathbb{M}$, i.e. for every $0 \leq s < t \leq l(\gamma_0, \gamma_1)$,

$$l(\bar{\gamma}_s, \bar{\gamma}_t) = t - s.$$

This naturally motivates the following definition.

Definition 2.38 (Geodesics). *A curve $\gamma: [0, 1] \rightarrow \mathbb{M}$ will be called l -geodesic if every $0 \leq s < t \leq 1$ satisfies*

$$l(\gamma_s, \gamma_t) = (t - s)l(\gamma_0, \gamma_1) > 0. \quad (2.4)$$

Remark 2.39 (About terminology). We point out the slight change of terminology of curves satisfying (2.4) to e.g. [21, 44] from “timelike geodesics” to “ l -geodesics”. This is in line with the nomenclature for probability measures in Definition 2.61. Moreover, it is debatable to call a rough null curve obeying the equality in (2.4) “geodesic” — every reparameterization again has the same property, and it is a priori unclear which parametrization of null maximizers is the “correct” one; see [92, Rem. 9] however. For this reason and since we mainly deal with curves satisfying (2.4), we have decided on this short-hand terminology.

For emphasis, we may term a curve satisfying (2.4) *affinely parametrized*. It is clear from the definition of Len_l that l -geodesics are maximizing.

For later use, we define the (a priori possibly empty) class of l -geodesics by

$$\begin{aligned} \text{TGeo}(\mathbb{M}) := \{ \gamma: [0, 1] \rightarrow \mathbb{M} : l(\gamma_s, \gamma_t) = (t - s)l(\gamma_0, \gamma_1) > 0 \\ \text{for every } 0 \leq s < t \leq 1 \}. \end{aligned} \quad (2.5)$$

2.5. Regularity. It is a priori unclear whether l -geodesics are continuous. Even worse, as remarked even their general existence may a priori fail. Both issues will be ruled out by our assumption of regularity introduced in this section.

2.5.1. Definition and a characterization.

Definition 2.40 (Regularity). *We term \mathbb{M} regular if every strictly causal maximizer [sic] $\gamma: [0, 1] \rightarrow \mathbb{M}$ with $\gamma_0 \ll \gamma_1$ is already timelike.*

Note the implicit continuity hypothesis in Definition 2.40 which is never void by Theorem 2.36.

In other words, Definition 2.40 asks every *maximizer* to be either timelike or null. By Lemma 2.3, the space \mathbb{M} is regular if and only if every maximizer $\gamma: [0, 1] \rightarrow \mathbb{M}$ with the property $\gamma_s \ll \gamma_t$ for some $0 \leq s < t \leq 1$ is already timelike. This notion is motivated by the main consequence of a similar concept from [87] called *regular localizability* [87, Def. 3.16], namely that every maximizer in a regularly localizable Lorentzian pre-length space à la [87] is either timelike or null [87, Thm. 3.18].

Hypothesis 2.41 (Second standing assumption). *From now on, in addition to Hypothesis 2.8 (length metric spacetime with compact diamonds), we assume \mathbb{M} to be regular according to Definition 2.40.*

As Hypothesis 2.41 is standing, for terminological convenience we henceforth mostly drop the addenda “globally hyperbolic”, “regular”, and “length”, and call the space under consideration simply a *metric spacetime*.

Proposition 2.37 provides the desired correspondence between geodesics and maximizers: every geodesic is clearly a maximizer, while every maximizer admits a reparametrization belonging to $\text{TGeo}(\mathbb{M})$.

For the convenience of the reader, we now provide an equivalent definition of our setting obeying Hypothesis 2.41 which collects the properties of metric spacetimes we most frequently use.

Theorem 2.42 (Characterization of reinforcement). *Let $l: \mathbb{M}^2 \rightarrow [0, \infty] \cup \{-\infty\}$ be a signed time separation function. The tuple (\mathbb{M}, l) is a globally hyperbolic, regular length metric spacetime if and only if the following conditions hold.*

- (i) **Nonempty chronology.** *For every $x \in \mathbb{M}$, the sets $I^\pm(x)$ are nonempty.*
- (ii) **Global hyperbolicity.** *The space \mathbb{M} has is globally hyperbolic according to the conclusion of Theorem 2.12.*
- (iii) **Regular timelike geodesy.** *Every points $x, y \in \mathbb{M}$ with $x \ll y$ can be connected by a continuous element of $\text{TGeo}(\mathbb{M})$ ³.*

We wrap up our discussion with examples obeying Hypothesis 2.41.

Example 2.43 (Lipschitz spacetimes). Let (\mathbb{M}, g) be a Lipschitz spacetime, i.e. \mathbb{M} is a connected smooth topological manifold endowed with a Lorentzian metric tensor g of Lipschitz regularity. Since \mathbb{M} is causally plain [49, Def. 1.16, Cor. 1.17], it naturally induces a length metric spacetime [87, Thm. 5.12]. By [69, Thm. 1.1], we know regularity of such spaces (and here Lipschitz regularity plays a crucial role). Consequently, the metric spacetime induced by a globally hyperbolic Lipschitz spacetime satisfies Hypothesis 2.41.

Example 2.44 (Convex neighborhoods). If \mathbb{M} admits convex neighborhoods in the sense of [98, Def. 6.10], then \mathbb{M} is regular [98, Prop. 6.11].

Example 2.45 (Causal reversal). We define the function $l^\leftarrow: \mathbb{M}^2 \rightarrow [0, \infty] \cup \{-\infty\}$ by $l^\leftarrow(x, y) := l(y, x)$. Evidently, this is a signed time separation function as in Definition 2.1, and the induced structure $(\mathbb{M}, l^\leftarrow)$ on \mathbb{M} is called *causal reversal* of (\mathbb{M}, l) . One readily verifies the former satisfies Hypothesis 2.41 if and only if the latter does. In particular, a path $\gamma: [0, 1] \rightarrow \mathbb{M}$ belongs to $\text{TGeo}(\mathbb{M})$ according to (2.5) if and only if its time reversal $\gamma^\leftarrow: [0, 1] \rightarrow \mathbb{M}$ with $\gamma_t^\leftarrow := \gamma_{1-t}$ belongs to

$$\begin{aligned} \text{TGeo}^\leftarrow(\mathbb{M}) := \{ \sigma: [0, 1] \rightarrow \mathbb{M} : l^\leftarrow(\sigma_s, \sigma_t) = (t - s) l^\leftarrow(\sigma_0, \sigma_1) > 0 \\ \text{for every } 0 \leq s < t \leq 1 \}. \end{aligned}$$

2.5.2. *Continuity of l -geodesics.* The additional assumption of regularity will be useful to “lift” the geometry of the base space \mathbb{M} to the space of Borel probability measures over it, cf. Proposition 2.66 and [21, 22, 92]. As a first step, we note the following fact established in [92], whose proof we repeat for convenience.

Recall (2.5) for the definition of $\text{TGeo}(\mathbb{M})$.

Lemma 2.46 (Continuity of l -geodesics). *All elements of $\text{TGeo}(\mathbb{M})$ are continuous.*

Proof. We first show right-continuity of a given $\gamma \in \text{TGeo}(\mathbb{M})$; left-continuity is argued analogously.

First, let $0 < s < 1$, suppose to the contrary that γ is not right-continuous at s . By compactness of $J(\gamma_0, \gamma_1)$, this means the existence of a sequence $(\varepsilon_n)_{n \in \mathbb{N}}$ of numbers $0 < \varepsilon_n < 1 - s$ such that $(\gamma_{s+\varepsilon_n})_{n \in \mathbb{N}}$ converges to some $\gamma_{s+} \in \mathbb{M}$ with $\gamma_{s+} \neq \gamma_s$. Proposition 2.10 shows $\gamma_s < \gamma_{s+}$. Since $\gamma_0 \ll \gamma_s$ and $\gamma_{s+} \ll \gamma_1$ by the definition of $\text{TGeo}(\mathbb{M})$ and continuity of l_+ , concatenating

³By Lemma 2.46 below, every rough strictly causal maximizer from x to y is timelike and continuous.

- a strictly causal maximizer $\sigma^1: [0, 1/3] \rightarrow M$ from γ_0 to γ_s ,
- a strictly causal maximizer $\sigma^2: [1/3, 2/3] \rightarrow M$ from γ_s to γ_{s+} , and
- a strictly causal maximizer $\sigma^3: [2/3, 1] \rightarrow M$ from γ_{s+} to γ_1

yields a strictly causal curve $\sigma: [0, 1] \rightarrow M$ from γ_0 to γ_1 . Then

$$\begin{aligned} l(\gamma_0, \gamma_1) &\geq \text{Len}_l(\sigma) \\ &= \text{Len}_l(\sigma^1) + \text{Len}_l(\sigma^2) + \text{Len}_l(\sigma^3) \\ &= l(\gamma_0, \gamma_s) + l(\gamma_s, \gamma_{s+}) + l(\gamma_{s+}, \gamma_1) \\ &= l(\gamma_0, \gamma_1); \end{aligned}$$

here, in the second line, we have used Lemma 2.6, while the last line employs the definition of $\text{TGeo}(M)$. In particular, σ is a strictly causal maximizer from γ_0 to γ_1 . On the other hand, affine parametrization implies $l(\gamma_s, \gamma_{s+}) = 0$. Thus, σ has a null subsegment, in contradiction with regularity.

Right-continuity at $s = 0$ is argued similarly. Here we directly start with a null strictly causal maximizer from γ_0 to γ_{0+} and continue with a strictly causal maximizer connecting the chronologically related points γ_{0+} and γ_1 . \square

Recall Theorem 2.18 and Remark 2.31 for how $C([0, 1]; M)$ appearing below can be naturally and unambiguously endowed with a topology of uniform convergence on compact subsets of M , even if τ is not globally metrizable.

The next result can be interpreted as a limit curve theorem for l -geodesics. By Lemma 2.46, its (independent) proof is somewhat more straightforward than that of the limit curve theorem (Theorem 2.32) for general strictly causal curves.

Corollary 2.47 (Compactness of uniformly l -geodesics). *Given any $r > 0$ and any compact sets $C_0, C_1 \subset M$, the set*

$$G_r := \{\gamma \in \text{TGeo}(M) : \gamma_0 \in C_0, \gamma_1 \in C_1, l(\gamma_0, \gamma_1) \geq r\}$$

is compact with respect to uniform convergence in $C([0, 1]; M)$.

Proof. Since the set G_r is closed by continuity of l_+ , it suffices to prove its precompactness. By Lemma 2.46, we have $G_r \subset C([0, 1]; M)$, so that the statement makes sense. Let $(\gamma^n)_{n \in \mathbf{N}}$ be a sequence in G_r .

Following an idea developed in [14], we construct a candidate $\gamma: [0, 1] \rightarrow M$ for a uniform limit of $(\gamma^n)_{n \in \mathbf{N}}$ along a nonrelabeled subsequence. By compactness, a diagonal procedure — recalling that $J(C_0, C_1)$ is Polish by Theorem 2.18 — and Proposition 2.10 there exists a nonrelabeled subsequence of $(\gamma^n)_{n \in \mathbf{N}}$ such that for every $t \in [0, 1] \cap \mathbf{Q}$, there exists $\gamma_t \in J(C_0, C_1)$ such that $\gamma_t^n \rightarrow \gamma_t$ as $n \rightarrow \infty$. The rest of the candidate γ is constructed by approximation from the right. For $t \in [0, 1] \setminus \mathbf{Q}$, let $(t_i)_{i \in \mathbf{N}}$ and $(s_j)_{j \in \mathbf{N}}$ be sequences in $(t, 1] \cap \mathbf{Q}$ which converge to t . Up to passing to subsequences, we may and will assume the existence of the limits $x_t := \lim_{i \rightarrow \infty} \gamma_{t_i}$ and $y_s := \lim_{j \rightarrow \infty} \gamma_{s_j}$, which evidently lie in $J(C_0, C_1)$. We claim $x_t = y_s$. Indeed, given any $i \in \mathbf{N}$ we have $\gamma_{s_j} \leq \gamma_{t_i}$ for every large enough $j \in \mathbf{N}$, and thus $y_s \leq x_t$ by first letting $j \rightarrow \infty$ and then $i \rightarrow \infty$. Analogously, we obtain $x_t \leq y_s$, and the claim follows from causality. In particular, this argument shows the limit γ_t can be unambiguously defined as $\lim_{i \rightarrow \infty} \gamma_{t_i}$, where $(t_i)_{i \in \mathbf{N}}$ is any sequence as above. This finishes the construction of γ .

Since $l(\gamma_0^n, \gamma_1^n) \geq r$ for every $n \in \mathbf{N}$, clearly $\gamma \in \text{TGeo}(M)$ — in particular, γ is continuous by Lemma 2.46.

Lastly, we claim uniform convergence of $(\gamma^n)_{n \in \mathbf{N}}$ to γ . Let d metrize the relative topology on $J(C_0, C_1)$. Given any $\varepsilon > 0$, for every $0 \leq t \leq 1$, by Proposition 2.28 (strong causality) there is an open neighborhood $V_t \subset M$ of γ_t such that every strictly causal curve with endpoints in V_t does not leave the metric ball $\mathbf{B}^d(\gamma_t, \varepsilon)$. Up to

shrinking V_t , we may and will assume it to be an open d -ball of radius $0 < \delta_t < \varepsilon$. We cover the compact set $\gamma_{[0,1]}$ by finitely many of these sets V_{t_1}, \dots, V_{t_m} , where $m \in \mathbf{N}$, and t_1, \dots, t_m may and will be chosen to be rational. By the compactness of $[0, 1]$ and uniform continuity of γ , we may and will further assume that $|t_i - t_{i+1}| \leq \delta$ for every $i = 1, \dots, m-1$, and that $0 = t_1 < t_2 < \dots < t_{m-1} < t_m = 1$ is an increasing δ -net in $[0, 1]$, where $\delta > 0$ is such that every $0 \leq s, t \leq 1$ with $|s - t| \leq \delta$ satisfies $d(\gamma_s, \gamma_t) \leq \varepsilon$. By our choice of t_1, \dots, t_m , we have

$$\sup_{t \in [t_i, t_{i+1}]} d(\gamma_{t_i}, \gamma_t) \leq 2\varepsilon. \quad (2.6)$$

Moreover, by pointwise convergence of $(\gamma^n)_{n \in \mathbf{N}}$ to γ on the rationals, there exists $n_i \in \mathbf{N}$ such that $\gamma_{t_i}^n \in V_{t_i}$ and $\gamma_{t_{i+1}}^n \in V_{t_{i+1}}$ for every $n \in \mathbf{N}$ with $n \geq n_i$. Since $t_1 < t_2 < \dots < t_{m-1} < t_m$, we have $V_{t_i} \cap V_{t_{i+1}} \neq \emptyset$. Therefore, for every $n \geq n_i$ the restriction $\gamma^n|_{[t_i, t_{i+1}]}$ is a strictly causal curve which starts in V_{t_i} , passes through $V_{t_i} \cap V_{t_{i+1}} \neq \emptyset$, and ends in $V_{t_{i+1}}$; strong causality then easily implies

$$\sup_{t \in [t_i, t_{i+1}]} d(\gamma_t^n, \gamma_{t_i}) \leq 2\varepsilon. \quad (2.7)$$

Whenever $n \geq \max\{n_1, \dots, n_{m-1}\}$, combining (2.6) and (2.7) finally yields

$$\sup_{t \in [0, 1]} d(\gamma_t^n, \gamma_t) \leq 4\varepsilon. \quad \square$$

2.5.3. Intermediate points. In order to study geodesics of probability measures later, given any $x, y \in \mathbf{M}$ and any $0 \leq t \leq 1$, and following [91, Ch. 2], we now investigate the regularity of the (possibly empty) t -intermediate point set

$$Z_t(x, y) := \{z \in J(x, y) : l(x, z) = t l(x, y), l(z, y) = (1 - t) l(x, y)\}.$$

Moreover, for $C \subset \mathbf{M}^2$ we define

$$\begin{aligned} Z_t(C) &:= \bigcup_{(x, y) \in C} Z_t(x, y), \\ Z(C) &:= \bigcup_{t \in [0, 1]} Z_t(C). \end{aligned} \quad (2.8)$$

Remark 2.48 (Initial and final points). If $x, y \in \mathbf{M}$ obey $x \ll y$ then $Z_0(x, y) = \{x\}$ [91, Rem. 2.6]. Indeed, if $x' \in Z_0(x, y) \setminus \{x\}$ then $l(x, x') = 0$, and together with $l(x', y) = l(x, y) > 0$ this implies the concatenation of a null maximizer from x to x' and an l -geodesic from x' to y to be a maximizer from x to y which intermediately changes its causal character from null to timelike. This contradicts regularity of \mathbf{M} . Analogously, we show $Z_1(x, y) = \{y\}$ if $l(x, y) > 0$.

The following is a nonsmooth analog to [91, Lem. 2.5].

Lemma 2.49 (Intermediate point sets inherit compactness I). *Let $0 \leq t \leq 1$. If $C \subset \mathbf{M}^2$ is precompact, so are $Z_t(C)$ and $Z(C)$. Moreover, if C is compact, so are $Z_t(C)$ and $Z(C)$.*

Proof. Let $C_1, C_2 \subset \mathbf{M}$ be the projections of C to the first and second coordinate, respectively. We first assume precompactness of C . By its definition, $Z(C)$ is a subset of $J(\bar{C}_1, \bar{C}_2)$ which is compact by Corollary 2.11, and hence $Z(C)$ is precompact. In turn, this implies precompactness of $Z_t(C) \subset Z(C)$.

The same argument combined with the continuity of l_+ and Corollary 2.14 gives the second statement provided C is compact. \square

2.6. Lorentzian optimal transport. Next, we review the optimal transport problem in our Lorentzian setting. We refer to [44, 51, 78, 91, 99, 125] for details. Their results carry over with no essential change to our slightly different setting.

From now on, we insist on τ being Polish. As outlined in Remark 2.29, under mild conditions this can be relaxed by working on a compact exhaustion, but to streamline the presentation we directly assume Polishness of τ . Whenever needed, d will denote a complete and separable metric inducing τ .

Fix a Radon measure \mathfrak{m} on M . To avoid pathologies, we assume \mathfrak{m} is nontrivial, but we do *not* assume \mathfrak{m} has full support in general. By Remark 2.29, \mathfrak{m} is σ -finite (which has to be assumed if τ is merely first-countable). We often write

$$\mathcal{M} := (M, l, \mathfrak{m}). \quad (2.9)$$

Definition 2.50 (Metric measure spacetime). *The triple \mathcal{M} in (2.9), i.e. a regular Polish length metric spacetime (M, l) with compact diamonds according to our second standing Hypothesis 2.41 endowed with a nontrivial Radon measure \mathfrak{m} , is called a metric measure spacetime.*

Also, we write

$$\mathcal{M}^{\leftarrow} := (M, l^{\leftarrow}, \mathfrak{m})$$

for the structure induced by the causal reversal from Example 2.45.

2.6.1. Suslin measurability. In Chapter 6 we need to discuss Suslin measurability, basic notions of which are briefly summarized here. See [118] for details.

A subset E of a Polish space X is called *Suslin measurable*, or briefly *Suslin*, if it is the continuous image of a Borel subset of a Polish space (possibly different from X). In the last clause, one can equivalently remove ‘‘Borel subset of a’’. Another characterization of Suslin measurability is that E is the projection of a Borel subset of a Polish product space, i.e. there are a Polish space Y and a Borel set $F \subset X \times Y$ such that $E = \text{pr}_1(F)$, where $\text{pr}_1: X \times Y \rightarrow X$ is defined as $\text{pr}_1(x, y) := x$.

The crucial property we use later is *universal measurability* of all Suslin sets. Here, a subset E of a Polish space X is universally measurable if, for every Borel probability measure μ on X , it is measurable with respect to the completion of the Borel σ -field of X with respect to μ .

2.6.2. Notation from probability theory. Let $\mathcal{P}(M)$ denote the space of all Borel probability measures on M , endowed with the narrow topology, as detailed below in Subsection 2.6.3. The Dirac measure at $x \in M$ is denoted δ_x . Let $\mathcal{P}_c(M)$ and $\mathcal{P}^{\text{ac}}(M, \mathfrak{m})$ denote the spaces of compactly supported and \mathfrak{m} -absolutely continuous measures in $\mathcal{P}(M)$, respectively. We set $\mathcal{P}_c^{\text{ac}}(M, \mathfrak{m}) := \mathcal{P}_c(M) \cap \mathcal{P}^{\text{ac}}(M, \mathfrak{m})$. When writing $\mu = \rho \mathfrak{m} + \mu_{\perp}$ for a given $\mu \in \mathcal{P}(M)$, we mean the Lebesgue decomposition of μ with respect to \mathfrak{m} , where ρ is the Radon–Nikodým density of its \mathfrak{m} -absolutely continuous part, and μ_{\perp} is its \mathfrak{m} -singular part.

The following terminology will be useful in the sequel. It has an evident analog for sequences of probability measures.

Definition 2.51 (Uniformly compact support). *A family $\mathcal{C} \subset \mathcal{P}(M)$ is called uniformly compactly supported if there exists a compact subset $C \subset M$ such that $\text{spt } \mu \subset C$ for every $\mu \in \mathcal{C}$.*

The *push-forward* of a measure $\mu \in \mathcal{P}(M)$ under a Borel map $F: M \rightarrow M'$, where M' is any given topological space, is the probability measure $F_{\#}\mu \in \mathcal{P}(M')$ defined by setting, for every Borel set $B \subset M'$,

$$F_{\#}\mu[B] := \mu[F^{-1}(B)].$$

A measure $\pi \in \mathcal{P}(M^2)$ is called *coupling* of given $\mu, \nu \in \mathcal{P}(M)$ if

$$\begin{aligned} (\text{pr}_1)_\# \pi &= \mu, \\ (\text{pr}_2)_\# \pi &= \nu. \end{aligned}$$

Here and in the sequel, given any $i = 1, 2$, $\text{pr}_i: M^2 \rightarrow M$ denotes the projection $\text{pr}_i(x) := x_i$; the projection $\text{pr}_i: M^n \rightarrow M$ is defined analogously for arbitrary $n \in \mathbf{N}$ and $i = 1, \dots, n$. Let $\Pi(\mu, \nu)$ be the set of all couplings of μ and ν . It is never empty as it contains at least the product measure $\mu \otimes \nu \in \mathcal{P}(M^2)$. Moreover, by *diagonal coupling* of μ we mean the probability measure $\pi := D_\# \mu$, where $D: M \rightarrow M^2$ is defined by $D(x) := (x, x)$.

We recall the following well-known result, cf. e.g. [7, Lem. 2.1].

Lemma 2.52 (Gluing). *Given any $n \in \mathbf{N}$, let $\mu_0, \dots, \mu_n \in \mathcal{P}(M)$, and let $\pi_{i-1, i} \in \Pi(\mu_{i-1}, \mu_i)$ for every $i = 1, \dots, n$. Then there exists a measure $\omega \in \mathcal{P}(M^{n+1})$ such that every $i = 1, \dots, n$ satisfies*

$$(\text{pr}_{i-1}, \text{pr}_i)_\# \omega = \pi_{i-1, i}.$$

We call a set $\mathcal{C} \subset \mathcal{P}(M)$ *tight* if for every $\varepsilon > 0$ there exists a compact set $C \subset M$ such that $\mu[C^c] \leq \varepsilon$ for every $\mu \in \mathcal{C}$, where $C^c := M \setminus C$ denotes set complementation; we also agree on the evident meaning of a tight sequence. Tightness is a useful property in view of Prokhorov's theorem (Theorem 2.54 below). Recall that tightness transfers to couplings [129, Lem. 4.4]: if $\mathcal{C}_0, \mathcal{C}_1 \subset \mathcal{P}(M)$ are tight, so is the set $\Pi(\mathcal{C}_0 \times \mathcal{C}_1)$ of all couplings of elements of \mathcal{C}_0 and \mathcal{C}_1 . This argument does not require τ to be Polish. The converse implication holds as well by the continuity of the projection maps and Prokhorov's theorem.

Lastly, let $C([0, 1]; M)$ be the space of continuous maps $\gamma: [0, 1] \rightarrow M$, endowed with the topology of uniform convergence. For $0 \leq t \leq 1$, define the continuous *evaluation map* $\mathbf{e}_t: C([0, 1]; M) \rightarrow M$ by

$$\mathbf{e}_t(\gamma) := \gamma_t. \quad (2.10)$$

The continuous *restriction map* $\text{restr}_s^t: C([0, 1]; M) \rightarrow C([0, 1]; M)$, where $0 \leq s < t \leq 1$, is the map which restricts $\gamma \in C([0, 1]; M)$ to the segment on $[s, t]$ and then “stretches” this segment affinely to $[0, 1]$, i.e.

$$\text{restr}_s^t(\gamma)_r := \gamma_{(1-r)s+rt}. \quad (2.11)$$

2.6.3. Narrow convergence. A sequence $(\mu_n)_{n \in \mathbf{N}}$ in $\mathcal{P}(M)$ converges *narrowly*⁴ to $\mu \in \mathcal{P}(M)$ [17, p. 13] if every continuous bounded test function $\varphi \in C_b(M)$ satisfies

$$\lim_{n \rightarrow \infty} \int_M \varphi d\mu_n = \int_M \varphi d\mu.$$

When τ is Polish, narrow convergence is induced by a metric [8, Rem. 5.1.1]. In the more general case, narrow convergence of directed nets is defined in the evident way.

By continuity of the projection, it is clear that the property of being a coupling is narrowly closed with respect to narrow convergence of the marginals.

Given an interval $I \subset \mathbf{R}$, by narrow continuity of a collection $(\mu_t)_{t \in I}$ in $\mathcal{P}(M)$ we mean continuity of the map $t \mapsto \mu_t$ on I with respect to narrow convergence.

We recall two results on convergence and compactness in the narrow topology we occasionally employ. The first is Alexandrov's theorem⁵ [17, Thm. 2.1],

⁴Some authors prefer to call this *weak* convergence; note that for noncompact M , this usage of “weak” is inconsistent with its meaning in Banach space theory.

⁵This theorem seems to have been proven first in [5]. A more common name for it is *Portmanteau's theorem*, a terminology originating in [17]. The references of [17] quote the article “Espoir pour l'ensemble vide?”, viz. “Hope for the empty set?”, of Jean-Pierre Portmanteau in the *Annales de l'Université de Fellein*. This appears to be a hoax, since we find no historical record

which relates narrow convergence to convergence of set evaluations. The second is Prokhorov's theorem [17, Thms. 5.1, 5.2], which is a useful characterization of narrow precompactness in $\mathcal{P}(M)$.

Theorem 2.53 (Alexandrov's theorem). *Narrow convergence of a sequence $(\mu_n)_{n \in \mathbf{N}}$ in $\mathcal{P}(M)$ to $\mu \in \mathcal{P}(M)$ is equivalent to any of the following.*

(i) **Lower semicontinuity on open sets.** *For every open set $U \subset M$,*

$$\mu[U] \leq \liminf_{n \rightarrow \infty} \mu_n[U].$$

(ii) **Upper semicontinuity on closed sets.** *For every closed set $C \subset M$,*

$$\mu[C] \geq \limsup_{n \rightarrow \infty} \mu_n[C].$$

(iii) **Continuity sets.** *For every Borel set $E \subset M$ with $\mu[\partial E] = 0$,*

$$\mu[E] = \lim_{n \rightarrow \infty} \mu_n[E].$$

Theorem 2.54 (Prokhorov's theorem). *A given set $\mathcal{C} \subset \mathcal{P}(M)$ is narrowly precompact if and only if it is tight.*

Remark 2.55 (Extension beyond Polish spaces). Even if τ is not Polish, Theorem 2.53 and Theorem 2.54 extend to merely first-countable τ in our setting, and at this point there is no need to restrict ourselves to uniformly compactly supported measures. The crucial property is *complete regularity* of τ , i.e. for every closed set $C \subset M$ and every $x \in C^c$ in its complement C^c , there exists a continuous function $f: M \rightarrow \mathbf{R}$ vanishing on C and $f(x) = 1$. Every locally compact Hausdorff topology is completely regular, hence τ is by Corollary 2.17.

Theorem 2.53 then follows from [18, Thm. 8.2.3, Cor. 8.2.4].

Theorem 2.54 is found in [18, Thm. 8.6.7]. Indeed, Theorem 2.18 implies every compact subset of M is metrizable.

2.6.4. *Causal and chronological couplings.* Given any $\mu, \nu \in \mathcal{P}(M)$, the (possibly empty) set $\Pi_{\leq}(\mu, \nu)$ of *causal couplings* of μ and ν is the set of all $\pi \in \Pi(\mu, \nu)$ with $\pi[l \geq 0] = 1$. Since $\{l \geq 0\}$ is closed, the latter condition holds if and only if $\text{spt } \pi \subset \{l \geq 0\}$. By replacing “ \leq ” and “ \geq ” by “ \ll ” and “ $>$ ”, respectively, in an analogous way we define the set $\Pi_{\ll}(\mu, \nu)$ of *chronological couplings* of μ and ν . However, note that the property $\pi \in \Pi_{\ll}(\mu, \nu)$ does not imply $\text{spt } \pi \subset \{l > 0\}$ in general, while the converse is trivially true.

2.6.5. *The ℓ_p -optimal transport problem.* Given an exponent $0 < p \leq 1$, besides (2.1) we will adopt the conventions

$$(-\infty)^p := (-\infty)^{1/p} := -\infty.$$

We define the function $\ell_p: \mathcal{P}(M)^2 \rightarrow [0, \infty] \cup \{-\infty\}$ by

$$\ell_p(\mu, \nu) := \sup_{\pi \in \Pi_{\leq}(\mu, \nu)} \|l_+\|_{L^p(M^2, \pi)}, \quad (2.12)$$

and call it total transport cost. We agree on the usual convention $\sup \emptyset := -\infty$, in which case we can equivalently use an unconditional supremum to infer

$$\ell_p(\mu, \nu) = \sup_{\pi \in \Pi(\mu, \nu)} \|l\|_{L^p(M^2, \pi)}. \quad (2.13)$$

The sets of maximizers of (2.12) and (2.13) coincide [44, Rem. 2.2]. A coupling $\pi \in \Pi(\mu, \nu)$ is called *ℓ_p -optimal* if it is causal and it attains the supremum in the

of mathematician of that name, nor of the University of Felletin (let alone the mentioned paper). To acknowledge Alexandrov's contribution, we have decided to call Theorem 2.53 “Alexandrov's theorem”, as suggested by Bogachev [18, p. 453].

definition of ℓ_p ; if $\ell_p(\mu, \nu) = -\infty$ no coupling of μ and ν is ℓ_p -optimal according to this terminology, even though $\ell_p(\mu, \nu) = \|l\|_{L^p(\mathbb{M}^2, \pi)}$ holds for every $\pi \in \Pi(\mu, \nu)$. On the other hand, by our definition every $\pi \in \mathcal{P}(\mathbb{M}^2)$ with $\|l\|_{L^p(\mathbb{M}^2, \pi)} = \infty$ is an ℓ_p -optimal coupling of its marginals.

Standard arguments of calculus of variations, combined with the upper boundedness of ℓ_+ on compact subsets of \mathbb{M}^2 , ensure the existence of ℓ_p -optimal couplings in all cases relevant for our purposes, cf. [44, Prop. 2.3] and [129, Thm. 4.1].

Lemma 2.56 (Existence of ℓ_p -optimal couplings). *If two measures $\mu, \nu \in \mathcal{P}_c(\mathbb{M})$ obey $\Pi_{\leq}(\mu, \nu) \neq \emptyset$, then μ and ν admit an ℓ_p -optimal coupling.*

Another feature is the reverse triangle inequality for ℓ_p akin to its counterpart (2.2) for l [44, Prop. 2.5], which is a standard application of the gluing lemma: for every $\mu, \nu, \sigma \in \mathcal{P}(\mathbb{M})$, we have

$$\ell_p(\mu, \sigma) \geq \ell_p(\mu, \nu) + \ell_p(\nu, \sigma). \quad (2.14)$$

These observations yield interesting facts about the causal structure of the space of probability measures on \mathbb{M} , which generalize those from [51] to our nonsmooth setting. We defer this discussion to Appendix B.

2.6.6. Timelike p -dualizability. Even if $\ell_p(\mu, \nu) > 0$ for two given $\mu, \nu \in \mathcal{P}(\mathbb{M})$, in general ℓ_p -optimal couplings need not vanish outside set $\{l > 0\}$, though they cannot vanish inside it (e.g. transport a Dirac mass halfway to a point in its null future and halfway to a point in its chronological future in Minkowski spacetime). Various notions for μ and ν to have “entirely chronological” ℓ_p -optimizers have been studied in the literature. Starting from McCann’s p -separation [91, Def. 4.1], Cavalletti–Mondino [44, Defs. 2.18, 2.27] introduced the notions of (strong) timelike p -dualizability recapitulated in Definition 2.57 as nonsmooth analogs for the more relaxed criteria of [91, Thms. 7.1, 7.4]. They allow for a good duality theory [44, Props. 2.19, 2.21, Thm. 2.26], see also [91, Thm. 4.3].

As usual, given two functions $a, b: \mathbb{M} \rightarrow \mathbf{R}$ we define $a \oplus b: \mathbb{M}^2 \rightarrow \mathbf{R}$ by

$$(a \oplus b)(x, y) := a(x) + b(y).$$

Definition 2.57 (Timelike p -dualizability). *Let $\mu, \nu \in \mathcal{P}(\mathbb{M})$ be given.*

- a. *We call the pair (μ, ν) timelike p -dualizable if $\ell_p(\mu, \nu) > 0$, some ℓ_p -optimal coupling of μ and ν belongs to $\Pi_{\ll}(\mu, \nu)$, and there are lower semicontinuous functions $a, b: \mathbb{M} \rightarrow [0, \infty)$ with $a \oplus b \in L^1(\mathbb{M}^2, \mu \otimes \nu)$ and*

$$l^p \leq a \oplus b \quad \text{on } \text{spt } \mu \times \text{spt } \nu. \quad (2.15)$$

If $\pi \in \Pi_{\ll}(\mu, \nu)$ is such a coupling, we call (μ, ν) timelike p -dualizable by π , and the coupling π timelike p -dualizing.

- b. *Furthermore, we call the pair (μ, ν) strongly timelike p -dualizable if it is timelike p -dualizable, and there exists an l^p -cyclically monotone Borel set $\Gamma \subset \{l > 0\} \cap (\text{spt } \mu \times \text{spt } \nu)$ such that every ℓ_p -optimal coupling $\pi \in \Pi_{\leq}(\mu, \nu)$ is concentrated on Γ .*

Here, we call a set $\Gamma \subset \{l \geq 0\}$ l^p -cyclically monotone [44, Def. 2.6] if for every $n \in \mathbf{N}$ and every $(x_1, y_1), \dots, (x_n, y_n) \in \Gamma$, we have

$$\sum_{i=1}^n l^p(x_i, y_i) \geq \sum_{i=1}^n l^p(x_{i+1}, y_i),$$

where $x_{n+1} := x_1$. Accordingly, an element of $\mathcal{P}(\mathbb{M}^2)$ is l^p -cyclically monotone if it is concentrated on an l^p -cyclically monotone subset of $\{l \geq 0\}$.

Note that (2.15) forces $\ell_p(\mu, \nu) < \infty$ if (μ, ν) is timelike p -dualizable.

Clearly, strong timelike p -dualizability implies timelike p -dualizability. Moreover, any set Γ as in the preceding item **b.** actually characterizes ℓ_p -optimality. Indeed, if $\pi \in \Pi(\mu, \nu)$ satisfies $\pi[\Gamma] = 1$, it is chronological and l^p -cyclically monotone, hence ℓ_p -optimal by the hypothesized moment conditions (2.15) [44, Prop. 2.8].

We will mainly need these notions for compactly supported measures, in which case the definitions become easier to handle.

Remark 2.58 (Compact support simplifies timelike p -dualizability). Assume $\mu, \nu \in \mathcal{P}_c(M)$. Then (μ, ν) is timelike p -dualizable if and only if they admit a timelike p -dualizing coupling $\pi \in \Pi_{\ll}(\mu, \nu)$, i.e. a *chronological* ℓ_p -optimal coupling.

Example 2.59 (Chronological product measures). If $\mu, \nu \in \mathcal{P}_c(M)$ satisfy

$$\text{spt } \mu \times \text{spt } \nu \subset \{l > 0\}, \quad (2.16)$$

the pair (μ, ν) is strongly timelike p -dualizable for every $0 < p \leq 1$ [44, Cor. 2.29]. However, unlike strong timelike p -dualizability, cf. Lemma 2.67 below, the condition (2.16) does not propagate to the interior of ℓ_p -geodesics.

Remark 2.60 (Timelike p -dualizability in the smooth framework). The p -separation of $(\mu, \nu) \in \mathcal{P}_c(M)^2$ [91, Def. 4.1] implies strong timelike p -dualizability of (μ, ν) . Indeed, by definition the set S therein is contained in $\{l > 0\} \cap (\text{spt } \mu \times \text{spt } \nu)$, and it is evidently l^p -cyclically monotone. Moreover, standard Kantorovich duality [129, Thm. 5.10] ensures every ℓ_p -optimal coupling of μ and ν to be concentrated on S . In the smooth setting, (strong) timelike p -dualizability is encoded in the hypotheses of [91, Thm. 7.1], which does not assume p -separation.

2.7. Geodesics of probability measures. Following [91, Def. 1.1] we now introduce various notions of ℓ_p -geodesics, and study their properties in detail. Key parts of our analysis are sufficient conditions for

- their narrow continuity, and
- their admittance of a “lifting” to the space of Borel probability measures concentrated on affinely parametrized geodesics on M .

We refer to [14, 21, 22, 44, 92, 99] for further treatments of related notions; in a certain sense, our following discussion continues the study of the Lorentzian structure of $\mathcal{P}(M)$ initiated in [51] for smooth spacetimes. We point out that the notion of ℓ_p -geodesics in [21, 22, 44] is slightly more restrictive than that of [91]. Roughly speaking, in [21, 22, 44] ℓ_p -geodesics are represented by probability measures concentrated on appropriate curves, and hence given by liftings *by definition*. As we do not have Lipschitz causal curves at our disposal, some desired properties need to be reworked; especially compactness arguments become trickier. A byproduct of our results is that the restrictions from [21, 22, 44] are not really necessary, and a much weaker notion is sufficient to work with under Hypothesis 2.41. Also, Definition 2.61 is more flexible with respect to convergence questions, cf. e.g. Theorem 2.68.

2.7.1. Various notions of geodesics. Recall the definition of the evaluation map e_t from (2.10), where $0 \leq t \leq 1$. We say a collection $(\mu_t)_{t \in [0,1]}$ in $\mathcal{P}(M)$ is *represented* by $\pi \in \mathcal{P}(C([0,1]; M))$ if every $0 \leq t \leq 1$ satisfies

$$\mu_t = (e_t)_\# \pi.$$

We endow the path space $\mathcal{P}(C([0,1]; M))$ with its natural narrow topology. Given any $\mu_0, \mu_1 \in \mathcal{P}(M)$, we define the (possibly empty) class

$$\text{OptTGeo}_p(\mu_0, \mu_1) := \{\pi \in \mathcal{P}(C([0,1]; M)) : (e_0, e_1)_\# \pi \in \Pi(\mu_0, \mu_1) \text{ is } \ell_p\text{-optimal, } \pi[\text{TGeo}(M)] = 1\}. \quad (2.17)$$

Its elements will be called *timelike ℓ_p -optimal dynamical plans*, or simply *plans*; a plan can also be thought of as a traffic plan or a routing. This definition makes sense since $\text{TGeo}(\mathbb{M})$ is a Borel subset of $C([0, 1]; \mathbb{M})$: it is a σ -compact subset of $C([0, 1]; \mathbb{M})$ by Corollary 2.47.

Definition 2.61 (Geodesics of probability measures). *We call a curve $(\mu_t)_{t \in [0, 1]}$ of elements in $\mathcal{P}(\mathbb{M})$*

- a. *rough ℓ_p -geodesic if for every $0 \leq s < t \leq 1$,*

$$0 < \ell_p(\mu_s, \mu_t) = (t - s) \ell_p(\mu_0, \mu_1) < \infty, \quad (2.18)$$

- b. *ℓ_p -geodesic if it is a narrowly continuous rough ℓ_p -geodesic, and*
c. *displacement ℓ_p -geodesic if it is represented by a timelike ℓ_p -optimal dynamical plan $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$.*

Remark 2.62 (Displacement ℓ_p -geodesics are ℓ_p -geodesics). Every displacement ℓ_p -geodesic is an ℓ_p -geodesic. Indeed, let a curve $(\mu_t)_{t \in [0, 1]}$ be represented by a plan $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$. By Lebesgue's theorem and Lemma 2.46, it is clear that $(\mu_t)_{t \in [0, 1]}$ is narrowly continuous. Moreover, $(\mathbf{e}_0, \mathbf{e}_1)_{\#} \pi$ is an ℓ_p -optimal coupling of μ_0 and μ_1 . To show (2.18), let $0 \leq s < t \leq 1$. Note that $(\mathbf{e}_s, \mathbf{e}_t)_{\#} \pi$ is a causal coupling of μ_s and μ_t , and thus

$$\begin{aligned} \ell_p(\mu_s, \mu_t) &\geq \|l \circ (\mathbf{e}_s, \mathbf{e}_t)\|_{L^p(\text{TGeo}(\mathbb{M}), \pi)} \\ &= (t - s) \|l \circ (\mathbf{e}_0, \mathbf{e}_1)\|_{L^p(\text{TGeo}(\mathbb{M}), \pi)} \\ &= (t - s) \ell_p(\mu_0, \mu_1). \end{aligned} \quad (2.19)$$

This together with the reverse triangle inequality (2.2) implies

$$\begin{aligned} \ell_p(\mu_0, \mu_1) &\geq \ell_p(\mu_0, \mu_s) + \ell_p(\mu_s, \mu_t) + \ell_p(\mu_s, \mu_t) \\ &\geq [s + (t - s) + (1 - t)] \ell_p(\mu_0, \mu_1) \\ &= \ell_p(\mu_0, \mu_1). \end{aligned}$$

Consequently, equality holds throughout (2.19), and $(\mathbf{e}_s, \mathbf{e}_t)_{\#} \pi$ couples μ_s and μ_t optimally with respect to ℓ_p .

2.7.2. Narrow continuity and equivalence. The previous Remark 2.62 establishes a hierarchy between the various types of ℓ_p -geodesic (Definition 2.61). Under natural hypotheses, these are in fact equivalent, which is the content of Corollary 2.64. The latter is partly a measure version of Lemma 2.46.

This result requires the technical Lemma 2.63. Like [91, Cor. 2.10], it follows from the characterization of equality cases in the reverse triangle inequality (2.14) from [91, Prop. 2.9] as well as Lemmas 2.56 and 2.49.

Lemma 2.63 (Interpolants inherit compact support). *Let $(\mu_t)_{t \in [0, 1]}$ be a rough ℓ_p -geodesic, and let $0 < t < 1$. If μ_0 and μ_1 are compactly supported, so is μ_t . More precisely, its support $\text{spt } \mu_t \subset Z_t(\text{spt } \mu_0 \times \text{spt } \mu_1)$ is compact by Lemma 2.49.*

Corollary 2.64 (Chronology makes rough ℓ_p -geodesics narrowly continuous). *Let $(\mu_t)_{t \in [0, 1]}$ be a rough ℓ_p -geodesic with compactly supported endpoints μ_0 and μ_1 , where $0 < p < 1$. Suppose every ℓ_p -optimal coupling of μ_0 and μ_1 to be chronological. Then $(\mu_t)_{t \in [0, 1]}$ is narrowly continuous. In particular, if*

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\},$$

then $(\mu_t)_{t \in [0, 1]}$ is a displacement ℓ_p -geodesic by Proposition 2.66 below.

Proof. Our definition of geodesy (or the compactness of $\text{spt } \mu_0 \times \text{spt } \mu_1$) ensures $\ell_p(\mu_0, \mu_1) < \infty$. By Lemma 2.63 and compactness of $\text{spt } \mu_0$ and $\text{spt } \mu_1$, the measure μ_t has support in $Z(\text{spt } \mu_0 \times \text{spt } \mu_1)$ for every $0 \leq t \leq 1$. The latter is compact by

Lemma 2.49, and notably independent of t . Lemma 2.56 establishes existence of ℓ_p -optimal couplings of causally related, compactly supported marginals. The fact $\ell_p(\mu_0, \mu_1) > 0$ implied by our last hypothesis and our above observations will thus yield existence of ℓ_p -optimal couplings of all measures considered below.

We first claim narrow right-continuity of $(\mu_t)_{t \in (0,1)}$. Let $(t_n)_{n \in \mathbf{N}}$ be any sequence decreasing to a given $0 < t < 1$. By compactness of $Z(\text{spt } \mu_0 \times \text{spt } \mu_1)$, $(\mu_{t_n})_{n \in \mathbf{N}}$ has a narrow limit $\nu \in \mathcal{P}_c(\mathbf{M})$ along a nonrelabeled subsequence. We claim $\nu = \mu_t$. Suppose to the contrary $\nu \neq \mu_t$. By Theorem B.5, we have $\ell_p(\mu_t, \nu) \geq 0$, and by Lemma 2.56 an ℓ_p -optimal coupling π_2 of μ_t and ν exists. Now we construct a coupling π of μ_0 and μ_1 as follows. Let π_1 as well as π_3 be ℓ_p -optimal couplings of μ_0 and μ_t as well as ν and μ_1 , respectively. We glue these couplings π_1 , π_2 , and π_3 together, cf. e.g. [44, Prop. 2.5] or [91, Prop. 2.9] for this standard construction. We thus obtain a measure $\omega \in \mathcal{P}(\mathbf{M}^4)$ with the properties

$$\begin{aligned} (\text{pr}_1, \text{pr}_2)_\# \omega &= \pi_1, \\ (\text{pr}_2, \text{pr}_3)_\# \omega &= \pi_2, \\ (\text{pr}_3, \text{pr}_4)_\# \omega &= \pi_3. \end{aligned}$$

Finally, we define $\pi := (\text{pr}_1, \text{pr}_4)_\# \omega$. By construction, π is causal and obeys

$$\begin{aligned} \ell_p(\mu_0, \mu_1) &\geq \|l\|_{L^p(\mathbf{M}^2, \pi)} \\ &\geq \|l \circ (\text{pr}_1, \text{pr}_2) + l \circ (\text{pr}_2, \text{pr}_3) + l \circ (\text{pr}_3, \text{pr}_4)\|_{L^p(\mathbf{M}^4, \omega)} \\ &\geq \|l\|_{L^p(\mathbf{M}^2, \pi_1)} + \|l\|_{L^p(\mathbf{M}^2, \pi_2)} + \|l\|_{L^p(\mathbf{M}^2, \pi_3)} \\ &= \ell_p(\mu_0, \mu_t) + \ell_p(\mu_t, \nu) + \ell_p(\nu, \mu_1) \\ &\geq t \ell_p(\mu_0, \mu_1) + (1-t) \ell_p(\mu_0, \mu_1) \\ &= \ell_p(\mu_0, \mu_1); \end{aligned} \tag{2.20}$$

here we used (2.2) in the second, the reverse Hölder inequality in the third, and (2.18) (plus the inequality $\ell_p(\mu_t, \nu) \geq 0$) in the second last line. Therefore, all the inequalities in (2.20) must be equalities. In particular, π is an ℓ_p -optimal coupling of μ_0 and μ_1 and thus chronological by assumption. For ω -a.e. $(x_1, x_2, x_3, x_4) \in \mathbf{M}^4$ we thus have $x_1 \ll x_4$; moreover, since equality holds throughout (2.20), by strict convexity (implied by $p < 1$) again we obtain

$$\begin{aligned} l(x_1, x_2) &= t l(x_1, x_4), \\ l(x_2, x_3) &= 0, \\ l(x_3, x_4) &= (1-t) l(x_1, x_4), \end{aligned} \tag{2.21}$$

and therefore $x_1 \ll x_2$ and $x_3 \ll x_4$. On the other hand, π_2 necessarily places mass outside the diagonal of \mathbf{M}^2 since $\nu \neq \mu_t$. In particular $\omega[E_\varepsilon] > 0$ for some $\varepsilon > 0$ by monotone convergence, where E_ε is the set of points $x \in \mathbf{M}^4$ with $d(x_2, x_3) \geq \varepsilon$; cf. Theorem 2.18. For every $x \in E_\varepsilon$, by (2.21) following an l -geodesic from x_1 to x_2 , jumping from x_2 to x_3 , and going from there to x_4 by a further l -geodesic gives a discontinuous element of $\text{TGeo}(\mathbf{M})$ from x_1 to x_4 , a contradiction to Lemma 2.46.

The case $t = 0$ is dealt with similarly. Here the l -geodesics constructed above directly start with a jump. Thus $(\mu_t)_{t \in [0,1]}$ is right-continuous.

Analogously, we show narrow left-continuity of $(\mu_t)_{t \in (0,1]}$.

The last claim about displacement ℓ_p -geodesy follows from Proposition 2.66 once we have narrow continuity at our disposal. \square

Remark 2.65 (Propagation of moments). If μ_0 or μ_1 do not have compact support, it is less clear how the moment condition (2.15), ensuring existence of ℓ_p -optimal

couplings [44, Prop. 2.3], propagates throughout rough ℓ_p -geodesics; such a property might follow from the construction of a nonsmooth Hopf-Lax semigroup.

2.7.3. *Lifting the geometry on M to $\mathcal{P}(M)$.* We turn to the lifting procedure for ℓ_p -geodesics stated in Proposition 2.66. Some of the statements hold in higher generality, but we will not need these extensions. On the other hand, unlike [21] where (v) below was available for strongly timelike p -dualizable pairs (μ_0, μ_1) — based on uniform Lipschitz continuity of causal curves and reparametrization — our general compactness result from Corollary 2.47 only guarantees the necessary tightness properties away from the lightcone.

Under curvature and nonbranching assumptions, this lifting result holds for more general marginal measures compared to (2.22), cf. Corollary 3.25.

Recall the definition (2.11) of the restriction map restr_s^t for $0 \leq s < t \leq 1$.

Proposition 2.66 (Criteria for an ℓ_p -geodesic to be represented by a plan). *Let $\mu_0, \mu_1 \in \mathcal{P}(M)$. Then the following hold.*

- (i) *For every chronological ℓ_p -optimal coupling $\pi \in \Pi_{\ll}(\mu_0, \mu_1)$, there exists a timelike ℓ_p -optimal dynamical plan $\boldsymbol{\pi} \in \text{OptTGeo}_p(\mu_0, \mu_1)$ with*

$$\pi = (\mathbf{e}_0, \mathbf{e}_1)_{\#} \boldsymbol{\pi}.$$

- (ii) *If (μ_0, μ_1) is timelike p -dualizable, there exists at least one displacement ℓ_p -geodesic from μ_0 to μ_1 .*
 (iii) *Given any $0 \leq s < t \leq 1$, the inclusion $\boldsymbol{\pi} \in \text{OptTGeo}_p(\mu_0, \mu_1)$ implies the restriction $(\text{restr}_s^t)_{\#} \boldsymbol{\pi}$ to belong to $\text{OptTGeo}_p(\mu_s, \mu_t)$, where*

$$\mu_s := (\mathbf{e}_s)_{\#} \boldsymbol{\pi},$$

$$\mu_t := (\mathbf{e}_t)_{\#} \boldsymbol{\pi}.$$

- (iv) *Let $\boldsymbol{\pi} \in \text{OptTGeo}_p(\mu_0, \mu_1)$, and let $\boldsymbol{\sigma} \in \mathcal{P}(C([0, 1]; M))$ be a nonzero measure with $\boldsymbol{\sigma} \leq \boldsymbol{\pi}$. Define the probability measures*

$$\sigma_0 := \boldsymbol{\sigma}[C([0, 1]; M)]^{-1} (\mathbf{e}_0)_{\#} \boldsymbol{\sigma},$$

$$\sigma_1 := \boldsymbol{\sigma}[C([0, 1]; M)]^{-1} (\mathbf{e}_1)_{\#} \boldsymbol{\sigma}.$$

Then $\boldsymbol{\sigma}[C([0, 1]; M)]^{-1} \boldsymbol{\sigma} \in \text{OptTGeo}_p(\sigma_0, \sigma_1)$.

Now we additionally assume compactness of $\text{spt } \mu_0$ and $\text{spt } \mu_1$ as well as

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}, \tag{2.22}$$

and that $0 < p < 1$. Then the following hold.

- (v) *Every ℓ_p -geodesic from μ_0 to μ_1 is in fact a displacement ℓ_p -geodesic.*
 (vi) *Let $(\mu_0^n)_{n \in \mathbf{N}}$ and $(\mu_1^n)_{n \in \mathbf{N}}$ be sequences in $\mathcal{P}_c(M)$ converging narrowly to μ_0 and μ_1 , respectively. Then every sequence $(\boldsymbol{\pi}^n)_{n \in \mathbf{N}}$ of timelike ℓ_p -optimal dynamical plans $\boldsymbol{\pi}^n \in \text{OptTGeo}_p(\mu_0^n, \mu_1^n)$ has a narrow accumulation point.*
 (vii) *If, in addition, there are compact sets $C_0, C_1 \subset M$ such that $\text{spt } \mu_0^n \subset C_0$ and $\text{spt } \mu_1^n \subset C_1$ for every $n \in \mathbf{N}$ — in other words, $(\mu_0^n)_{n \in \mathbf{N}}$ and $(\mu_1^n)_{n \in \mathbf{N}}$ are uniformly compactly supported — and*

$$\ell_p(\mu_0, \mu_1) \leq \liminf_{n \rightarrow \infty} \ell_p(\mu_0^n, \mu_1^n),$$

then $\ell_p(\mu_0^n, \mu_1^n) \rightarrow \ell_p(\mu_0, \mu_1)$ as $n \rightarrow \infty$, and any accumulation point from (vi) belongs to $\text{OptTGeo}_p(\mu_0, \mu_1)$. This holds in particular if the endpoint marginal sequences are constant, in which case the previous statement implies narrow compactness of $\text{OptTGeo}_p(\mu_0, \mu_1)$.

Proof. We first establish (i). Using causal closedness (Proposition 2.10), we deduce the multivalued map sending a chronologically related pair $(x, y) \in M^2$ to the set of all $\gamma \in \text{TGeo}(M)$ connecting x to y to have closed graph. The statement then follows from standard measurable selection arguments, cf. e.g. [7, 129].

To see (ii), apply (i) to a timelike p -dualizing coupling of μ_0 and μ_1 .

Item (iii) easily follows from Remark 2.62.

We turn to (iv). The inequality $\sigma \leq \pi$ is stable under push-forwards, which shows the sub-probability measure $(\mathbf{e}_0, \mathbf{e}_1)_\# \sigma$ to be chronological. By the restriction property of ℓ_p -optimal couplings [44, Lem. 2.10], $(\mathbf{e}_0, \mathbf{e}_1)_\# \sigma$ is thus an ℓ_p -optimal coupling of its marginals σ_0 and σ_1 . Since $\sigma \leq \pi$, the measure σ is concentrated on $\text{TGeo}(M)$, and so is its normalized version.

We pass over to (vi). We claim tightness of $(\pi^n)_{n \in \mathbf{N}}$. Given any $\varepsilon > 0$, by narrow convergence of $(\mu_0^n)_{n \in \mathbf{N}}$ and $(\mu_1^n)_{n \in \mathbf{N}}$ there exist compact sets $C_0, C_1 \subset M$ with

$$\begin{aligned} \sup_{n \in \mathbf{N}} \mu_0^n [C_0^c] &\leq \varepsilon, \\ \sup_{n \in \mathbf{N}} \mu_1^n [C_1^c] &\leq \varepsilon. \end{aligned}$$

Recall the set $G := \text{TGeo}(M) \cap \mathbf{e}_0^{-1}(C_0) \cap \mathbf{e}_1^{-1}(C_1)$ is precompact thanks to Corollary 2.47 and (2.22). Observe that

$$\begin{aligned} \sup_{n \in \mathbf{N}} \pi^n [C([0, 1]; M) \setminus G] &\leq \sup_{n \in \mathbf{N}} \pi^n [\mathbf{e}_0^{-1}(C_0^c)] + \sup_{n \in \mathbf{N}} \pi^n [\mathbf{e}_1^{-1}(C_1^c)] \\ &= \sup_{n \in \mathbf{N}} \mu_0 [C_0^c] + \sup_{n \in \mathbf{N}} \mu_1 [C_1^c] \\ &\leq 2\varepsilon. \end{aligned}$$

This consideration shows the desired tightness by Corollary 2.47. By Prokhorov's theorem, $(\pi^n)_{n \in \mathbf{N}}$ has an accumulation point $\pi \in \mathcal{P}(C([0, 1]; M))$.

Next, we prove $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ under the assumptions of (vii). It suffices to show ℓ_p -optimality of $(\mathbf{e}_0, \mathbf{e}_1)_\# \pi$; indeed, π is concentrated on the uniform closure of $\text{TGeo}(M)$ by Alexandrov's theorem, and ℓ_p -optimality of $(\mathbf{e}_0, \mathbf{e}_1)_\# \pi$ would imply chronology of the latter by assumption, and thus $\pi[\text{TGeo}(M)] = 1$. To this aim, uniform boundedness of l_+ on $J(C_0, C_1)$ combined with [129, Lem. 4.3] yield

$$\begin{aligned} \ell_p(\mu_0, \mu_1) &\geq \|l\|_{L^p(M^2, (\mathbf{e}_0, \mathbf{e}_1)_\# \pi)} \\ &= \|l_+\|_{L^p(M^2, (\mathbf{e}_0, \mathbf{e}_1)_\# \pi)} \\ &\geq \limsup_{n \rightarrow \infty} \|l_+\|_{L^p(M^2, (\mathbf{e}_0, \mathbf{e}_1)_\# \pi^n)} \\ &= \limsup_{n \rightarrow \infty} \ell_p(\mu_0^n, \mu_1^n) \\ &\geq \liminf_{n \rightarrow \infty} \ell_p(\mu_0^n, \mu_1^n) \\ &\geq \ell_p(\mu_0, \mu_1). \end{aligned}$$

This confirms the limit asserted and proves the desired ℓ_p -optimality of $(\mathbf{e}_0, \mathbf{e}_1)_\# \pi$.

Lastly, item (v) requires most work, but is performed in analogy with the usual dyadic argument, cf. e.g. [7, Thm. 2.10] or [129, Thm. 7.21, Cor. 7.22]. We build a sequence $(\pi^n)_{n \in \mathbf{N}}$ in $\text{OptTGeo}_p(\mu_0, \mu_1)$ as follows. Given any $n \in \mathbf{N}$ and any $k = 1, \dots, 2^n$, we first observe that $\mu_{(k-1)2^{-n}}$ and $\mu_{k2^{-n}}$ admit a causal coupling since they lie on an ℓ_p -geodesic. These measures also have support in a compact set which is independent of k or n by Lemma 2.63. Thus, by [44, Prop. 2.3] these admit an ℓ_p -optimal coupling $\pi_k^n \in \Pi_{\leq}(\mu_{(k-1)2^{-n}}, \mu_{k2^{-n}})$.

Using Lemma 2.52 we find $\omega^n \in \mathcal{P}(M^{2^n+1})$ such that, for every $k = 1, \dots, 2^n$,

$$(\text{pr}_{k-1}, \text{pr}_k)_\# \omega^n = \pi_k^n. \quad (2.23)$$

The coupling $\pi^n := (\text{pr}_0, \text{pr}_{2^n})_{\#} \omega^n$ of μ_0 and μ_1 is causal and obeys

$$\begin{aligned}
\ell_p(\mu_0, \mu_1) &\geq \left\| l \right\|_{L^p(\mathbb{M}^2, \pi^n)} \\
&\geq \left\| \sum_{k=1}^{2^n} l \circ (\text{pr}_{k-1}, \text{pr}_k) \right\|_{L^p(\mathbb{M}^{2^n+1}, \omega^n)} \\
&\geq \sum_{k=1}^{2^n} \left\| l \right\|_{L^p(\mathbb{M}^2, \pi_k^n)} \\
&= \sum_{k=1}^{2^n} \ell_p(\mu_{(k-1)2^{-n}}, \mu_{k2^{-n}}) \\
&= \ell_p(\mu_0, \mu_1);
\end{aligned} \tag{2.24}$$

here we used (2.2) in the second line, the reverse Hölder inequality in the third line, and ℓ_p -geodesy in the last line. Therefore, equalities must hold throughout the above chain of inequalities. This implies π^n is an ℓ_p -optimal coupling of μ_0 and μ_1 , hence chronological. Moreover, by [91, Prop. 2.9] ω^n -a.e. $(x_0, x_1, \dots, x_{2^n}) \in \mathbb{M}^{2^n+1}$ satisfies the following identity for every $k = 1, \dots, 2^n$:

$$l(x_{k-1}, x_k) = 2^{-n} l(x_0, x_{2^n}). \tag{2.25}$$

In turn, this yields chronology of π_k^n for every such k ; by item (i), we consequently find a plan $\pi_k^n \in \text{OptTGeo}_p(\mu_{(k-1)2^{-n}}, \mu_{k2^{-n}})$ with

$$\pi_k^n = (\mathbf{e}_0, \mathbf{e}_1)_{\#} \pi_k^n.$$

Owing to (2.23) and (2.25) we thus construct $\pi^n \in \mathcal{P}(\mathcal{C}([0, 1]; \mathbb{M}))$ such that

- for every $k = 1, \dots, 2^n$, we have $(\text{res}_{(k-1)2^{-n}}^k)_{\#} \pi^n = \pi_k^n$,
- $(\mathbf{e}_0, \mathbf{e}_1)_{\#} \pi$ is a chronological ℓ_p -optimal coupling of μ_0 and μ_1 , and
- π^n is concentrated on $\text{TGeo}(\mathbb{M})$.

In other words, we have $\pi^n \in \text{OptTGeo}_p(\mu_0, \mu_1)$.

For every $n \in \mathbf{N}$, by construction the timelike ℓ_p -optimal dynamical plan π^n is concentrated on the set of all $\gamma \in \text{TGeo}(\mathbb{M})$ with $\gamma_0 \in \text{spt } \mu_0$ and $\gamma_1 \in \text{spt } \mu_1$. An argument as for (vi) above (note that singletons of probability measures are automatically tight) shows tightness of the sequence $(\pi^n)_{n \in \mathbf{N}}$, and thus a nonrelabeled subsequence of it converges narrowly to some $\pi \in \mathcal{P}(\mathcal{C}([0, 1]; \mathbb{M}))$. In particular, the endpoint couplings $(\mathbf{e}_0, \mathbf{e}_1)_{\#} \pi^n$ converge narrowly to $(\mathbf{e}_0, \mathbf{e}_1)_{\#} \pi$, which is a coupling of μ_0 and μ_1 . Alexandrov's theorem yields π to be concentrated on causal curves. Thus, stability of ℓ_p -optimal couplings in the form [44, Thm. 2.14] in fact saturates ℓ_p -optimality and hence chronology of $(\mathbf{e}_0, \mathbf{e}_1)_{\#} \pi$.

We claim that π represents $(\mu_t)_{t \in [0, 1]}$. By construction, we have $(\mathbf{e}_0)_{\#} \pi = \mu_0$ and $(\mathbf{e}_1)_{\#} \pi = \mu_1$. Given any dyadic number $0 < t < 1$, passing to the limit in the identity $\mu_t = (\mathbf{e}_t)_{\#} \pi^n$, valid for every sufficiently large $n \in \mathbf{N}$, yields $\mu_t = (\mathbf{e}_t)_{\#} \pi$ for every such t . Since $(\mu_t)_{t \in [0, 1]}$ is narrowly continuous, and since π is concentrated on continuous curves, both sides of the previous identity depend continuously on t with respect to narrow convergence. This yields $\mu_t = (\mathbf{e}_t)_{\#} \pi$ for every $0 < t < 1$.

Since $\pi[\text{TGeo}(\mathbb{M})] = 1$ by (vii), this terminates the proof. \square

By using the first part of Proposition 2.66, the proof of the following result is analogous to [22, Lem. 3.1], see also [91, Prop. 5.5]. The difference to point out here compared to [22] is that with our notion of ℓ_p -geodesics, propagation of bare timelike p -dualizability is unclear, but not expected to hold in general.

Lemma 2.67 (Criteria for propagation of (strong) timelike p -dualizability). *Let $(\mu_t)_{t \in [0, 1]}$ be an ℓ_p -geodesic.*

- (i) If the pair (μ_0, μ_1) is strongly timelike p -dualizable, so is (μ_s, μ_t) for every $0 \leq s < t \leq 1$.
- (ii) If $(\mu_t)_{t \in [0,1]}$ constitutes a displacement ℓ_p -geodesic, then timelike p -dualizability of (μ_0, μ_1) implies that of (μ_s, μ_t) for every $0 \leq s < t \leq 1$.

2.7.4. Compactness.

Theorem 2.68 (Accumulation of rough ℓ_p -geodesics at rational times). *Assume Hypothesis 2.41. For every $n \in \mathbf{N}$, let $(\mu_t^n)_{t \in \mathbf{N}}$ be a given rough ℓ_p -geodesic, such that the sequences $(\mu_0^n)_{n \in \mathbf{N}}$ and $(\mu_1^n)_{n \in \mathbf{N}}$ are uniformly compactly supported, say in compact sets $C_0, C_1 \subset M$, respectively. Assume accumulation points $\mu_0, \mu_1 \in \mathcal{P}_c(M)$ of the two sequences satisfy*

$$0 < \ell_p(\mu_0, \mu_1) \leq \liminf_{n \rightarrow \infty} \ell_p(\mu_0^n, \mu_1^n). \quad (2.26)$$

Then there is a rough ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ from μ_0 to μ_1 such that up to a nonre-labeled subsequence, $(\mu_t^n)_{n \in \mathbf{N}}$ converges narrowly to μ_t for every rational $0 \leq t \leq 1$.

As in Proposition 2.66, the lower semicontinuity (2.26) in Theorem 2.68 implies the limit exists and equality holds [44, Lem. 2.11, Thm. 2.14].

Proof of Theorem 2.68. The construction is similar to the first part of the proof of Corollary 2.47, thus we only outline it. By Prokhorov's theorem and Lemma 2.49, we may and will assume the existence of a family $(\mu_t)_{t \in [0,1] \cap \mathbf{Q}}$ in $\mathcal{P}(M)$ with support in the common compact subset $Z(C_0 \times C_1) \subset M$ such that up to a nonre-labeled subsequence, the sequence $(\mu_t^n)_{n \in \mathbf{N}}$ converges narrowly to μ_t for every rational $0 \leq t \leq 1$. Given any irrational $0 \leq t \leq 1$, let $(t_i)_{i \in \mathbf{N}}$ be defined by $t_i := \lceil t 2^i \rceil 2^{-i}$. As above, we can thus define $\mu_t \in \mathcal{P}_c(M)$ as a subsequential narrow limit of $(\mu_{t_i})_{i \in \mathbf{N}}$ — in fact, the definition of μ_t does not depend on the choice of monotone approximating sequence by antisymmetry, cf. Proposition B.4; compare with the proof of Corollary 2.47.

We claim that the collection $(\mu_t)_{t \in [0,1]}$ is a rough ℓ_p -geodesic from μ_0 to μ_1 . To this aim, we first show every fixed rational $0 \leq s < t \leq 1$ to satisfy

$$\ell_p(\mu_s, \mu_t) \geq (t - s) \ell_p(\mu_0, \mu_1). \quad (2.27)$$

The construction of μ_s and μ_t , tightness of couplings, Alexandrov's theorem, and closedness of $\{l \geq 0\}$ ensure $\Pi_{\leq}(\mu_s, \mu_t) \neq \emptyset$. Since l_+ is bounded on $Z(C_0 \times C_1)$, again by Alexandrov's theorem and tightness of couplings some $\pi_{s,t} \in \Pi_{\leq}(\mu_s, \mu_t)$ is the narrow limit of a sequence $(\pi_{s,t}^n)_{n \in \mathbf{N}}$ of ℓ_p -optimal couplings $\pi_{s,t}^n \in \Pi_{\leq}(\mu_s^n, \mu_t^n)$, up to a nonre-labeled subsequence. This already yields (2.27) because

$$\begin{aligned} \ell_p(\mu_s, \mu_t) &\geq \|l\|_{L^p(M^2, \pi_{s,t})} \\ &= \lim_{n \rightarrow \infty} \|l\|_{L^p(M^2, \pi_{s,t}^n)} \\ &\geq (t - s) \liminf_{n \rightarrow \infty} \ell_p(\mu_0^n, \mu_1^n) \\ &\geq (t - s) \ell_p(\mu_0, \mu_1). \end{aligned}$$

Finally, we show (2.27) to hold for every $0 \leq s < t \leq 1$ — as in Remark 2.62, this will imply ℓ_p -geodesy of $(\mu_t)_{t \in [0,1]}$. But by “causally monotone continuity of ℓ_p ” [44, Thm. 2.14], it is clear that (2.27) extends to every $0 \leq s < t \leq 1$. \square

Remark 2.69 (Narrow continuity). If every ℓ_p -optimal coupling of the endpoint marginals μ_0 and μ_1 in Theorem 2.68 is chronological, by Corollary 2.64 the limit rough ℓ_p -geodesic is an ℓ_p -geodesic, and it is uniquely determined by its values on rational $0 \leq t \leq 1$.

More generally, if a limit curve $(\mu_t)_{t \in [0,1]}$ of rough ℓ_p -geodesics $(\mu_t^n)_{t \in [0,1]}$, where $n \in \mathbf{N}$, according to Theorem 2.68 above is narrowly continuous, then $(\mu_t^n)_{n \in \mathbf{N}}$




FIGURE 1. Timelike (forward and backward) branching.

converges narrowly to μ_t for every $0 \leq t \leq 1$. This follows from antisymmetry of the causal relation \preceq on $\mathcal{P}(\mathcal{M})$ discussed in Section B.4; compare with the proof of Corollary 2.47. Indeed, every subsequence of $(\mu_t^n)_{n \in \mathbf{N}}$ has a further nonrelabeled subsequence converging to some $\nu \in \mathcal{P}_c(\mathcal{M})$ by Prokhorov's theorem. For $s < t < r$ with $s, r \in \mathbf{Q}$, passing to the limit as $n \rightarrow \infty$ in the chain $\mu_s^n \preceq \mu_t^n \preceq \mu_r^n$ yields $\mu_s \preceq \nu \preceq \mu_r$. Narrow continuity of $(\mu_t)_{t \in [0,1]}$ gives $\mu_t \preceq \nu \preceq \mu_t$, and thus $\mu_t = \nu$ by antisymmetry of \preceq (Proposition B.4).

2.8. Taking the reference measure \mathbf{m} into account.

2.8.1. *Timelike p -essential nonbranching.* The following condition, which relaxes timelike nonbranching [44, Def. 1.10] to an ‘‘a.e. requirement’’, has been introduced in [21, Def. 2.21, Rem. 2.22] following [111]. It involves the reference measure \mathbf{m} from the metric measure spacetime \mathcal{M} discussed at (2.9).

A subset $G \subset \text{TGeo}(\mathcal{M})$ is *timelike nonbranching* if for every $0 \leq s < t \leq 1$, the map $\text{restr}_s^t|_G$ is injective, where restr_s^t is from (2.11).

Definition 2.70 (Timelike essential nonbranching). *We term the space \mathcal{M} timelike p -essentially nonbranching if for every $\pi \in \text{OptTGeo}_p(\mathcal{P}_c^{\text{ac}}(\mathcal{M}, \mathbf{m})^2)$ there exists a timelike nonbranching Borel set $G \subset \text{TGeo}(\mathcal{M})$ with $\pi[G] = 1$.*

Note that timelike p -essential nonbranching is a property of \mathbf{m} rather than \mathcal{M} (and that $\mathcal{M} = \text{spt } \mathbf{m}$ is not assumed).

Despite our more general approaches to geodesics on \mathcal{M} and $\mathcal{P}_c(\mathcal{M})$, respectively, the proof of [22, Lem. 2.23] carries over with no essential change.

Lemma 2.71 (Mutually singular interpolants). *Assume \mathcal{M} is timelike p -essentially nonbranching. Decompose $\pi \in \text{OptTGeo}_p(\mathcal{P}_c^{\text{ac}}(\mathcal{M}, \mathbf{m})^2)$ as $\pi = \lambda_1 \pi^1 + \dots + \lambda_n \pi^n$ for certain $0 < \lambda_1, \dots, \lambda_n < 1$ and $n \in \mathbf{N}$. Lastly, assume $(\mathbf{e}_t)_\# \pi^1, \dots, (\mathbf{e}_t)_\# \pi^n$ to be \mathbf{m} -absolutely continuous for some $0 < t < 1$, and mutual singularity of π^1, \dots, π^n . Then the measures $(\mathbf{e}_t)_\# \pi^1, \dots, (\mathbf{e}_t)_\# \pi^n$ are mutually singular as well.*

2.8.2. *Boltzmann entropy.* We define the Boltzmann (or Boltzmann–Shannon) entropy $\text{Ent}_{\mathbf{m}}: \mathcal{P}(\mathcal{M}) \rightarrow [-\infty, \infty]$ by decomposing $\mu = \rho \mathbf{m} + \mu_{\perp}$ into its absolutely continuous and singular parts with respect to \mathbf{m} and setting

$$\text{Ent}_{\mathbf{m}}(\mu) := \begin{cases} \int_{\mathcal{M}} \rho \log \rho \, d\mathbf{m} & \text{if } (\rho \log \rho)_{-} \in L^1(\mathcal{M}, \mathbf{m}) \text{ and } \mu_{\perp}(\mathcal{M}) = 0, \\ +\infty & \text{if } (\rho \log \rho)_{-} \in L^1(\mathcal{M}, \mathbf{m}) \text{ and } \mu_{\perp}(\mathcal{M}) > 0, \\ -\infty & \text{if } (\rho \log \rho)_{-} \notin L^1(\mathcal{M}, \mathbf{m}). \end{cases}$$

We denote by $\mathcal{D}(\text{Ent}_{\mathbf{m}})$ the finiteness domain of $\text{Ent}_{\mathbf{m}}$. We will mostly consider $\text{Ent}_{\mathbf{m}}$ on compactly supported measures. In particular, on $\mathcal{P}_c(\mathcal{M})$ the Boltzmann

entropy does not attain the value $-\infty$, since

$$\text{Ent}_{\mathfrak{m}}(\mu) \geq -\log \mathfrak{m}[\text{spt } \mu] > -\infty$$

for every $\mu \in \mathcal{P}_c(\mathbf{M})$ by Jensen's inequality and the Radon property of \mathfrak{m} .

We occasionally use the following narrow lower semicontinuity property of $\text{Ent}_{\mathfrak{m}}$, cf. e.g. [8, Lem. 9.4.3] or [120, Lem. 4.1]. If $(\mu_n)_{n \in \mathbf{N}}$ is a given uniformly compactly supported sequence in $\mathcal{P}(\mathbf{M})$ which converges narrowly to $\mu \in \mathcal{P}(\mathbf{M})$, then

$$\text{Ent}_{\mathfrak{m}}(\mu) \leq \liminf_{n \rightarrow \infty} \text{Ent}_{\mathfrak{m}}(\mu_n).$$

Given any $N > 0$, we also define the functional $\mathcal{U}_N: \mathcal{P}(\mathbf{M}) \rightarrow [0, \infty]$ by

$$\mathcal{U}_N := e^{-\text{Ent}_{\mathfrak{m}}(\mu)/N}. \quad (2.28)$$

We employ the conventions $e^{-\infty} := 0$ and $e^{\infty} := \infty$. Note that \mathcal{U}_N is narrowly upper semicontinuous in the same sense as $\text{Ent}_{\mathfrak{m}}$, and for every $\mu \in \mathcal{P}_c(\mathbf{M})$,

$$\mathcal{U}_N(\mu) \leq \mathfrak{m}[\text{spt } \mu]^{1/N}. \quad (2.29)$$

3. TIMELIKE CURVATURE-DIMENSION CONDITION

From now on, unless explicitly stated otherwise we fix an exponent $0 < p < 1$, a lower semicontinuous function $k: \mathbf{M} \rightarrow \mathbf{R}$, and a real number $N > 1$ ⁶. The lower boundedness of k is *not* assumed.

3.1. Distortion coefficients along l -geodesics. Now we describe how to generalize the distortion coefficients employed to extend the timelike-curvature dimension condition [21, 44] to the variable case. We refer to Ketterer [79, 80] for more details.

3.1.1. Continuous potentials on real intervals. Given any $L > 0$, let $\kappa: [0, L] \rightarrow \mathbf{R}$ be a continuous function. The *generalized sine function* $\sin_{\kappa}: [0, L] \rightarrow \mathbf{R}$ is the unique solution to the ODE $f'' + \kappa f = 0$ satisfying $f(0) = 0$ and $f'(0) = 1$. This solution depends uniformly continuously on κ .

Definition 3.1 (Distortion coefficients on the half-line $[0, \infty)$). *For $0 \leq t \leq 1$, the generalized distortion coefficient $\sigma_{\kappa}^{(t)}: [0, L] \rightarrow [0, \infty]$ is defined by*

$$\sigma_{\kappa}^{(t)}(\theta) := \begin{cases} \frac{\sin_{\kappa}(t\theta)}{\sin_{\kappa}(\theta)} & \text{if } \sin_{\kappa}(r) > 0 \text{ for every } 0 < r \leq \theta, \\ \infty & \text{otherwise.} \end{cases} \quad (3.1)$$

Example 3.2 (Constant potentials). If κ is constant, the distortion coefficients from Definition 3.1 take the following customary form:

$$\sigma_{\kappa}^{(t)}(\theta) = \begin{cases} \frac{\sin(\sqrt{\kappa} t \theta)}{\sin(\sqrt{\kappa} \theta)} & \text{if } 0 < \kappa \theta^2 < \pi^2, \\ \infty & \text{if } \kappa \theta^2 \geq \pi^2, \\ t & \text{if } \kappa \theta^2 = 0, \\ \frac{\sinh(\sqrt{-\kappa} t \theta)}{\sinh(\sqrt{-\kappa} \theta)} & \text{if } \kappa < 0. \end{cases}$$

Remark 3.3 (Scaling). Given any $0 \leq \theta \leq L$, we have either $\sigma_{\kappa}^{(t)}(\theta) = \infty$ for every $0 \leq t \leq 1$, or $\sigma_{\kappa}^{(t)}(\theta) < \infty$ for every $0 \leq t \leq 1$, cf. Lemma 3.6 below.

In the latter case, the function $u: [0, 1] \rightarrow [0, \infty)$ given by $u(t) := \sigma_{\kappa}^{(t)}(\theta)$ obeys $u(0) = 0$, $u(1) = 1$, and the following ODE for every $0 < t < 1$:

$$u''(t) + \kappa(t\theta)\theta^2 u(t) = 0. \quad (3.2)$$

⁶The entropic condition we will introduce actually only requires $N > 0$. For other timelike curvature-dimension conditions such as in [21] the situation $N = 1$ is technically slightly different from $N > 1$, but can be understood as the limiting case of the latter; cf. e.g. Proposition 3.18.

Uniqueness implies the useful scaling property

$$\sigma_{\kappa}^{(t)}(\theta) = \sigma_{\kappa\theta^2}^{(t)}(1), \quad (3.3)$$

where $\kappa\theta^2$ is understood as the function $\tilde{\kappa}: [0, 1] \rightarrow \mathbf{R}$ with

$$\tilde{\kappa}(t) := \kappa(t\theta)\theta^2.$$

3.1.2. Continuous potentials along l -geodesics. Now we adapt these notions to a continuous function $k: \mathbf{M} \rightarrow \mathbf{R}$. The natural approach is to consider k along an appropriate timelike curve.

Let $\gamma \in \text{TGeo}(\mathbf{M})$ be fixed, and define

$$|\dot{\gamma}| := l(\gamma_0, \gamma_1). \quad (3.4)$$

We use this notation only for formal reasons. However, it should be compared to the *metric speed* of absolutely continuous curves in metric spaces, cf. e.g. [8, Thm. 1.1.2]. In fact, for causal curves γ which need not be geodesic, we define a notion $|\dot{\gamma}|: [0, 1] \rightarrow \mathbf{R}$ of *causal speed* in [14] which reduces to (3.4) if $\gamma \in \text{TGeo}(\mathbf{M})$; in this more general setting, $|\dot{\gamma}|$ need not be constant; only its average over $[0, 1]$ is dominated by $l(\gamma_0, \gamma_1)$, and we can then consider the l -arclength reparametrization $\bar{\gamma}: [0, |\dot{\gamma}|] \rightarrow \mathbf{R}$ of a rectifiable curve γ given by Proposition 2.37. For $\gamma \in \text{TGeo}(\mathbf{M})$ — the only relevant case here — it has the simple form

$$\bar{\gamma}_r := \gamma_{r/|\dot{\gamma}|}.$$

In view of Remark 3.3, define the continuous forward and backward functions $k_{\gamma}^{\pm}: [0, |\dot{\gamma}|] \rightarrow \mathbf{R}$ by

$$\begin{aligned} k_{\gamma}^{-}(t|\dot{\gamma}|) &:= k(\gamma_{1-t}), \\ k_{\gamma}^{+}(t|\dot{\gamma}|) &:= k(\gamma_t), \end{aligned} \quad (3.5)$$

which depend continuously, with respect to uniform convergence, on $\gamma \in \text{TGeo}(\mathbf{M})$ for fixed $0 \leq t \leq 1$. Equivalently, we can express (3.5) as

$$\begin{aligned} k_{\gamma}^{-}(r) &:= k(\bar{\gamma}_{|\dot{\gamma}|-r}), \\ k_{\gamma}^{+}(r) &:= k(\bar{\gamma}_r). \end{aligned}$$

3.1.3. Lower semicontinuous potentials along l -geodesics. Now we assume $k: M \rightarrow \mathbf{R}$ to be only lower semicontinuous (and we still fix $\gamma \in \text{TGeo}(\mathbf{M})$). In this case, we use the previous constructions and argue by approximation. The scheme can be found in [8, Sec. 5.1]. It has proven useful in the general study of metric measure spaces with variable Ricci curvature bounds [25, 79, 80, 122, 123].

Let $E \subset \mathbf{M}$ be any set containing $\gamma_{[0,1]}$ such that $\inf k(E) > -\infty$. Typically, we only consider $E := \gamma_{[0,1]}$ when γ is fixed, but when γ varies it will be convenient to choose E larger; the precise choice of E does not matter, cf. Remark 3.5. We approximate k pointwise by a nondecreasing sequence of functions as follows. For $n \in \mathbf{N}$, define $k_n: E \rightarrow \mathbf{R}$ by

$$k_n(x) := \inf\{\min\{k(y), n\} + n \mathbf{d}(x, y) : y \in E\}. \quad (3.6)$$

It is not difficult to check that

- k_n is a bounded Lipschitz function that satisfies $\inf k(E) \leq k_n \leq k$ on E for every $n \in \mathbf{N}$, and
- $(k_n)_{n \in \mathbf{N}}$ is nondecreasing and converges pointwise to k on E .

For the next definition, we consider the function $k_{n,\gamma}^{\pm}: [0, |\dot{\gamma}|] \rightarrow \mathbf{R}$ defined by (3.5) with k replaced by k_n from (3.6).

Definition 3.4 (Distortion coefficients along l -geodesics). *Given any $0 \leq \theta \leq |\dot{\gamma}|$, we define*

- a. the generalized sine functions $\sin_{k_\gamma^\pm}$ associated with k_γ^\pm by

$$\sin_{k_\gamma^\pm}(\theta) := \lim_{n \rightarrow \infty} \sin_{k_{n,\gamma}^\pm}(\theta),$$

- b. the distortion coefficients $\sigma_{k_\gamma^\pm}^{(t)}$, where $0 \leq t \leq 1$, associated with k_γ^\pm by

$$\sigma_{k_\gamma^\pm}^{(t)}(\theta) := \lim_{n \rightarrow \infty} \sigma_{k_{n,\gamma}^\pm}^{(t)}(\theta).$$

Remark 3.5 (Well-definedness). Definition 3.4 is reasonable since the involved sequences are monotone by the nondecreasingness of $(k_n)_{n \in \mathbf{N}}$ and (Jacques Charles François) Sturm's comparison theorem [80, Thm. 3.1], and since $(k_{n,\gamma}^\pm)_{n \in \mathbf{N}}$ converges pointwise to k_γ^\pm . By a similar argument, these definitions are independent of the choice of the approximating sequence [79, Rem. 3.6], and in particular of E .

The ODE (3.2) still holds for $\kappa := k_\gamma^\pm$ in the distributional sense.

If k is constant(ly equal to $\kappa \in \mathbf{R}$) along γ in Definition 3.4, then

$$\sigma_{k_\gamma^\pm}^{(t)}(\theta) = \sigma_\kappa^{(t)}(\theta),$$

where $\sigma_\kappa^{(t)}(\theta)$ is as in (3.2) with constant potential κ .

We refer to [79, Lem. 3.9] and [80, Props. 3.4, 3.15, Lems. 3.10, 3.14] for proofs of the subsequent facts. Let $\text{LSC}(\mathbf{M})$ denote the convex space of real-valued lower semicontinuous functions on \mathbf{M} .

Lemma 3.6 (Some properties of distortion coefficients). *Let $E \subset \mathbf{M}$ be as above. The generalized distortion coefficients from Definition 3.4 satisfy the following for every $k_\gamma \in \{k_\gamma^-, k_\gamma^+\}$ and every $0 \leq \theta \leq |\dot{\gamma}|$.*

- (i) **Nondegeneracy.** *Either $\sigma_{k_\gamma}^{(t)}(\theta) = \infty$ for every $0 < t < 1$ or $\sigma_{k_\gamma}^{(t)}(\theta) < \infty$ for every $0 < t < 1$. In the latter case,*

$$\sigma_{k_\gamma}^{(t)}(\theta) = \frac{\sin_{k_\gamma}(t\theta)}{\sin_{k_\gamma}(\theta)}.$$

- (ii) **Monotonicity.** *If a lower semicontinuous function $k' : E \rightarrow \mathbf{R}$ satisfies $k' \leq k$ on E , then $\sigma_{k'}^{(t)}(\theta) \leq \sigma_k^{(t)}(\theta)$ for every $0 \leq t \leq 1$.*
- (iii) **Lower semicontinuity.** *Let $(k_i)_{i \in \mathbf{N}}$ be a sequence of lower semicontinuous functions $k_i : E \rightarrow \mathbf{R}$ such that for every $x \in E$,*

$$k(x) \leq \liminf_{i \rightarrow \infty} k_i(x).$$

Define $k_{i,\gamma}$ by (3.5) with k replaced by k_i . Then for every $0 \leq t \leq 1$,

$$\sigma_{k_\gamma}^{(t)}(\theta) \leq \liminf_{i \rightarrow \infty} \sigma_{k_{i,\gamma}}^{(t)}(\theta).$$

In particular, $\sigma_{k_\gamma}^{(t)}(|\dot{\gamma}|)$ is lower semicontinuous in $\gamma \in \text{TGeo}(\mathbf{M})$ with respect to uniform convergence.

- (iv) **Convexity.** *For every $0 \leq t \leq 1$, $\sigma_{k_\gamma}^{(t)}(1)$ is logarithmically convex, hence convex in $k \in \text{LSC}(\mathbf{M})$. Moreover, the function $\mathbf{G}_t : \mathbf{R}^2 \times \text{LSC}(\mathbf{M}) \rightarrow \mathbf{R} \cup \{\infty\}$ defined by*

$$\mathbf{G}_t(r, s, k) := \log[\sigma_{k_\gamma}^{(1-t)}(1) e^r + \sigma_{k_\gamma}^{(t)}(1) e^s]$$

is jointly convex.

3.1.4. *Timelike (k, N) -convexity.* We continue to assume lower semicontinuity of $k: \mathbf{M} \rightarrow \mathbf{R}$. Based on the distortion coefficients from Definition 3.4, we introduce *timelike (k, N) -convexity* for functionals on metric spacetimes (without a reference measure); see [91, Def. 6.5] for a related definition for functionals defined on $\mathcal{P}(\mathbf{M})$ in the smooth case. We follow the approach initiated for metric measure spaces with constant k in [54], and further developed in [79] with variable k .

Recall our conventions $e^{-\infty} := 0$ and $e^{\infty} := \infty$. Let $g: [0, 1]^2 \rightarrow \mathbf{R}$ denote the usual Green's function

$$g(s, t) := \min\{s(1-t), t(1-s)\} \quad (3.7)$$

on $[0, 1]$ with Dirichlet boundary conditions. As usual, the finiteness domain of a functional $S: \mathbf{M} \rightarrow \mathbf{R} \cup \{\infty\}$ is denoted by $\mathcal{D}(S)$.

Definition 3.7 (Timelike k -convexity). *We say that S is*

- a. *weakly timelike k -convex if for every $x_0, x_1 \in \mathcal{D}(S)$ there is $\gamma \in \text{TGeo}(\mathbf{M})$ from x_0 to x_1 passing through $\mathcal{D}(S)$ such that for every $0 \leq t \leq 1$,*

$$S(\gamma_t) \leq (1-t)S(x_0) + tS(x_1) - \int_0^1 g(s, t) k(\gamma_s) |\dot{\gamma}|^2 ds, \quad (3.8)$$

- b. *strongly timelike k -convex if it is weakly timelike k -convex, and for every $x_0, x_1 \in \mathcal{D}(S)$, every l -geodesic $\gamma \in \text{TGeo}(\mathbf{M})$ from x_0 to x_1 passing through $\mathcal{D}(S)$ satisfies (3.8) for every $0 \leq t \leq 1$.*

The following modifies the previous definition by taking a further ‘‘dimensional parameter’’ N into account. Indeed, it is a stronger property than mere timelike k -convexity, cf. Corollary 3.19. Given a functional S as above, analogously to (2.28) we define $U_N: \mathbf{M} \rightarrow [0, \infty)$ by

$$U_N(x) := e^{-S(x)/N}.$$

Definition 3.8 (Timelike (k, N) -convexity). *We say that S is*

- a. *weakly timelike (k, N) -convex if for every $x_0, x_1 \in \mathcal{D}(S)$, there exists an l -geodesic $\gamma \in \text{TGeo}(\mathbf{M})$ from x_0 to x_1 passing through $\mathcal{D}(S)$ such that for every $0 \leq t \leq 1$,*

$$U_N(\gamma_t) \geq \sigma_{k\bar{\gamma}}^{(1-t)}(|\dot{\gamma}|) U_N(x_0) + \sigma_{k\bar{\gamma}}^{(t)}(|\dot{\gamma}|) U_N(x_1), \quad (3.9)$$

- b. *strongly timelike (k, N) -convex if it is weakly timelike (k, N) -convex, and for every $x_0, x_1 \in \mathcal{D}(S)$, every $\gamma \in \text{TGeo}(\mathbf{M})$ from x_0 to x_1 passing through $\mathcal{D}(S)$ satisfies (3.9) for every $0 \leq t \leq 1$.*

The following properties are elementary. The proofs from the metric case carry over with no noteworthy changes with the help of Lemma 3.6, cf. [54, Lems. 2.9, 2.10, 2.12] and [79, Lem. 3.15, Cor. 3.16]. Moreover, to get the idea for (ii) and (iii), the reader is invited to consult the proofs of Proposition 3.20 and Corollary 3.19 below, respectively.

Lemma 3.9 (Dependence on parameters). *Suppose $k': \mathbf{M} \rightarrow \mathbf{R}$ is lower semicontinuous, $N' > 0$, and let $S, S': \mathbf{M} \rightarrow \mathbf{R} \cup \{\infty\}$ be functionals. Then the following properties hold.*

- (i) **Scaling.** *If S is weakly timelike (k, N) -convex, then λS is weakly timelike $(\lambda k, \lambda N)$ -convex for every $\lambda > 0$.*
- (ii) **Addition.** *If S is weakly timelike (k, N) -convex and S' is strongly timelike (k', N') -convex, then $S + S'$ is weakly timelike $(k+k', N+N')$ -convex relative to the base space $\mathcal{D}(S) \cap \mathcal{D}(S')$.*

- (iii) **Monotonicity.** *If S is weakly timelike (k, N) -convex, $k' \leq k$ on M , and $N' \geq N$, then S is weakly timelike (k', N') -convex. In particular, S is weakly timelike k -convex.*

These claims also hold by replacing every occurrence of “weakly” by “strongly”.

3.2. Distortion coefficients along plans. In order to formulate Definition 3.13 below, we need to extend Definition 3.4 from functions along l -geodesics to functions along the dynamical plans $\text{OptTGeo}_p(\mathcal{P}_c(M)^2)$ of (2.17) (for more general plans concentrated on causal curves; adaptations of the definitions given below are straightforward). To this aim, we require some further notation, reflecting the fact that unless k is nonnegative, distorted convexity of the entropy along ℓ_q -geodesics needs to be measured in an L^2 -way rather than by ℓ_q -arclength, analogous to [44].

3.2.1. L^2 -cost. The L^2 -cost of a dynamical plan π is defined by

$$\mathbf{c}_\pi := \|l \circ (\mathbf{e}_0, \mathbf{e}_1)\|_{L^2(\text{TGeo}(M), \pi)}.$$

Denoting the endpoint marginals of π by $\mu_0 := (\mathbf{e}_0)_\# \pi$ and $\mu_1 := (\mathbf{e}_1)_\# \pi$, Jensen’s inequality and ℓ_p -optimality of $(\mathbf{e}_0, \mathbf{e}_1)_\# \pi \in \Pi(\mu_0, \mu_1)$ imply

$$\mathbf{c}_\pi \geq \ell_p(\mu_0, \mu_1). \quad (3.10)$$

3.2.2. *Lower semicontinuous potentials along plans.* Let $(k_n)_{n \in \mathbf{N}}$ be the monotone sequence from (3.6) approximating k on $E := J(\text{spt } \mu_0, \text{spt } \mu_1)$. For $n \in \mathbf{N}$, define the continuous forward and backward “superpositions” $k_{n, \pi}^\pm : [0, \mathbf{c}_\pi] \rightarrow \mathbf{R}$ by

$$\begin{aligned} k_{n, \pi}^-(t \mathbf{c}_\pi) \mathbf{c}_\pi^2 &:= \int k_{n, \gamma}^-(t |\dot{\gamma}|) |\dot{\gamma}|^2 d\pi(\gamma) \\ &= \int k_n(\gamma_{1-t}) |\dot{\gamma}|^2 d\pi(\gamma), \\ k_{n, \pi}^+(t \mathbf{c}_\pi) \mathbf{c}_\pi^2 &:= \int k_{n, \gamma}^+(t |\dot{\gamma}|) |\dot{\gamma}|^2 d\pi(\gamma) \\ &= \int k_n(\gamma_t) |\dot{\gamma}|^2 d\pi(\gamma); \end{aligned} \quad (3.11)$$

recall (3.5) and compare with Remark 3.3. In other words,

$$\begin{aligned} k_{n, \pi}^-(r) &= \frac{1}{\mathbf{c}_\pi^2} \int k_n(\bar{\gamma}_{\mathbf{c}_\pi - r}) |\dot{\gamma}|^2 d\pi(\gamma), \\ k_{n, \pi}^+(r) &= \frac{1}{\mathbf{c}_\pi^2} \int k_n(\bar{\gamma}_r) |\dot{\gamma}|^2 d\pi(\gamma), \end{aligned}$$

where $\bar{\gamma} : [0, \mathbf{c}_\pi] \rightarrow M$ designates the proper-time reparametrization $\bar{\gamma}_r := \gamma_{r/\mathbf{c}_\pi}$ with respect to \mathbf{c}_π . Since the involved integrands are all bounded, (3.11) is well-defined. The sequence $(k_{n, \pi}^\pm)_{n \in \mathbf{N}}$ converges pointwise to a lower semicontinuous function $k_\pi^\pm : [0, \mathbf{c}_\pi] \rightarrow \mathbf{R} \cup \{\infty\}$. By Levi’s monotone convergence theorem, the latter is related to the pointwise limit k_γ^\pm of $(k_{n, \gamma}^\pm)_{n \in \mathbf{N}}$ by the obvious analog of (3.11).

Similar arguments as those leading to Definition 3.4 yield the following notions and their well-definedness.

Definition 3.10 (Distortion coefficients along plans). *Given $0 \leq \theta \leq \mathbf{c}_\pi$ for $\pi \in \text{OptTGeo}_p(\mathcal{P}_c(M)^2)$, as in (3.1) we define*

- a. *the generalized sine functions $\sin_{k_\pi^\pm}$ associated with k_π^\pm by*

$$\sin_{k_\pi^\pm}(\theta) := \lim_{n \rightarrow \infty} \sin_{k_{n, \pi}^\pm}(\theta),$$

- b. *the distortion coefficients $\sigma_{k_\pi^\pm}^{(t)}$, where $0 \leq t \leq 1$, associated with k_π^\pm by*

$$\sigma_{k_\pi^\pm}^{(t)}(\theta) := \lim_{n \rightarrow \infty} \sigma_{k_{n, \pi}^\pm}^{(t)}(\theta).$$

It is not difficult to generalize the following lemma beyond $\text{OptTGeo}_p(\mathcal{P}_c(\mathbb{M})^2)$. We will only need it in this form. Furthermore, the implicit lower boundedness assumption on k is technical and will hold in all relevant cases since we mostly work with compactly supported measures.

Lemma 3.11 (Narrow lower semicontinuous dependence on plans). *Let $(\pi^i)_{i \in \mathbb{N}}$ be a sequence in $\text{OptTGeo}_p(\mathcal{P}_c(\mathbb{M})^2)$ converging narrowly to $\pi \in \mathcal{P}(C([0, 1]; \mathbb{M}))$. Furthermore, assume the endpoint marginal sequences $((e_0)_\# \pi^i)_{i \in \mathbb{N}}$ and $((e_1)_\# \pi^i)_{i \in \mathbb{N}}$ are uniformly compactly supported, say on the compact sets $C_0, C_1 \subset \mathbb{M}$, respectively. Then for every $0 \leq t \leq 1$,*

$$\sigma_{k_\pi^\pm}^{(t)}(\mathbf{c}_\pi) \leq \liminf_{i \rightarrow \infty} \sigma_{k_{\pi^i}^\pm}^{(t)}(\mathbf{c}_{\pi^i}).$$

Proof. By Alexandrov's theorem and causal closedness (Proposition 2.10), π is concentrated on causal curves contained in the compact set $J(C_0, C_1)$. Let $(k_n)_{n \in \mathbb{N}}$ be the sequence from (3.6) approximating the function k on $E := J(C_0, C_1)$. Given any $n \in \mathbb{N}$ and any $0 \leq t \leq 1$, continuity of k_n and l_+ , respectively, imply the lower semicontinuous dependence of the quantity $k_n(\gamma_t) |\dot{\gamma}|^2$ on $\gamma \in G$. Here, G is the uniform closure of the set of $\gamma \in \text{TGeo}(\mathbb{M})$ having $\gamma_0 \in C_0$ and $\gamma_1 \in C_1$. Hence, [129, Lem. 4.3] — here we use G being Polish, and uniform lower boundedness — and the definition (3.11) entail

$$k_{n,\pi}^+(t \mathbf{c}_\pi) \mathbf{c}_\pi^2 \leq \liminf_{i \rightarrow \infty} k_{n,\pi^i}^+(t \mathbf{c}_{\pi^i}) \mathbf{c}_{\pi^i}^2. \quad (3.12)$$

An analogous argument yields the same inequality for $k_{n,\pi}^+$ replaced by $k_{n,\pi}^-$.

Therefore, in view of Remark 3.3 we obtain

$$\begin{aligned} \sigma_{k_\pi^\pm}^{(t)} \mathbf{c}_\pi^2(1) &= \sup_{n \in \mathbb{N}} \sigma_{k_{n,\pi}^\pm}^{(t)} \mathbf{c}_\pi^2(1) \\ &\leq \sup_{n \in \mathbb{N}} \liminf_{i \rightarrow \infty} \sigma_{k_{n,\pi^i}^\pm}^{(t)} \mathbf{c}_{\pi^i}^2(1) \\ &\leq \liminf_{i \rightarrow \infty} \sup_{n \in \mathbb{N}} \sigma_{k_{n,\pi^i}^\pm}^{(t)} \mathbf{c}_{\pi^i}^2(1) \\ &\leq \liminf_{i \rightarrow \infty} \sigma_{k_{\pi^i}^\pm}^{(t)} \mathbf{c}_{\pi^i}^2(1). \end{aligned}$$

Here we have used analogs of Lemma 3.6 for general potentials, namely lower semicontinuity in the second inequality and monotonicity in the last one, as stated in [79, Lem. 3.9] and [80, Props. 3.4, 3.5]. The claim follows from (3.3). \square

The other parts of Lemma 3.6 carry over with appropriate modifications, taking into account that, unlike Subsection 3.1.3, here the potential can be infinite.

3.3. Definition and basic properties. We start with the infinite-dimensional analog of the variable curvature-dimension condition for metric measure spaces [122] à la Lott–Sturm–Villani [89, 120]. For constant k , this has been adapted to the Lorentzian context in [22]. Let $g: [0, 1]^2 \rightarrow \mathbf{R}$ be the Green's function with Dirichlet boundary conditions, cf. (3.7) of Subsection 3.1.4.

Definition 3.12 (Dimensionless variable timelike curvature-dimension condition). *We term \mathcal{M} to satisfy the timelike curvature-dimension condition $\text{TCD}_p(k, \infty)$ if for every timelike p -dualizable pair $(\mu_0, \mu_1) \in \mathcal{P}_c^{\text{ac}}(\mathbb{M}, \mathfrak{m})^2$, there exist*

- an ℓ_p -geodesic $(\mu_t)_{t \in [0, 1]}$ connecting μ_0 to μ_1 , and
- a plan $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$

such that for every $0 \leq t \leq 1$,

$$\text{Ent}_{\mathfrak{m}}(\mu_t) \leq (1 - t) \text{Ent}_{\mathfrak{m}}(\mu_0) + t \text{Ent}_{\mathfrak{m}}(\mu_1) - \int_0^1 \int g(s, t) k(\gamma_s) |\dot{\gamma}|^2 d\pi(\gamma) ds.$$

If the previous statement holds for merely every strongly timelike p -dualizable pair (μ_0, μ_1) as above, the space \mathcal{M} is said to satisfy the weak timelike curvature-dimension condition $\text{wTCD}_p(k, \infty)$.

Our next definition, taking into account an “upper dimension bound” N , is based on an adaptation of the abstract convexity from of the previous sections to the functionals \mathcal{U}_N and $\text{Ent}_{\mathfrak{m}}$ in place of U_N and S , respectively.

Definition 3.13 (Variable entropic timelike curvature-dimension condition). *We term \mathcal{M} to satisfy the entropic timelike curvature-dimension condition $\text{TCD}_p^e(k, N)$ if for every timelike p -dualizable pair $(\mu_0, \mu_1) \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathfrak{m})^2$, there exist*

- an ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ ⁷ connecting μ_0 to μ_1 , and
- a plan $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$

such that for every $0 \leq t \leq 1$, the functional \mathcal{U}_N from (2.28) satisfies

$$\mathcal{U}_N(\mu_t) \geq \sigma_{\frac{k^-}{\pi}/N}^{(1-t)}(\mathbf{c}_\pi) \mathcal{U}_N(\mu_0) + \sigma_{\frac{k^+}{\pi}/N}^{(t)}(\mathbf{c}_\pi) \mathcal{U}_N(\mu_1). \quad (3.13)$$

If the previous statement holds for merely every strongly timelike p -dualizable pair (μ_0, μ_1) as above, the space \mathcal{M} is said to satisfy the weak entropic timelike curvature-dimension condition $\text{wTCD}_p^e(k, N)$.

Remark 3.14 (Nonnegative potentials). If k is nonnegative, by (3.3), (3.10), and the monotonicity property stated in Lemma 3.6, the inequality (3.13) yields

$$\mathcal{U}_N(\mu_t) \geq \sigma_{\frac{k^-}{\pi}/N}^{(1-t)}(\ell_p(\mu_0, \mu_1)) \mathcal{U}_N(\mu_0) + \sigma_{\frac{k^+}{\pi}/N}^{(t)}(\ell_p(\mu_0, \mu_1)) \mathcal{U}_N(\mu_1).$$

In other words, $\text{Ent}_{\mathfrak{m}}$ is *weakly* (k, N, p) -convex on $\mathcal{P}_c(\mathbf{M})$ with respect to the obvious adaptation of [91, Def. 6.5] for variable curvature bounds.

Remark 3.15 (Time reversal invariance). By time reversal in (3.13), \mathcal{M} satisfies $\text{TCD}_p^e(k, N)$ if and only if its causal reversal does; analogously for $\text{wTCD}_p^e(k, N)$.

Remark 3.16 (Only the support matters). The support of \mathfrak{m} might fail to be a length metric spacetime in the strict sense of Hypothesis 2.8 by the possible lack of nontrivial chronology. In fact, as shown in Theorem 5.1 below, if \mathfrak{m} has an atom at $x \in \mathbf{M}$ then the chronological future and the chronological past of x relative to $\text{spt } \mathfrak{m}$ are empty. In addition, regularity of $\text{spt } \mathfrak{m}$ does not imply regularity of \mathbf{M} .

On the other hand, $\text{spt } \mathfrak{m}$ is l -geodesic under a $\text{wTCD}_p^e(k, N)$ condition, in the following sense. For every $x, y \in \text{spt } \mathfrak{m}$ with $x \ll y$ there is $\gamma \in \text{TGeo}(\mathbf{M})$ with $\gamma_0 = x$, $\gamma_1 = y$, and $\gamma_t \in \text{spt } \mathfrak{m}$ for every $0 \leq t \leq 1$; compare also with [44, Rem. 3.10]. Indeed, choose sufficiently small neighborhoods $U_x, U_y \subset \mathbf{M}$ around x and y , respectively, such that $U_x \times U_y \Subset \{l > 0\}$, apply Proposition 2.66 to the ℓ_p -geodesic between the uniform distributions of U_x and U_y certifying (3.13), and then argue as in [39, Rem. 6.11] for the obtained lifting.

Moreover, it still makes sense to say the (weak) entropic timelike curvature-dimension condition is satisfied by \mathbf{M} if and only if it holds for $\text{spt } \mathfrak{m}$ — indeed, every $(\mu_t)_{t \in [0,1]}$ as in Definition 3.13 obeys $\text{spt } \mu_t \subset \text{spt } \mathfrak{m}$ for every $0 \leq t \leq 1$. In this sense, $\text{TCD}_p^e(k, N)$ and $\text{wTCD}_p^e(k, N)$ are properties of $\text{spt } \mathfrak{m}$ rather than \mathbf{M} .

We now show further elementary properties of the notion from Definition 3.13. They are only stated for the weak timelike entropic curvature-dimension condition, but are valid analogously if every occurrence of “ wTCD^e ” is replaced by “ TCD^e ”.

⁷One could alternately start with *rough* ℓ_p -geodesics here. For $\text{wTCD}_p^e(k, N)$, this makes no difference by Corollary 2.64. For $\text{TCD}_p^e(k, N)$, by Theorem 3.27, a posteriori one can always choose a narrowly continuous collection $(\mu_t)_{t \in [0,1]}$ certifying (3.13) if \mathcal{M} is timelike p -essentially nonbranching.

We start by stating the following natural Lorentzian adaptation of isomorphisms of metric measure spacetimes, cf. e.g. [66, 70, 71, 120] for the metric measure case. This will be studied in detail in [27]; in particular, in some cases the inherent — quite strong — topological hypotheses can be dropped.

Definition 3.17 (Metric measure spacetime isomorphism). *Let \mathcal{M} and \mathcal{M}' be two metric measure spacetimes. A map $\iota: \text{spt } \mathfrak{m} \rightarrow \text{spt } \mathfrak{m}'$ is called isomorphism if it is a homeomorphism, it is measure-preserving — i.e. $\iota_{\#}\mathfrak{m} = \mathfrak{m}'$ —, and for every $x, y \in \text{spt } \mathfrak{m}$ we have*

$$l'_+(\iota(x), \iota(y)) = l_+(x, y).$$

If such an isomorphism exists, we call \mathcal{M} and \mathcal{M}' isomorphic.

Proposition 3.18 (Basic properties). *Let $k': \mathbb{M} \rightarrow \mathbf{R}$ be lower semicontinuous, and let $N' > 0$ as well as $\alpha, \beta > 0$. Then the following properties hold.*

- (i) **Isomorphism.** *Assume \mathcal{M} and \mathcal{M}' to be isomorphic through ι according to Definition 3.17. If \mathcal{M} obeys $\text{wTCD}_p^e(k, N)$, then \mathcal{M}' obeys $\text{wTCD}_p^e(\iota_{\#}k, N)$, where $\iota_{\#}k := k \circ \iota$.*
- (ii) **Consistency.** *If $k' \geq k$ on \mathbb{M} and $N' \leq N$, and \mathcal{M} obeys $\text{wTCD}_p^e(k', N')$, then it also satisfies $\text{wTCD}_p^e(k, N)$.*
- (iii) **Scaling.** *If \mathcal{M} obeys $\text{wTCD}_p^e(k, N)$, then its rescaling $\mathcal{M}_{\alpha, \beta} := (\mathbb{M}, \alpha l, \beta \mathfrak{m})$ satisfies $\text{wTCD}_p^e(k/\alpha^2, N)$.*

These statements hold analogously by setting $N = \infty$.

Proof. The claim (i) follows from the subsequent properties of isomorphy whose straightforward proofs are left to the reader. It is clear from Definition 3.17 that $\iota_{\#}k$ is still lower semicontinuous on $\text{spt } \mathfrak{m}'$.

- The inverse ι^{-1} pushes forward every \mathfrak{m}' -absolutely continuous measures $\mu'_0, \mu'_1 \in \mathcal{P}_c^{\text{ac}}(\mathbb{M}', \mathfrak{m}')$ to \mathfrak{m} -absolutely continuous measures μ_0 and μ_1 on J , and it preserves compactness of supports.
- If π' is a chronological ℓ'_p -optimal coupling of μ'_0 and μ'_1 , then $(\iota^{-1}, \iota^{-1})_{\#}\pi$ is a chronological ℓ_p -cyclically monotone coupling of μ_0 and μ_1 , and thus ℓ_p -optimal [44, Prop. 2.8].
- If the pair (μ'_0, μ'_1) is strongly timelike p -dualizable relative to an $(\ell')^p$ -cyclically monotone set $\Gamma' \subset \mathbb{M}'$ as in Definition 2.57, then (μ_0, μ_1) is strongly timelike p -dualizable relative to the ℓ_p -cyclically monotone set $\Gamma := (\iota^{-1}, \iota^{-1})(\Gamma')$.
- If $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ then $\pi' := (\iota_{\#})_{\#}\pi \in \text{OptTGeo}'_p(\mu'_0, \mu'_1)$, where we set $\iota_{\#}(\gamma) := \iota \circ \gamma$. The respective L^2 -costs coincide,

$$\|l \circ (\mathbf{e}_0, \mathbf{e}_1)\|_{L^2(\text{TGeo}(\mathbb{M}), \pi)} = \|l' \circ (\mathbf{e}_0, \mathbf{e}_1)\|_{L^2(\text{TGeo}(\mathbb{M}'), \pi')}.$$

In particular, recalling that $\iota_*k := k \circ \iota$, for every $0 \leq t \leq 1$ we have

$$\sigma_{k_{\#}^{\pm}/N}^{(t)}(\mathbf{c}_{\pi}) = \sigma_{(\iota_*k)_{\#}^{\pm}/N}^{(t)}(\mathbf{c}'_{\pi'}).$$

- If $(\mu_t)_{t \in [0,1]}$ is an ℓ_p -geodesic consisting of \mathfrak{m} -absolutely continuous measures, so is $(\mu'_t)_{t \in [0,1]}$ relative to ℓ'_p , where $\mu'_t := \iota_{\#}\mu_t$, and

$$\text{Ent}_{\mathfrak{m}}(\mu_t) = \text{Ent}_{\mathfrak{m}'}(\mu'_t).$$

Item (ii) follows from the monotonicity properties of the distortion coefficients asserted in Lemma 3.6.

Finally, (iii) is a consequence of the identities $\text{Ent}_{\beta \mathfrak{m}}(\mu) = \text{Ent}_{\mathfrak{m}}(\mu) - \log \beta$ for every $\mu \in \mathcal{P}_c(\mathbb{M})$, 1-homogeneity of the L^2 -norm, and (3.3). \square

Corollary 3.19 (Dimensional consistency). *The condition $\text{wTCD}_p^e(k, N)$ implies $\text{wTCD}_p(k, \infty)$.*

Proof. By Proposition 3.18, Definition 3.10, as well as Levi's monotone convergence theorem we may and will assume that k is continuous and bounded.

Let $\mu_0, \mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$ be as in Definition 3.12. Assume $\mu_0, \mu_1 \in \mathcal{D}(\text{Ent}_{\mathbf{m}})$; otherwise the claim is clear. Again by Proposition 3.18, let $(\mu_t)_{t \in [0,1]}$ and $\pi \in \Pi_{\ll}(\mu_0, \mu_1)$ certify the displacement semiconvexity of $\mathcal{U}_{N'}$ according to Definition 3.13 for every $N' \geq N$ (noting Lemma 3.9). Jensen's inequality and Lemma 2.63 imply $\text{Ent}_{\mathbf{m}}(\mu_t) > -\infty$ uniformly in $0 \leq t \leq 1$, while $\text{wTCD}_p^e(k, N)$, finiteness of \mathbf{m} on compact sets, boundedness of k , and Lemma 3.6 yield $\text{Ent}_{\mathbf{m}}(\mu_t) < \infty$ uniformly in $0 \leq t \leq 1$. Therefore,

$$\begin{aligned} & (1-t)\text{Ent}_{\mathbf{m}}(\mu_0) + t\text{Ent}_{\mathbf{m}}(\mu_1) - \text{Ent}_{\mathbf{m}}(\mu_t) \\ &= (1-t) \lim_{N \rightarrow \infty} N [1 - \mathcal{U}_N(\mu_0)] + t \lim_{N \rightarrow \infty} N [1 - \mathcal{U}_N(\mu_1)] \\ & \quad + \lim_{N \rightarrow \infty} N [\mathcal{U}_N(\mu_t) - 1] \\ & \geq \liminf_{N \rightarrow \infty} N [\sigma_{k_{\pi}^- c_{\pi}^2/N}^{(1-t)}(\mathbf{c}_{\pi}) + \sigma_{k_{\pi}^+ c_{\pi}^2/N}^{(t)}(\mathbf{c}_{\pi}) - 1]. \end{aligned}$$

In the last line, the entropy terms have been dropped since they converge to 1.

The function $u: [0, 1] \rightarrow \mathbf{R}$ defined by

$$u(t) := N [\sigma_{k_{\pi}^- c_{\pi}^2/N}^{(1-t)} + \sigma_{k_{\pi}^+ c_{\pi}^2/N}^{(t)} - 1]$$

satisfies Dirichlet boundary conditions $u(0) = u(1) = 0$. Moreover, by Remark 3.3, the hypotheses on k , and (3.11), for every $t \in (0, 1)$ it satisfies

$$\begin{aligned} u''(t) &= -k_{\pi}^+(t \mathbf{c}_{\pi}) c_{\pi}^2 [\sigma_{k_{\pi}^- c_{\pi}^2/N}^{(1-t)} + \sigma_{k_{\pi}^+ c_{\pi}^2/N}^{(t)}] \\ &= - \int k(\gamma_t) |\dot{\gamma}|^2 [\sigma_{k_{\pi}^- c_{\pi}^2/N}^{(1-t)} + \sigma_{k_{\pi}^+ c_{\pi}^2/N}^{(t)}] d\pi(\gamma), \end{aligned}$$

and therefore, e.g. by [79, Prop. 3.8],

$$u(t) = \int_0^1 \int g(s, t) k(\gamma_s) |\dot{\gamma}|^2 [\sigma_{k_{\pi}^- c_{\pi}^2/N}^{(1-s)} + \sigma_{k_{\pi}^+ c_{\pi}^2/N}^{(s)}] d\pi(\gamma) ds.$$

By the discussion from Subsection 3.1.1 we finally have

$$\lim_{N \rightarrow \infty} \sup_{s \in [0,1]} |\sigma_{k_{\pi}^- c_{\pi}^2/N}^{(1-s)} + \sigma_{k_{\pi}^+ c_{\pi}^2/N}^{(s)} - 1| = 0,$$

from which the claim follows. \square

Lastly, we study measure perturbations. Unfortunately, since the arguments of the distortion coefficients in Definition 3.13 depend on the entire plan π and not only on its endpoint marginals, the more restrictive additional assumption of timelike p -essential nonbranching seems needed; see also Corollary 3.25.

Proposition 3.20 (Variable timelike curvature dimension conditions with densities). *Assume $\text{wTCD}_p^e(k, N)$, and assume \mathcal{M} is timelike p -essentially nonbranching. Moreover, let $V: \text{spt } \mathbf{m} \rightarrow \mathbf{R}$ be a Borel function which is locally bounded and (k', N') -convex in the sense of Definition 3.8 for a lower semicontinuous function $k': \mathbf{M} \rightarrow \mathbf{R}$ and $N' > 0$. Lastly, define the measure $\mathbf{m}_V := e^{-V} \mathbf{m}$. Then the weighted globally hyperbolic, regular length metric measure spacetime $\mathcal{M}_V := (\mathbf{M}, l, \mathbf{m}_V)$ satisfies $\text{wTCD}_p^e(k + k', N + N')$.*

Proof. We first prove a (k, N) -convexity-type inequality similar to Definition 3.8 and Definition 3.13 for the functional $V^*: \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m}) \rightarrow \mathbf{R}$ given by

$$V^*(\mu) := \int_{\mathbf{M}} V d\mu.$$

We note that $\mathcal{D}(V^*) = \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$. Given any $0 \leq t \leq 1$, recall the convex function G_t from Lemma 3.6. Let $(\mu_t)_{t \in [0,1]}$ be any ℓ_p -geodesic interpolating two given measures $\mu_0, \mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$, passing through $\mathcal{D}(V^*)$, and represented by a plan $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$, cf. Corollary 3.25 below. Thanks to our hypothesis on V and Jensen's inequality,

$$\begin{aligned} -\frac{1}{N'} V^*(\mu_t) &= -\frac{1}{N'} \int V(\gamma_t) d\pi(\gamma) \\ &\geq \int G_t \left[-\frac{1}{N'} V(\gamma_0), -\frac{1}{N'} V(\gamma_1), \frac{1}{N'} k'_\gamma |\dot{\gamma}|^2 \right] d\pi(\gamma) \\ &\geq G_t \left[-\frac{1}{N'} V^*(\mu_0), -\frac{1}{N'} V^*(\mu_1), \frac{1}{N'} k'_\pi \mathbf{c}_\pi^2 \right]. \end{aligned}$$

Now we turn to the actual statement. Assume μ_0 and μ_1 to be strongly timelike p -dualizable in addition. Let $(\mu_t)_{t \in [0,1]}$ be the ℓ_p -geodesic connecting μ_0 to μ_1 represented by $\pi' \in \text{OptTGeo}_p(\mu_0, \mu_1)$, shown to exist and be unique in Corollary 3.25 below, and also let $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ certify the displacement semiconvexity of \mathcal{U}_N defining $\text{wTCD}_p^e(k, N)$. Recall that

$$\text{Ent}_{\mathbf{m}_V}(\mu) = \text{Ent}_{\mathbf{m}}(\mu) + V^*(\mu) \quad (3.14)$$

for every $\mu \in \mathcal{P}_c(\mathbf{M})$, and define $\mathcal{U}_{N+N'}^V : \mathcal{P}_c(\mathbf{M}) \rightarrow [0, \infty]$ by

$$\mathcal{U}_{N+N'}^V(\mu) := e^{-\text{Ent}_{\mathbf{m}_V}(\mu)/(N+N')}.$$

Given any $0 \leq t \leq 1$, (3.14) and the previous considerations yield

$$\begin{aligned} \log \mathcal{U}_{N+N'}^V(\mu_t) &= -\frac{N}{N+N'} \frac{1}{N} \text{Ent}_{\mathbf{m}}(\mu_t) - \frac{N'}{N+N'} \frac{1}{N'} V^*(\mu_t) \\ &\geq \frac{N}{N+N'} G_t \left[-\frac{1}{N} \text{Ent}_{\mathbf{m}}(\mu_0), -\frac{1}{N} \text{Ent}_{\mathbf{m}}(\mu_1), \frac{1}{N} k_\pi \mathbf{c}_\pi^2 \right] \\ &\quad + \frac{N'}{N+N'} G_t \left[-\frac{1}{N'} V^*(\mu_0), -\frac{1}{N'} V^*(\mu_1), \frac{1}{N'} k_{\pi'} \mathbf{c}_{\pi'}^2 \right] \\ &\geq G_t \left[-\frac{1}{N+N'} \text{Ent}_{\mathbf{m}_V}(\mu_0), -\frac{1}{N+N'} \text{Ent}_{\mathbf{m}_V}(\mu_1), \right. \\ &\quad \left. \frac{1}{N+N'} (k+k')_\pi \mathbf{c}_\pi^2 \right]. \end{aligned}$$

In the last estimate, we have used Jensen's inequality, the linearity (3.11) of

$$k_\pi^\pm \mathbf{c}_\pi^2 + k_{\pi'}^\pm \mathbf{c}_{\pi'}^2 = (k+k')_\pi^\pm \mathbf{c}_\pi^2 \quad (3.15)$$

in k , and the fact that $\pi = \pi'$ ensured by Theorem 3.24 below. \square

3.4. Qualitative properties from local(ly uniform) curvature bounds. Unlike balls in metric geometry, causal diamonds are always (causally) convex in a metric spacetime. Hence, by local lower boundedness of k and Proposition 3.18, every $\text{wTCD}_p^e(k, N)$ space is locally a $\text{wTCD}_p^e(K, N)$ space for some $K \in \mathbf{R}$; the same is true with both w 's removed.

Lemma 3.21 (From variable to constant bounds). *Assume $\text{wTCD}_p^e(k, N)$. Let $C_0, C_1 \subset \mathbf{M}$ be two compact sets. Then $K := \inf k(J(C_0, C_1)) \in \mathbf{R}$, and the metric measure spacetime structure on $J(C_0, C_1)$ induced by \mathcal{M} obeys $\text{wTCD}_p^e(K, N)$.*

An analogous claim holds clearly for the nonweak entropic timelike curvature-dimension condition, using Remark 2.15 to make sense of both weak and nonweak entropic timelike curvature-dimension condition for subsets of \mathbf{M} , in spite of the fact that Definition 3.13 is stated only for length metric spacetimes.

This result allows us to transfer qualitative properties from the variable to the constant case. The relevant ones for our purposes treated now are

- the uniqueness of chronological ℓ_p -optimal couplings,
- the uniqueness of ℓ_p -geodesics, and
- the existence of ℓ_p -geodesics with uniformly bounded densities.

Although k might not be uniformly bounded from below on M , the proofs of all results below from [21, 22, 44] are local in nature and can thus be proven with a local uniform lower bound on k , as in Lemma 3.21. We will extend the above mentioned results in two ways.

- First, recall our ℓ_p -geodesics are defined differently than in [21, 22]. Since [21, 44] essentially establish uniqueness of *displacement* ℓ_p -geodesics, and in our case not every ℓ_p -geodesic is lifted by a plan, this requires additional care (although the proofs will be similar to [21, 44]).
- Second, slightly generalizing [22, 44] we show the above mentioned results to hold under a mere local entropic timelike curvature-dimension condition, as stated in Definition 3.22 below.

3.4.1. Uniqueness of chronological ℓ_p -optimal couplings from local curvature bounds.

Definition 3.22 (Local timelike curvature-dimension bounds). *We say \mathcal{M} obeys the local entropic timelike curvature-dimension condition $\text{TCD}_{p,\text{loc}}^e(k, N)$ if every point $o \in \text{spt } \mathbf{m}$ has a neighborhood $U \subset M$ with the subsequent property. For every $(\mu_0, \mu_1) \in \mathcal{P}_c^{\text{ac}}(M, \mathbf{m})^2$ with $\text{spt } \mu_0, \text{spt } \mu_1 \subset U$ and*

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}, \quad (3.16)$$

there exist

- an ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ connecting μ_0 to μ_1 , and
- a plan $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$

such that for every $0 \leq t \leq 1$, the functional \mathcal{U}_N from (2.28) satisfies

$$\mathcal{U}_N(\mu_t) \geq \sigma_{\frac{k^-}{N}}^{(1-t)}(\mathbf{c}\pi) \mathcal{U}_N(\mu_0) + \sigma_{\frac{k^+}{N}}^{(t)}(\mathbf{c}\pi) \mathcal{U}_N(\mu_1).$$

Note that the above ℓ_p -geodesic is defined with respect to the ambient structure from \mathcal{M} and thus, in particular, allowed to leave U .

By Example 2.59, $\text{TCD}_{p,\text{loc}}^e(k, N)$ is clearly implied by $\text{wTCD}_p^e(k, N)$. Moreover, time reversal (Remark 3.15) implies the \mathbf{m} -absolute continuity of μ_1 can be substituted for that of μ_0 in the hypotheses of the following theorem.

Theorem 3.23 (Maps and uniqueness for chronological ℓ_p -optimal couplings). *We assume \mathcal{M} is a timelike p -essentially nonbranching $\text{TCD}_{p,\text{loc}}^e(k, N)$ space, and let $(\mu_0, \mu_1) \in \mathcal{P}(M)^2$ be timelike p -dualizable by $\pi \in \Pi_{\ll}(\mu_0, \mu_1)$. Lastly, assume μ_0 to be \mathbf{m} -absolutely continuous. Then there is a μ_0 -measurable map $T: M \rightarrow M$ with*

$$\pi = (\text{Id}, T)_{\#} \mu_0.$$

In particular, π is the unique chronological ℓ_p -optimal coupling of μ_0 and μ_1 .

Proof. We assume \mathbf{m} -absolute continuity of μ_0 with density ρ_0 . By the local compactness of τ , cf. Corollary 2.17, and the usual contradiction argument, cf. the proofs of [22, Thm. 4.15] and [44, Thm. 3.20], it is not restrictive to assume the compactness of $C_0 := \text{spt } \mu_0$ and $C_1 := \text{spt } \mu_1$. Furthermore, by modifying π and [44, Lem. 2.10], we may and will assume without restriction $\|\rho_0\|_{L^\infty(M, \mathbf{m})} < \infty$, so that $\mu_0 \in \mathcal{D}(\text{Ent}_{\mathbf{m}})$, as well as the existence of $r > 0$ with

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l \geq r\}.$$

In turn, up to replacing M with $J(\text{spt } \mu_0, \text{spt } \mu_1)$, for notational convenience we assume compactness of M (which does not conflict with Remark 2.15).

By $\text{TCD}_{p,\text{loc}}^e(k, N)$, there exist $n \in \mathbb{N}$ and relatively open sets $U_1, \dots, U_n \subset M$ covering M such that for every $i = 1, \dots, n$, the statement of the theorem is true

for every pair of timelike p -dualizable measures with support in U_i . This can be shown by solely relying on *displacement* ℓ_p -geodesics. Indeed, first note that every ℓ_p -geodesic as in Definition 3.22 is a displacement ℓ_p -geodesic by Proposition 2.66. Lemma 3.21, a local form of Proposition A.3, and again Proposition 2.66 allow us to show a local (entropic, i.e. in terms of \mathcal{U}_N) version of [22, Thm. 4.12] within U_i for every $i = 1, \dots, n$, where “timelike proper-time parametrized ℓ_p -geodesic” is replaced by “displacement ℓ_p -geodesic” in our case; cf. [21, Thm. 4.11]. Following the arguments for [22, Lems. 4.13, 4.14, Thm. 4.15] we obtain the claim.

Consequently, by following the proof of [22, Thm. 4.16], for every $i = 1, \dots, n$ we also get uniqueness of *displacement* ℓ_p -geodesics connecting timelike p -dualizable pairs both supported in U_i with \mathfrak{m} -absolutely continuous initial distribution.

We now follow the argument for [42, Cor. 5.4] to get prove the claim for μ_0 and μ_1 as above. By Lebesgue’s number lemma, there exists $\varepsilon > 0$ such that if any subset $E \subset \mathbb{M}$ has diameter less than ε , then $E \subset U_i$ for some $i = 1, \dots, n$. Let G_r denote the compact set of all $\gamma \in \text{TGeo}(\mathbb{M})$ with $\gamma_0 \in C_0$ and $\gamma_1 \in C_1$, cf. Corollary 2.47. By Arzelà–Ascoli’s theorem, there exists $0 < \delta < \varepsilon/5$ such that if $0 \leq s, t \leq 1$ satisfies $|s - t| < \delta$, then every $\gamma \in G_r$ obeys $\mathbf{d}(\gamma_s, \gamma_t) < \varepsilon/5$.

Let $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ satisfy $\pi = (\mathbf{e}_0, \mathbf{e}_1)_{\#} \pi$, whose existence is granted by Proposition 2.66. Set $\mu_t := (\mathbf{e}_t)_{\#} \pi$, where $0 \leq t \leq 1$. Given any $0 \leq t < 1$, let $\pi^t : \mathbb{M} \rightarrow \mathcal{P}(C([0, 1]; \mathbb{M}))$ denote the disintegration of π with respect to \mathbf{e}_t , i.e.

$$\mathbf{d}\pi(\gamma) = \mathbf{d}\pi_x^t(\gamma) \, \mathrm{d}\mu_t(x).$$

By the triangle inequality, given any $z \in \mathbb{M}$ we obtain

$$\text{diam} \left[\text{spt } \mu_t \llcorner \mathbf{B}^{\mathbf{d}}(z, \delta) \cup \text{spt}(\mathbf{e}_{t+\delta/2})_{\#} \left[\int_{\mathbf{B}^{\mathbf{d}}(z, \delta)} \pi_x^t \, \mathrm{d}\mu_t(x) \right] \right] < \varepsilon. \quad (3.17)$$

Now we argue successively to get the desired map $T : \mathbb{M} \rightarrow \mathbb{M}$ which satisfies $\pi = (\text{Id}, T)_{\#} \mu_0$. We start the process from $t = 0$. Covering $\text{spt } \mu_0$ with finitely many δ -balls with respect to \mathbf{d} , by (3.17), the mentioned Lebesgue number lemma, and pasting optimal maps together we construct a μ_0 -measurable map $T_{\delta/3} : \mathbb{M} \rightarrow \mathbb{M}$ such that $(\mathbf{e}_0, \mathbf{e}_{\delta/3})_{\#} \pi = (\text{Id}, T_{\delta/3})_{\#} \mu_0$. We repeat this process finitely many times — say $n \in \mathbb{N}$, chosen to be the smallest integer with $(n+1)\delta/3 \geq 1$ — by starting from $t = i\delta/3$ for every $i = 1, \dots, n$, thereby obtaining a $\mu_{i\delta/3}$ -measurable map $T_{(i+1)\delta/3} : \mathbb{M} \rightarrow \mathbb{M}$ such that $(\mathbf{e}_{i\delta/3}, \mathbf{e}_{(i+1)\delta/3})_{\#} \pi = (\text{Id}, T_{(i+1)\delta/3})_{\#} \mu_{i\delta/3}$. Here, (3.17) and local versions of Proposition A.3 as well as the qualitative result [22, Lem. 4.14] — extended to timelike p -dualizable endpoints as in Step 4 in the proof of [44, Prop. 3.19] using Lemma 2.71 — inductively yield that for every $0 = i_1, \dots, i_n$, there exists at least one *displacement* ℓ_p -geodesic $(\nu_t)_{t \in [0, 1]}$ from $\mu_{i\delta/3}$ to $\mu_{(i+3/2)\delta/3}$ with the following properties.

- The measure ν_t is \mathfrak{m} -absolutely continuous for every $0 \leq t < 1$.
- For every $\eta > 0$ the respective \mathfrak{m} -density of ν_t belongs to $L^\infty(\mathbb{M}, \mathfrak{m})$, uniformly in $0 \leq t \leq 1 - \eta$.

By the above uniqueness results, it is necessarily represented by $(\text{restr}_{i\delta/3}^{(i+3/2)\delta/3})_{\#} \pi$. Thus, μ_t is \mathfrak{m} -absolutely continuous for every $i \in 0, \dots, n$ and every $i\delta/3 \leq t \leq (i+1)\delta/3$ with $t < 1$, and the respective \mathfrak{m} -densities are uniformly in $L^\infty(\mathbb{M}, \mathfrak{m})$. The maps obtained that way readily yield the claimed T by composition. \square

3.4.2. Maps and uniqueness of ℓ_p -geodesics from local curvature bounds. With a similar reasoning based upon [44, Prop. 3.19, Thm. 3.21] and again [21, Lem. 4.14, Thms. 4.16, 4.19] with Lemma 2.63, the subsequent Theorem 3.24 is readily verified.

Theorem 3.24 (Maps, uniqueness, and regularity of displacement ℓ_p -geodesics). *Retain the assumptions and notation from Theorem 3.23. Then $\text{OptTGeo}_p(\mu_0, \mu_1) =$*

$\{\pi\}$ and there is a μ_0 -measurable map $\mathfrak{T}: \mathbb{M} \rightarrow \text{TGeo}(\mathbb{M})$ with

$$\pi = \mathfrak{T}_\# \mu_0.$$

In particular, μ_0 and μ_1 are connected by exactly one displacement ℓ_p -geodesic; the latter is represented by the plan π .

Lastly, if $\mu_0 \in \mathcal{P}_c(\mathbb{M}) \cap \mathcal{D}(\text{Ent}_m)$ and $\mu_1 \in \mathcal{P}_c(\mathbb{M})$ above, then we have $(e_t)_\# \pi \in \mathcal{P}_c(\mathbb{M}) \cap \mathcal{D}(\text{Ent}_m)$ for every $0 \leq t < 1$.

Corollary 3.25 (Strongly timelike p -dualizable, (compactly supported) densities have unique (rough) ℓ_p -geodesics). *Retain the assumptions of Theorem 3.23 for \mathcal{M} . Moreover, let the pair $(\mu_0, \mu_1) \in \mathcal{P}(\mathbb{M})^2$ be timelike p -dualizable, and assume all ℓ_p -optimal couplings of μ_0 and μ_1 to be chronological. Moreover, suppose \mathfrak{m} -absolute continuity of μ_0 or μ_1 . Then there exists exactly one ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ from μ_0 to μ_1 ; it is in fact a displacement ℓ_p -geodesic.*

If in addition μ_0 and μ_1 have compact support, then there exists exactly one rough ℓ_p -geodesic joining them.

Proof. By the first part of Proposition 2.66, there exists at least one displacement ℓ_p -geodesic from μ_0 to μ_1 . By Remark 2.62, this is also an ℓ_p -geodesic.

Let $(\mu_t)_{t \in [0,1]}$ be an ℓ_p -geodesic from μ_0 to μ_1 . To show its uniqueness, we revisit the proof of (v) in Proposition 2.66. Retaining its notation, we note that since every ℓ_p -optimal coupling of μ_0 and μ_1 is chronological and by Theorem 3.24, the constructed sequence $(\pi^n)_{n \in \mathbb{N}}$ interpolating $(\mu_t)_{t \in [0,1]}$ at dyadic points in $[0,1]$ is constant, say equal to $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$. By narrow continuity of $(\mu_t)_{t \in [0,1]}$, the latter is represented by π . The claim follows from Theorem 3.24.

The claim about uniqueness of rough ℓ_p -geodesics is clear, since Corollary 2.64 asserts every rough ℓ_p -geodesic from μ_0 to μ_1 to be an ℓ_p -geodesic. \square

We do not know whether uniqueness of rough ℓ_p -geodesics holds under the first set of assumptions, i.e. without compact supports; recall Remark 2.65.

3.4.3. *Good geodesics.* The following adaptation of [21, Thm. 1.2] is not used in our work, but is a crucial technical ingredient to develop a rich calculus in our setting as initiated in [14].

Theorem 3.26 (Existence of good geodesics). *Assume $\text{wTCD}_p^c(k, N)$. Then every pair $(\mu_0, \mu_1) = (\rho_0 \mathfrak{m}, \rho_1 \mathfrak{m}) \in \mathcal{P}_c^{\text{ac}}(\mathbb{M}, \mathfrak{m})^2$ of strongly timelike p -dualizable measures with $\rho_0, \rho_1 \in L^\infty(\mathbb{M}, \mathfrak{m})$ are joined by an ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ such that*

- (i) μ_t is \mathfrak{m} -absolutely continuous with density ρ_t for every $0 \leq t \leq 1$, with

$$\sup_{t \in [0,1]} \|\rho_t\|_{L^\infty(\mathbb{M}, \mathfrak{m})} < \infty,$$

- (ii) for every $0 \leq t \leq 1$, there exists $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ such that

$$\mathcal{U}_N(\mu_t) \geq \sigma_{K/N}^{(1-t)}(\mathfrak{c}_\pi) \mathcal{U}_N(\mu_0) + \sigma_{K/N}^{(t)}(\mathfrak{c}_\pi) \mathcal{U}_N(\mu_1),$$

where $K := \inf k(J(\text{spt } \mu_0, \text{spt } \mu_1))$.

Proof. We only outline the proof; the strategy of [22] carries over almost verbatim. However, our framework a priori lacks a lifting theorem for general strongly timelike p -dualizable measures (unless \mathcal{M} is timelike p -essentially nonbranching, cf. Corollary 3.25) — the latter has been relied on in [22] quite extensively. The necessary modification is to consider certain sets of t -intermediate points introduced in Section B.4 below. Following [110] (which has inspired [22]), the minimization procedure from [22] is then performed on this set instead of sets of plans. Strong timelike p -dualizability propagates thanks to an analog of Lemma 2.67, and the limit geodesic is constructed from countably many intermediate measures as in the proof of Corollary 2.47 by using antisymmetry (Theorem B.5). \square

3.5. Pathwise characterization. Our goal now to show that $\text{TCD}_p^e(k, N)$ admits a pathwise characterization along ℓ_p -geodesics under timelike p -essential nonbranching. In this setting, we also show the equivalence of $\text{TCD}_p^e(k, N)$ and $\text{wTCD}_p^e(k, N)$. For constant k , this is due to [21, Thm. 3.35], see also [54, Thm. 3.12] and [91, Thm. 7.1]. In the smooth framework, the corresponding estimate (3.18) is precisely the one obtained from Raychaudhuri's equation for the transport Jacobian [26, 91], see also Section 4.1.

Theorem 3.27 (Distorted displacement concavity of mean free path). *Assume \mathcal{M} to be timelike p -essentially nonbranching. Then the following are equivalent.*

- (i) *The condition $\text{TCD}_p^e(k, N)$ holds.*
- (ii) *The condition $\text{wTCD}_p^e(k, N)$ holds.*
- (iii) *For every timelike p -dualizable pair $(\mu_0, \mu_1) = (\rho_0 \mathbf{m}, \rho_1 \mathbf{m}) \in \mathcal{P}^{\text{ac}}(\mathbf{M}, \mathbf{m})^2$, there is a displacement ℓ_p -geodesic from μ_0 to μ_1 which is represented by $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ such that for every $0 \leq t \leq 1$, we have $(e_t)_\# \pi = \rho_t \mathbf{m} \in \mathcal{P}^{\text{ac}}(\mathbf{M}, \mathbf{m})$, and π -a.e. $\gamma \in \text{TGeo}(\mathbf{M})$ satisfies*

$$\rho_t(\gamma_t)^{-1/N} \geq \sigma_{k\gamma/N}^{(1-t)}(|\dot{\gamma}|) \rho_0(\gamma_0)^{-1/N} + \sigma_{k\gamma/N}^{(t)}(|\dot{\gamma}|) \rho_1(\gamma_1)^{-1/N}. \quad (3.18)$$

Proof. It is evident that (i) implies (ii).

Next, we show (iii) implies (i). Let $(\mu_0, \mu_1) = (\rho_0 \mathbf{m}, \rho_1 \mathbf{m}) \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})^2$, and assume $\mu_0 \in \mathcal{D}(\text{Ent}_{\mathbf{m}})$. Furthermore, letting $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ be as given by (iii) as well as $0 < t < 1$, compactness of $J(\text{spt } \mu_0, \text{spt } \mu_1)$ implies the existence of $K \leq 0$ with $k \geq K$ on $J(\text{spt } \mu_0, \text{spt } \mu_1)$. By (3.18), π -a.e. $\gamma \in \text{TGeo}(\mathbf{M})$ thus satisfies the estimate

$$\rho_t(\gamma_t) \leq \sup \sigma_{K/N}^{(1-t)}(l_+(x, y))^{-N} \rho_0(\gamma_0).$$

The supremum is taken over all $x \in \text{spt } \mu_0$ and $y \in \text{spt } \mu_1$. In fact, it is finite by Lemma 3.6 as well as compactness of $\text{spt } \mu_0$ and $\text{spt } \mu_1$. In particular, we obtain $\mu_t \in \mathcal{D}(\text{Ent}_{\mathbf{m}})$. Letting G_t denote the convex function from Lemma 3.6, from (3.18) again and Jensen's inequality we get

$$\begin{aligned} -\frac{1}{N} \text{Ent}_{\mathbf{m}}(\mu_t) &= -\frac{1}{N} \int \log \rho_t(\gamma_t) \, d\pi(\gamma) \\ &\geq \int G_t \left[-\frac{1}{N} \log \rho_0(\gamma_0), -\frac{1}{N} \log \rho_1(\gamma_1), k_\gamma |\dot{\gamma}|^2 \right] \, d\pi(\gamma) \\ &\geq G_t \left[-\frac{1}{N} \text{Ent}_{\mathbf{m}}(\mu_0), -\frac{1}{N} \text{Ent}_{\mathbf{m}}(\mu_1), \frac{1}{N} k_\pi c_\pi^2 \right]. \end{aligned}$$

Exponentiating both sides yields the desired inequality.

The case $\mu_1 \in \mathcal{D}(\text{Ent}_{\mathbf{m}})$ is argued analogously. If neither μ_0 nor μ_1 belong to $\mathcal{D}(\text{Ent}_{\mathbf{m}})$, necessarily $\mathcal{U}_N(\mu_0) = \mathcal{U}_N(\mu_1) = 0$ by (2.29), in which case (3.13) is clear.

The implication from (ii) to (iii) if the pair (μ_0, μ_1) from (iii) satisfies

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}$$

is shown in the same way as in the proof of [21, Prop. 3.38] with the modifications described in the proof of [21, Prop. 3.40], employing Lemma 2.71 and Lemma 3.6.

The general case is argued as for e.g. [21, Prop. 3.38] or [91, Thm. 7.1]. The idea is to decompose the support of a timelike p -dualizing coupling into countably many rectangles compactly contained in $\{l > 0\}$, and to restrict π to these, modulo successive removal of overlapping supports. Then we apply (iii) to these restrictions. Taking an appropriate convex combination of the plans constructed in that way yields the desired plan. This works since Lemma 2.71 ensures that the individual t -slices are mutually singular for every $0 \leq t \leq 1$. Thus, the intersection of the supports of any two \mathbf{m} -densities of these restrictions is \mathbf{m} -negligible, and hence the local inequality (3.18) globalizes straightforwardly. \square

Remark 3.28 (Parameter independence of negligible set). By imitating the proof of [39, Cor. 9.5], we get the following *prima facie* stronger, but — under timelike p -essential nonbranching of \mathcal{M} — equivalent version of (iii) in Theorem 3.27. For every timelike p -dualizable pair (μ_0, μ_1) as in (iii) there exists $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ with the following properties. For every $0 \leq t \leq 1$, we have $(\mathbf{e}_t)_\# \pi = \rho_t \mathbf{m} \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$, and there exist \mathbf{m} -versions of the respective \mathbf{m} -densities such that for π -a.e. $\gamma \in \text{TGeo}(\mathbf{M})$, $\rho_t(\gamma_t)$ depends locally Lipschitz continuously on $0 < t < 1$, and (3.18) holds for every $0 \leq t \leq 1$.

Thus, the exceptional set in (3.18) can be chosen to be independent of t .

Remark 3.29 (Variants). The proof of Theorem 3.27 shows that to verify either of its inherent conditions, if \mathcal{M} is timelike p -essentially nonbranching it suffices to check any of the following properties for every $\mu_0, \mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$ with

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}. \quad (3.19)$$

- The existence of $(\mu_t)_{t \in [0,1]}$ and π as in Definition 3.13 certifying (3.13).
- The existence of π as in (iii) in Theorem 3.27.

However, unlike (strong) timelike p -dualizability, cf. Lemma 2.67, the condition (3.19) does not propagate along ℓ_p -geodesics.

3.6. Local-to-global property. Now we prove the local-to-global property of the entropic timelike curvature-dimension condition. Roughly speaking, it states that to verify synthetic timelike Ricci curvature lower bounds, it suffices to check this condition locally (which is clearly true in the smooth case if such bounds are interpreted in terms of the ordinary Bakry–Émery–Ricci tensor).

Recall Definition 3.22 for the condition $\text{TCD}_{p,\text{loc}}^e(k, N)$.

Theorem 3.30 (Globalization). *If \mathcal{M} is timelike p -essentially nonbranching, the properties $\text{TCD}_{p,\text{loc}}^e(k, N)$ and $\text{TCD}_p^e(k, N)$ are equivalent.*

Proof. We only need to prove the forward implication. Let $\mu_0, \mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$ satisfy (3.19), i.e. there exists $r > 0$ such that

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l \geq r\}.$$

By Remark 3.29, it suffices to verify the existence of a plan $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ obeying (iii) in Theorem 3.27 to obtain $\text{TCD}_p^e(k, N)$. By restriction, truncation, and an eventual limit argument, we may and will assume $\mu_0, \mu_1 \in \mathcal{D}(\text{Ent}_{\mathbf{m}})$. Recall from Theorems 3.23 and 3.24 and Corollary 3.25 that under $\text{TCD}_{p,\text{loc}}^e(k, N)$, chronological ℓ_p -optimal couplings and ℓ_p -geodesics [sic] between strongly timelike p -dualizable pairs with endpoints in $\mathcal{D}(\text{Ent}_{\mathbf{m}})$ are indeed unique, and consist of \mathbf{m} -absolutely continuous measures. This is used several times in this proof without further note.

For notational convenience, up to replacing \mathbf{M} by $J(\text{spt } \mu_0, \text{spt } \mu_1)$ we may and will assume compactness of \mathbf{M} (which does not counteract Remark 2.15). Owing to (3.6), Lemma 3.6, and Levi’s monotone convergence theorem, we may and will assume continuity of k on \mathbf{M} . Then there exist $n \in \mathbf{N}$, $0 < \lambda < \text{diam } \mathbf{M}$, mutually disjoint Borel sets $L_1, \dots, L_n \subset \mathbf{M}$ with nonzero \mathbf{m} -measure, and relatively open sets $U_1, \dots, U_n \subset \mathbf{M}$ certifying Definition 3.22 with $\mathbf{B}^d(L_i, \lambda) \subset U_i$ for every $i = 1, \dots, n$; cf. [11, Thm. 5.1]. Finally, let G_r be the set of all $\gamma \in \text{TGeo}(\mathbf{M})$ with $\gamma_0 \in \text{spt } \mu_0$ and $\gamma_1 \in \text{spt } \mu_1$ with $l(\gamma_0, \gamma_1) \geq r$ as defined in Corollary 2.47. The latter states compactness of G_r in $\mathbf{C}([0, 1]; \mathbf{M})$. In particular, Arzelà–Ascoli’s theorem implies uniform equicontinuity of G_r , thus there exists $\delta > 0$ such that if $0 \leq s, t \leq 1$ obey $|s - t| < \delta$, every $\gamma \in G_r$ obeys $d(\gamma_s, \gamma_t) < \lambda$.

Let $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ be the unique timelike ℓ_p -optimal dynamical plan from μ_0 to μ_1 . Let ρ_t denote the \mathbf{m} -density of $(\mathbf{e}_t)_\# \pi$ for every $0 \leq t \leq 1$. Fix

$0 < s < t < \delta$. Given any $0 \leq r \leq 1$, we define

$$\begin{aligned}\boldsymbol{\sigma} &:= (\text{restr}_s^t)_\# \boldsymbol{\pi}, \\ \nu_r &:= (\mathbf{e}_r)_\# \boldsymbol{\sigma} = \varrho_r \mathbf{m},\end{aligned}$$

and recall that $\boldsymbol{\sigma}$ is the only element of $\text{OptTGeo}_p(\nu_0, \nu_1)$. Furthermore, given any $i = 1, \dots, n$ and any $0 \leq r \leq 1$ we define

$$\begin{aligned}\boldsymbol{\sigma}^i &:= (\mathbf{e}_0)_\# \boldsymbol{\sigma} [L_i]^{-1} \boldsymbol{\sigma} \llcorner (\mathbf{e}_0)^{-1} (L_i) \\ \nu_r^i &:= (\mathbf{e}_r)_\# \boldsymbol{\sigma}^i = \varrho_r^i \mathbf{m},\end{aligned}$$

where here and henceforth, we ignore indices i with $(\mathbf{e}_0)_\# \boldsymbol{\sigma} [L_i] = 0$. By restriction, cf. Proposition 2.66, $\boldsymbol{\sigma}^i$ is the only element of $\text{OptTGeo}_p(\nu_0^i, \nu_1^i)$. In particular, the pair (ν_0^i, ν_1^i) is timelike p -dualizable, and $\boldsymbol{\sigma}^i$ is concentrated on l -geodesics that do not leave U_i by our choice of s and t . Consequently, by $\text{TCD}_{p,\text{loc}}^e(k, N)$, a localization argument and relaxation of the condition (3.16) to mere timelike p -dualizability as outlined in the proof of Theorem 3.27 — making Definition 3.22 applicable to (ν_0^i, ν_1^i) by construction —, for every $i = 1, \dots, n$ and every $0 \leq r \leq 1$, $\boldsymbol{\sigma}^i$ -a.e. $\alpha \in \text{TGeo}(\mathbf{M})$ satisfies

$$\varrho_r^i(\alpha_r)^{-1/N} \geq \sigma_{k_\alpha^-/N}^{(1-r)}(|\dot{\alpha}|) \varrho_0^i(\alpha_0)^{-1/N} + \sigma_{k_\alpha^+/N}^{(r)}(|\dot{\alpha}|) \varrho_1^i(\alpha_1)^{-1/N}. \quad (3.20)$$

We point out that the exceptional set depends on i and r .

Now observe that $\boldsymbol{\sigma}^1, \dots, \boldsymbol{\sigma}^n$ are mutually singular. For every $0 < r < 1$, since ν_i is \mathbf{m} -absolutely continuous, ν_r^1, \dots, ν_r^n are thus mutually singular by Lemma 2.71. Since ν_0^1, \dots, ν_0^n are also mutually singular by construction, so are ν_1^1, \dots, ν_1^n since every chronological ℓ_p -optimal coupling is induced by a map by Theorem 3.23. In particular, for every $0 \leq r \leq 1$ and every $i, j = 1, \dots, n$ with $i \neq j$ we have

$$\mathbf{m}[\varrho_r^i > 0, \varrho_r^j > 0] = 0.$$

Hence, by (3.20) for every $0 \leq r \leq 1$, $\boldsymbol{\sigma}$ -a.e. $\alpha \in \text{TGeo}(\mathbf{M})$ satisfies

$$\varrho_r(\alpha_r)^{-1/N} \geq \sigma_{k_\alpha^-/N}^{(1-r)}(|\dot{\alpha}|) \varrho_0(\alpha_0)^{-1/N} + \sigma_{k_\alpha^+/N}^{(r)}(|\dot{\alpha}|) \varrho_1(\alpha_1)^{-1/N}.$$

In other words, for every $0 \leq r \leq 1$, $\boldsymbol{\pi}$ -a.e. $\gamma \in \text{TGeo}(\mathbf{M})$ obeys

$$\begin{aligned}\rho_{\tau(r)}(\gamma_{\tau(r)})^{1/N} &\geq \sigma_{k_{\gamma_{\circ\tau}}^-/N}^{(1-\tau(r))}((t-s)|\dot{\gamma}|) \rho_s(\gamma_s)^{-1/N} \\ &\quad + \sigma_{k_{\gamma_{\circ\tau}}^+/N}^{(\tau(r))}((t-s)|\dot{\gamma}|) \rho_t(\gamma_t)^{-1/N},\end{aligned} \quad (3.21)$$

where $\tau: [0, 1] \rightarrow [0, \delta]$ is defined by $\tau(r) := (1-r)s + rt$. We observe that the exceptional set still depends on r .

From here, we only outline the argument of obtaining (3.18), leaving out technicalities and referring to [24, 39, 80] for details. From (3.21), following the proof of [39, Cor. 9.5], see also Remark 3.28, we get the existence of nonrelabeled \mathbf{m} -versions of $\rho_{\tau(r)}$, where $0 \leq r \leq 1$, such that $\boldsymbol{\pi}$ -a.e. $\gamma \in \text{TGeo}(\mathbf{M})$ satisfies (3.21) for every $0 \leq r \leq 1$, hence the exceptional set in (3.21) can be chosen to be independent of r . Consequently, for $\boldsymbol{\pi}$ -a.e. $\gamma \in \text{TGeo}(\mathbf{M})$ the function $h: [s, t] \rightarrow \mathbf{R}$ defined by $h(u) := \rho_u(\gamma_u)$ is a $\text{CD}(\kappa, N+1)$ density in an evident variable adaptation of [39, Def. A.1], where $\kappa: [s, t] \rightarrow M$ is given by $\kappa(u) := k(\gamma_u)$. By Lemma 3.6, it is a $\text{CD}(K, N+1)$ density for some $K \in \mathbf{R}$ independent of γ , thus twice differentiable at \mathcal{L}^1 -a.e. $u \in [s, t]$. Up to removing another \mathcal{L}^1 -negligible set, by a variable version of [39, Lem. A.3], cf. [80, Cor. 3.13], for \mathcal{L}^1 -a.e. $u \in [s, t]$ we obtain

$$\frac{d^2}{du^2} h(u)^{1/N} \leq -\frac{|\dot{\gamma}|^2}{N} \kappa(u) h(u)^{1/N}.$$

This differential inequality globalizes straightforwardly to \mathcal{L}^1 -a.e. $u \in [0, 1]$ by covering $[0, 1]$ with intervals $[s, t]$, where s and t are as initially chosen above, leading to (3.18) again by [39, Lem. A.3] and [80, Cor. 3.13]. \square

Remark 3.31 (About chronological ℓ_p -geodesy [22, Def. 3.44]). As a byproduct, Theorem 3.30 shows the condition of chronological ℓ_p -geodesy of $\mathcal{P}_c^{\text{ac}}(\mathbb{M}, \mathfrak{m})$ [22, Def. 3.44] in the local-to-global property [22, Thm. 3.45] is in fact redundant.

4. EXAMPLES

Now we discuss examples satisfying our variable timelike curvature-dimension condition, inspired by corresponding results for constant k [22, 23, 26, 91, 99] and the recent observation [92].

In all cases, it is clear that the induced structure constitutes a metric measure spacetime \mathbb{M} satisfying Hypothesis 2.41, cf. e.g. [23, Sec. 2.2] or [87, Thm. 5.16].

4.1. Timelike convergence condition and stability.

4.1.1. *High regularity Lorentzian spacetimes.* As a warm-up, let us first discuss the most regular case which is well-established by now for constant k . Theorem 4.1 below includes the famous strong energy condition of Hawking and Penrose for solutions to the Einstein equations — it is also termed *timelike convergence condition* in the physics literature — synthetic versions of which have been initiated in [44] following [91, 99] and studied further in [21, 92]. However, even if the function k therein is uniformly bounded from below, Theorem 4.1 generally entails a quantitatively better statement.

Let g be a fixed smooth, globally hyperbolic Lorentzian metric with signature $+, -, \dots, -$ on a smooth, connected, noncompact topological manifold \mathbb{M} of dimension $n \in \mathbf{N}_{\geq 2}$. We write $|v|^2 := g(v, v)$ for $v \in T\mathbb{M}$, but stress the usual danger of confusion that the quantity $|v|^2$ may be negative despite the square. Let $V : \mathbb{M} \rightarrow \mathbf{R}$ be twice continuously differentiable. Then, given any $N > 0$ with $N \neq n$, consider the *Bakry–Émery–Ricci tensor*

$$\text{Ric}^{N,V} := \text{Ric} + \text{Hess } V - \frac{1}{N-n} dV \otimes dV$$

associated with the measure $\mathfrak{m} := e^{-V} \text{vol}$. Here vol denotes the Lorentzian volume measure associated with g . By convention, we set $\text{Ric}^{n,V} := \text{Ric}$, i.e. we assume V to be constant in the case $N = n$.

Given any function $k : \mathbb{M} \rightarrow \mathbf{R}$, we will say

$$\text{Ric}^{N,V} \geq k \quad \text{in all timelike directions} \quad (4.1)$$

if $\text{Ric}^{N,V}(v, v) \geq k(x) |v|^2$ for every $x \in \mathbb{M}$ and every $v \in T_x \mathbb{M}$ with $|v|^2 > 0$.

Although the argument for Theorem 4.1 is well-established, we give some more details to familiarize the reader with our variable setting, the way that the distortion coefficients in Definition 3.13 arise from basic ODE theory, etc.

Theorem 4.1 (Equivalence with the timelike convergence condition). *Let $k : \mathbb{M} \rightarrow \mathbf{R}$ be a continuous function and $N > 0$. Then*

$$\text{Ric}^{N,V} \geq k \quad \text{in all timelike directions} \quad (4.2)$$

if and only if $N \geq n$ and the $\text{wTCD}_p^e(k, N)$ condition holds for every $0 < p < 1$.

Furthermore, the preceding equivalence is still valid if $\text{wTCD}_p^e(k, N)$ therein is replaced by $\text{TCD}_p^e(k, N)$.

Proof. Since k is real-valued, the last statement follows from Theorem 3.27.

The forward direction of the first claim follows from straightforward adaptations of the usual arguments, cf. [21, Prop. A.3], [91, Thm. 6.4, Cor. 6.6], [92, Thm. 25], and [99, Thm. 4.3], together with a comparison theorem, cf. e.g. [80, Cor. 3.13]. For the convenience of the reader, we include some details here.

Given any $0 < p < 1$, the Cauchy–Lipschitz theorem for the geodesic equation implies \mathbb{M} to be timelike p -essentially nonbranching. Consequently, by Remark 3.29 it suffices to verify (3.18) for every $\mu_0, \mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbb{M}, \mathfrak{m})$ with

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}.$$

Hence, the pair (μ_0, μ_1) is p -separated according to [91, Def. 4.1] by [91, Lem. 4.4]. By [91, Thm. 5.8, Cor. 5.9] the measures μ_0 and μ_1 admit a unique chronological ℓ_p -optimal coupling $\pi \in \Pi_{\ll}(\mu_0, \mu_1)$ and are joined through a unique ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ consisting of \mathfrak{m} -absolutely continuous measures. Furthermore, the same results from [91] imply the existence of an \mathfrak{m} -a.e. twice approximately differentiable [91, Def. 3.8] Borel function $u: \mathbb{M} \rightarrow \mathbb{M}$ with the subsequent properties. Setting $q := p/(p-1) < 0$, the map $T: [0,1] \times \mathbb{M} \rightarrow \mathbb{M}$ given by

$$T_t(x) := \exp_x(t |\nabla u|^{q-2} \nabla u)$$

satisfies

$$\begin{aligned} \mu_t &= (T_t)_\# \mu_0, \\ \pi &= (\text{Id}, T_1)_\# \mu_0 \end{aligned} \tag{4.3}$$

for every $0 \leq t \leq 1$. Furthermore, for \mathfrak{m} -a.e. $x \in \mathbb{M}$ the approximate derivative $dT_t: T_x \mathbb{M} \rightarrow T_{T_t(x)} \mathbb{M}$ depends smoothly on $0 \leq t \leq 1$, and $(A_t)_{t \in [0,1]}$ is a matrix of Jacobi fields along the l -geodesic $(T_t(x))_{t \in [0,1]}$, where $A_t := dT_t$. Thus, [91, Prop. 6.1] ensures that \mathfrak{m} -a.e., for every $0 \leq t \leq 1$ the quantities

$$\begin{aligned} \phi_t &:= -\log |\det A_t|, \\ B_t &:= A_t' A_t^{-1} \end{aligned}$$

satisfy

$$\begin{aligned} \phi_t' &= -\text{tr } B_t, \\ \phi_t'' &= \text{Ric}(T_t', T_t') + \text{tr } B_t^2, \\ |\text{tr } B_t|^2 &\leq n \text{tr } B_t^2. \end{aligned} \tag{4.4}$$

Now we come to the proof. First, [92, Thm. 25] implies $N \geq n$. For simplicity, we then assume $N > n$, since the case $N = n$ is deduced along the same lines by ignoring V . Given any $0 \leq t \leq 1$, setting

$$\varphi_t := \phi_t + V \circ T_t,$$

by Cauchy's inequality for $\varepsilon := (N-n)/n > 0$, from (4.4) we deduce

$$\frac{1}{N} |\varphi_t'|^2 \leq \frac{1+\varepsilon}{N} |\text{tr } B_t|^2 + \frac{1+\varepsilon^{-1}}{N} dV(T_t')^2 \leq \text{tr } B_t^2 + \frac{1}{N-n} dV(T_t')^2 \quad \mathfrak{m}\text{-a.e.}$$

Setting $\theta := |T_t'| = l(\cdot, T_1)$, again by (4.4) this yields

$$\varphi_t'' + \frac{1}{N} \varphi_t' \leq -\text{Ric}^{N,V}(T_t, T_t') \leq -(k \circ T_t) \theta^2 \quad \mathfrak{m}\text{-a.e.}$$

A standard comparison theorem, cf. e.g. [80, Prop. 3.8], thus yields the transport Jacobian $J_t := e^{\varphi_t}$ [91, Cor. 5.11] to satisfy

$$J_t^{1/N} \geq \sigma_{\kappa^-/N}^{(1-t)}(\theta) J_0^{1/N} + \sigma_{\kappa^+/N}^{(t)}(\theta) J_1^{1/N} \quad \mathfrak{m}\text{-a.e.}$$

Here, analogously to (3.5) the functions $\kappa^\pm: [0, \theta] \rightarrow \mathbf{R}$ are defined by

$$\kappa^-(t\theta) := k \circ T_{1-t},$$

$$\kappa^+(t\theta) := k \circ T_t.$$

By (4.3) and the change of variables formula, this yields (3.18), as desired.

The backward direction is shown along the lines of [91, Thm. 8.5]. We shortly outline the necessary modifications in the variable case. Given any $z \in M$ and any $\varepsilon > 0$, fix $K \in \mathbf{R}$ and a neighborhood $U \subset M$ of z such that $K - \varepsilon \leq k \leq K + \varepsilon$ on U . By Lemma 3.6 and the local argument for [91, Thm. 8.5], we get

$$\text{Ric}^{N,V}(v, v) \geq (K - \varepsilon) |v|^2 \geq (k(x) - 2\varepsilon) |v|^2$$

for every $x \in U$ and every $v \in T_x M$ with $|v|^2 > 0$. Covering M with such sets U and sending $\varepsilon \rightarrow 0$, the claim follows. \square

4.1.2. Low regularity spacetimes. Following [23], we can reduce the regularity of g and V from the previous subsection to local $C^{1,1}$ -regularity, which we assume in this section — otherwise, we retain the notation and the hypotheses therein.

The Bakry–Émery–Ricci tensor $\text{Ric}^{N,V}$ is still well-defined \mathfrak{m} -a.e., where again $\mathfrak{m} := e^{-V}$ vol. Accordingly, given any $N > 0$ we will say

$$\text{Ric}^{N,V} \geq k \quad \text{in all timelike directions}$$

if for every smooth timelike vector field X on M , $\text{Ric}^{N,V}(X, X) \geq k(x) |X|^2$ holds for \mathfrak{m} -a.e. $x \in M$. Following the sophisticated approximation argument from [23, Lem. 2.8, Ch. 3] then leads to the following.

Theorem 4.2 (Timelike convergence condition in low regularity). *Let $k: M \rightarrow \mathbf{R}$ be a continuous function, assume g is of regularity $C^{1,1}$, and suppose $V \in C^{1,1}(M)$. Then for every $N \geq n$ the property*

$$\text{Ric}^{N,V} \geq k \quad \text{in all timelike directions}$$

implies the $\text{TCD}_p^e(k, N)$ condition for every $0 < p < 1$.

Remark 4.3 (Converse implication). Even for constant k , it is still open whether the converse statement from Theorem 4.2 holds. In the positive signature case, in this precise regularity this is only known under additional hypotheses about the stability of optimal maps thus far [86, Thm. 6.3, Rem. 6.4].

Remark 4.4 (Stability). As Theorem 4.2 is shown through approximation of g by smooth metrics, it can be seen as an instance of stability of our variable timelike curvature-dimension condition.

We do not address the general question of stability, as a clear candidate for measured Lorentz–Gromov–Hausdorff convergence (e.g. allowing for a precompactness theorem) is missing as of today. One proposal comes from [44, Thm. 3.15], and further possibilities will be studied in [27]. By using Lemma 3.11, it is not hard to adapt the proof of [44, Thm. 3.15], see also [22, Thm. 3.29], to show Theorem 4.5.

Theorem 4.5 (Weak stability à la Cavalletti–Mondino). *Let a sequence $(\mathcal{M}_n)_{n \in \mathbf{N}}$ converge to \mathcal{M}_∞ in the pointed sense of [44, Thm. 3.15], all spaces satisfying Hypothesis 2.41. Let $(k_n)_{n \in \mathbf{N}}$ and $(N_n)_{n \in \mathbf{N}}$ be sequences of lower semicontinuous functions $k_n: \mathcal{M}_n \rightarrow \mathbf{R}$ and numbers $1 \leq N_n < \infty$ such that \mathcal{M}_n satisfies $\text{TCD}_p^e(k_n, N_n)$, respectively. Lastly, let $k_\infty: \mathcal{M}_\infty \rightarrow \mathbf{R}$ be lower semicontinuous and $1 \leq N_\infty < \infty$ be such that*

$$\begin{aligned} \liminf_{n \rightarrow \infty} k_n &\geq k_\infty, \\ \limsup_{n \rightarrow \infty} N_n &\leq N_\infty. \end{aligned} \tag{4.5}$$

Then \mathcal{M}_∞ satisfies $\text{wTCD}_p^e(k_\infty, N_\infty)$.

Here the first inequality in (4.5) means for every $\varepsilon > 0$ there exists $n' \in \mathbf{N}$ such that for every $n \in \mathbf{N}$ with $n \geq n'$ and every $x \in M_n$,

$$k_n(x) \geq k \circ \iota_n(x) - \varepsilon,$$

where ι_n is the isometric embedding defined on M_n from [44, Thm. 3.15].

If \mathcal{M}_n is merely a timelike p -essentially nonbranching $\text{wTCD}_p^e(k_n, N_n)$ space for every $n \in \mathbf{N}$, by the previous result and Theorem 3.27 we still obtain \mathcal{M}_∞ obeys the $\text{wTCD}_p^e(k_\infty, N_\infty)$ condition.

A similar statement entails weak stability of the dimensionless variable timelike curvature-dimension condition from Definition 3.12.

If the local uniform lower boundedness in (4.5) is dropped, stability cannot be expected in general [92, Ex. 33].

4.1.3. Finsler spacetimes. Next, we adapt the results from [26] for constant k on Finsler spacetimes to variable k . We refer to [26, 96] for details and references.

Let M be a smooth, connected, noncompact topological manifold of dimension $n \in \mathbf{N}_{\geq 2}$. A *Lorentz–Finsler structure* on M [26, Def. 2.1] is a positively 2-homogeneous function $L: TM \rightarrow \mathbf{R}$, which outside the zero section is smooth and has non-degenerate Hessian with signature $+, -, \dots, -$. Analogously to the metric tensor g in Subsection 4.1.1, one can use L to build all elements of causality theory to define *Finsler spacetimes*, i.e. Lorentz–Finsler structures with a time orientation [26, Def. 2.3]. Let \mathfrak{m} be a *smooth measure* on M , i.e. \mathfrak{m} is mutually absolutely continuous to the Lebesgue measure \mathcal{L}^n on each chart with smooth density.

Uniqueness results for chronological ℓ_p -optimal couplings and ℓ_p -geodesics still hold on Finsler spacetimes without curvature assumptions [26, Thm. 4.17, Cor. 4.18], in analogy to the results from [91] mentioned in Subsection 4.1.1.

As Finsler spacetimes come with several candidates for a “volume measure” in general [104, Sec. 9.1], there is also no canonical Ricci tensor to consider. Given \mathfrak{m} , the correct choice is the weighted Ricci tensor Ric^N it induces, as explained in [26, Sec. 5.1]. We then say

$$\text{Ric}^N \geq k \quad \text{in all timelike directions}$$

if $\text{Ric}^N(v, v) \geq 2k(x)L(v)$ for every $x \in M$ and every $v \in T_x M$ with $L(v) > 0$.

The following theorem echoes the analogous characterization for constant k [26, Thm. 5.9, Thm. 6.1]. The proof is performed along the same lines.

Theorem 4.6 (Finslerian timelike convergence condition). *Let $k: M \rightarrow \mathbf{R}$ be a continuous function. Then*

$$\text{Ric}^N \geq k \quad \text{in all timelike directions}$$

if and only if $N \geq n$ and the $\text{wTCD}_p^e(k, N)$ condition holds for every $0 < p < 1$.

This equivalence still holds when $\text{wTCD}_p^e(k, N)$ is replaced by $\text{TCD}_p^e(k, N)$.

4.2. Regular solutions to the Einstein equations. Akin to the setting from Subsection 4.1.1, let g be a smooth, globally hyperbolic Lorentzian metric which solves the *Einstein equations*

$$\text{Ric} - \frac{1}{2} \text{scal } g + \Lambda g = 8\pi T. \quad (4.6)$$

Here $\Lambda \in \mathbf{R}$ is the cosmological constant, and T is the energy-momentum tensor from the least action principle for the Einstein–Hilbert action. In order to simplify the presentation, cf. (4.7) below, we assume M has dimension $n \in \mathbf{N}_{\geq 3}$.

In the interpretation of (4.1), the *weak energy condition* asserts

$$T \geq 0 \quad \text{in all timelike directions.}$$

It clearly holds in vacuum, where $T = 0$, in which case (4.6) is equivalent to

$$\text{Ric} = \frac{2\Lambda}{n-2}g \quad (4.7)$$

by [99, Lem. 4.1]. We thus land in the framework of timelike curvature-dimension conditions for constant k again [21, 44]. In the general non-vacuum case, assuming (4.6) the weak energy condition — which is believed under all physically reasonable non-quantum circumstances [130] — is equivalent to

$$\text{Ric} \geq \frac{1}{2} \text{scal } g - \Lambda g \quad \text{in all timelike directions.}$$

This motivates the consideration of the continuous function $k_g: \mathbf{M} \rightarrow \mathbf{R}$ with

$$k_g(x) := \frac{1}{2} \text{scal } g(x) - \Lambda.$$

The proof of the following characterization is analogous to Theorem 4.1.

Theorem 4.7 (Equivalence with the weak energy condition). *Assume \mathcal{M} smooth satisfies general relativity, i.e. it comes from a solution to the Einstein equations (4.6). Then the weak energy condition is equivalent to $\text{TCD}_p^e(k_g, n)$ for every $0 < p < 1$.*

This equivalence is still valid if $\text{TCD}_p^e(k_g, n)$ is replaced by $\text{wTCD}_p^e(k_g, n)$.

4.3. Null convergence condition. Unlike the positive signature case — by the noncompactness of $\{|\cdot|^2 = 1\}$ in $T_x\mathbf{M}$, where $x \in \mathbf{M}$ — existence of a *real-valued* function k obeying (4.2) is an extra condition on \mathbf{M} and not automatically satisfied by every smooth Lorentzian spacetime. In fact, the existence of such a function k is characterized by the null energy condition (or *null convergence condition* if general relativity is not assumed). This has been observed in [92]; we briefly recapitulate its findings and translate the main result from [92] to our variable language.

Define the function $c: \mathbf{M} \rightarrow \mathbf{R} \cup \{-\infty\}$ by

$$c(x) := \inf |v|^{-2} \text{Ric}^{N,V}(v, v), \quad (4.8)$$

where the infimum is taken over all $v \in T_x\mathbf{M}$ such that $|v|^2 > 0$. This function is only upper semicontinuous in general. Therefore, we have to consider its (possibly degenerate) lower semicontinuous envelope $k: \mathbf{M} \rightarrow \mathbf{R} \cup \{-\infty\}$ defined by

$$k(x) := \inf \liminf_{n \rightarrow \infty} c(x_n), \quad (4.9)$$

where the infimum is taken over all sequences $(x_n)_{n \in \mathbf{N}}$ in \mathbf{M} converging to x .

In the sequel, we will say

$$\text{Ric}^{N,V} \geq 0 \quad \text{in all null directions}$$

if $\text{Ric}^{N,V}(v, v) \geq 0$ for every $x \in \mathbf{M}$ and every $v \in T_x\mathbf{M}$ with $|v|^2 = 0$.

Theorem 4.8 (Equivalence with the null convergence condition). *The condition*

$$\text{Ric}^{N,V} \geq 0 \quad \text{in all null directions}$$

holds if and only if the following are satisfied.

- (i) *The function k from (4.9) does not attain the value $-\infty$.*
- (ii) *We have $N \geq n$.*
- (iii) *For every $0 < p < 1$, the $\text{wTCD}_p^e(k, N)$ condition holds.*

Moreover, the preceding equivalence still holds if $\text{wTCD}_p^e(k, N)$ is replaced by $\text{TCD}_p^e(k, N)$ in item (iii).

Proof. Again the last statement follows from Theorem 3.27 since in either case, (i) ensures k is real-valued.

As concerns the rest, we first show the forward implication. The hypothesized weighted null energy condition implies the function c from (4.8) is bounded from below on every compact subset of M [92, Thm. 26]. Thus, k does not attain the value $-\infty$. The rest is then argued as in the proof of Theorem 4.1.

The key point of the converse implication is again (i). Indeed, this hypothesis paired with lower semicontinuity of k ensures k is bounded from below on every compact subset of M . Thus for every compact subset $C \subset M$, there exists $K \in \mathbf{R}$ such that the metric measure spacetime structure on $J(C, C)$ induced by \mathcal{M} obeys $\text{wTCD}_p^e(K, N)$ for every $0 < p < 1$. According to [92, Cor. 27] this yields the weighted null energy condition. \square

5. GEOMETRIC INEQUALITIES

We turn to standard geometric applications of timelike Ricci curvature bounds from below. In this chapter, from $\text{wTCD}_p^e(k, N)$ we deduce

- the timelike Brunn–Minkowski inequality,
- the timelike Bonnet–Myers diameter inequality, and
- the timelike Bishop–Gromov volume growth inequality, along with several properties of \mathfrak{m} .

The constants obtained in this chapter are not sharp (e.g. Corollary 5.11), but this will be rectified in Chapter 7 using the needle decompositions of Chapter 6.

In Theorem 5.6 we also state a related result known as *Schneider’s theorem* in Riemannian geometry [114], see also [10, 62], considered in the setting of smooth spacetimes by Frankel–Galloway [58]. It is transferred from [80] to our setting. In terms of general relativity, it predicts timelike geodesic incompleteness if, in the far past (or future) of the point of reference, the Ricci curvature is positive and does not decay too quickly.

5.1. Timelike Brunn–Minkowski inequality. For any Borel sets $A_0, A_1 \subset M$ and any $t \in (0, 1)$, we define the set

$$A'_t := Z_t(A_0 \times A_1), \quad (5.1)$$

where the latter is defined according to (2.8). Furthermore, let $G(A_0, A_1)$ denote the class of all $\gamma \in \text{TGeo}(M)$ with $\gamma_0 \in A_0$ and $\gamma_1 \in A_1$.

Theorem 5.1 (Timelike Brunn–Minkowski inequality). *Assume $\text{TCD}_p^e(k, N)$. Let $A_0, A_1 \subset M$ be Borel subsets of M with positive and finite \mathfrak{m} -measure such that $(\mu_0, \mu_1) \in \mathcal{P}^{\text{ac}}(M, \mathfrak{m})^2$ is timelike p -dualizable, where*

$$\begin{aligned} \mu_0 &:= \mathfrak{m}[A_0]^{-1} \mathfrak{m} \llcorner A_0, \\ \mu_1 &:= \mathfrak{m}[A_1]^{-1} \mathfrak{m} \llcorner A_1. \end{aligned} \quad (5.2)$$

Then for every $0 \leq t \leq 1$,

$$\mathfrak{m}[A'_t]^{1/N} \geq \inf \sigma_{k_\gamma^-/N}^{(1-t)}(|\dot{\gamma}|) \mathfrak{m}[A_0]^{1/N} + \inf \sigma_{k_\gamma^+/N}^{(t)}(|\dot{\gamma}|) \mathfrak{m}[A_1]^{1/N},$$

where both infima are taken over all $\gamma \in G(A_0, A_1)$. Here, $\mathfrak{m}[A'_t]$ is interpreted as the outer measure of A'_t in case the latter is not \mathfrak{m} -measurable.

Proof. By inner regularity of \mathfrak{m} , it suffices to prove the claim for compact A_0 and A_1 . In this case, A'_t is compact by Lemma 2.49, and thus \mathfrak{m} -measurable.

By $\text{TCD}_p^e(k, N)$ the pair (μ_0, μ_1) are connected by an ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ from μ_0 to μ_1 such that for some $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$, we have

$$\mathcal{U}_N(\mu_t) \geq \sigma_{k_\pi^-/N}^{(1-t)}(\mathbf{c}_\pi) \mathcal{U}_N(\mu_0) + \sigma_{k_\pi^+/N}^{(t)}(\mathbf{c}_\pi) \mathcal{U}_N(\mu_1).$$

To estimate the left-hand side from below, we use the identities

$$\begin{aligned}\mathcal{U}_N(\mu_0) &= \mathfrak{m}[A_0]^{1/N}, \\ \mathcal{U}_N(\mu_1) &= \mathfrak{m}[A_1]^{1/N}\end{aligned}\tag{5.3}$$

and that, by Lemma 3.6 and Jensen's inequality,

$$\sigma_{k_{\bar{\gamma}}^-/N}^{(1-t)}(\mathbf{c}\pi) \geq \int \sigma_{k_{\bar{\gamma}}^-/N}^{(1-t)}(|\dot{\gamma}|) d\pi(\gamma) \geq \inf \sigma_{k_{\bar{\gamma}}^-/N}^{(1-t)}(|\dot{\gamma}|).$$

where the infimum is taken over all $\gamma \in G(A_0, A_1)$; analogously,

$$\sigma_{k_{\bar{\pi}}^+/N}^{(t)}(\mathbf{c}\pi) \geq \inf \sigma_{k_{\bar{\pi}}^+/N}^{(t)}(|\dot{\gamma}|)$$

with the infimum again running over all $\gamma \in \text{spt } \pi$. Here we have implicitly used (3.3) twice. On the other hand, from (2.29) and Lemma 2.63 we obtain

$$\mathcal{U}_N(\mu_t) \leq \mathfrak{m}[\text{spt } \mu_t]^{1/N} \leq \mathfrak{m}[A'_t]^{1/N}. \quad \square$$

The same strategy, together with the lifting of chronological ℓ_p -optimal couplings from Proposition 2.66, lead to a Brunn–Minkowski-type inequality in the infinite-dimensional case. Given any $0 \leq t \leq 1$ we define $k_t: M^2 \rightarrow \mathbf{R}$ by

$$k_t(x, y) := \inf \int_0^1 g(s, t) k(\gamma_s) ds,$$

where the infimum is taken over all $\gamma \in \text{TGeo}(M)$ with $\gamma_0 = x$ and $\gamma_1 = y$, and g is from (3.7). This function should be compared to the function \underline{k} from [25], which has been extensively used to formulate and investigate transport estimates for the heat flow on metric measure spaces with variable curvature bounds.

Corollary 5.2 (Dimensionless timelike Brunn–Minkowski inequality). *We assume $\text{TCD}_p(k, \infty)$ and retain the notation from Theorem 5.1. Then for every $0 \leq t \leq 1$,*

$$\log \mathfrak{m}[A'_t] \geq (1-t) \log \mathfrak{m}[A_0] + t \log \mathfrak{m}[A_1] - \inf_{M^2} \int k_t l^2 d\pi,$$

where the infimum is taken over all ℓ_p -optimal couplings $\pi \in \Pi_{\ll}(\mu_0, \mu_1)$. Again, $\mathfrak{m}[A'_t]$ is interpreted as the outer measure of A'_t if the latter is not \mathfrak{m} -measurable.

Remark 5.3 (Variants). Of course, Theorem 5.1 and Corollary 5.2 hold analogously under $\text{wTCD}_p^e(k, N)$ and $\text{wTCD}_p(k, \infty)$. In both situations, however, one has to assume *strong* timelike p -dualizability of the pair defined in (5.2).

In the case that \mathcal{M} is a timelike p -essentially nonbranching $\text{wTCD}_p^e(k, N)$ space, better inequalities can be achieved if (μ_0, μ_1) is strongly timelike p -dualizable. Indeed, by Theorem 3.24 and Corollary 3.25 the unique ℓ_p -geodesic connecting μ_0 to μ_1 is in fact a displacement ℓ_p -geodesic, represented by $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$. For every $0 \leq t \leq 1$, arguing as for Theorem 5.1 then yields

$$\mathfrak{m}[\text{spt}(\mathbf{e}_t)_\# \pi]^{1/N} \geq \inf \sigma_{k_{\bar{\gamma}}^-/N}^{(1-t)}(|\dot{\gamma}|) \mathfrak{m}[A_0]^{1/N} + \inf \sigma_{k_{\bar{\pi}}^+/N}^{(t)}(|\dot{\gamma}|) \mathfrak{m}[A_1]^{1/N},$$

where both infima are taken over all $\gamma \in \text{spt } \pi$. Thus, considering the Suslin set

$$A_t := \mathbf{e}_t[G(A_0, A_1)], \tag{5.4}$$

which is no larger than A'_t by Lemma 2.63, we have

$$\mathfrak{m}[A_t]^{1/N} \geq \inf \sigma_{k_{\bar{\gamma}}^-/N}^{(1-t)}(|\dot{\gamma}|) \mathfrak{m}[A_0]^{1/N} + \inf \sigma_{k_{\bar{\pi}}^+/N}^{(t)}(|\dot{\gamma}|) \mathfrak{m}[A_1]^{1/N},$$

where both infima are taken over all $\gamma \in G(A_0, A_1)$.

If $A_0 \times A_1 \Subset \{l > 0\}$, the above assumption of timelike p -essential nonbranching can be dropped by Proposition 2.66.

5.2. Timelike Bonnet–Myers inequality. In this discussion, we will denote by $\text{OptTGeo}_p^{\text{ac}}(\mathcal{P}_c(\mathbf{M})^2, \mathbf{m})$ the set of all $\pi \in \text{OptTGeo}_p(\mathcal{P}_c(\mathbf{M})^2)$ such that $(\mathbf{e}_t)_\# \pi$ is \mathbf{m} -absolutely continuous for every $0 \leq t \leq 1$. We also recall the definition

$$\mathbf{c}_\pi := \|l \circ (\mathbf{e}_0, \mathbf{e}_1)\|_{L^2(\text{TGeo}(\mathbf{M}), \pi)}$$

of the L^2 -cost associated with π .

Definition 5.4 (Effective diameter). *The effective l -diameter of \mathbf{M} with respect to k and N is defined by*

$$\text{diam}_{k/N} \mathbf{M} := \sup \mathbf{c}_\pi,$$

where the supremum is taken over all $\pi \in \text{OptTGeo}_p^{\text{ac}}(\mathcal{P}_c(\mathbf{M})^2, \mathbf{m})$ such that for some, or equivalently by Lemma 3.6, every $0 \leq t \leq 1$, we have $\sigma_{k_\pi^\pm/N}^{(t)}(\mathbf{c}_\pi) < \infty$.

By definition of the effective l -diameter, we have

$$\text{diam}_{k/N} \mathbf{M} \leq \text{diam}^l \text{spt } \mathbf{m}, \quad (5.5)$$

where the latter is the largest value l can take on $(\text{spt } \mathbf{m})^2$.

Theorem 5.5 (Timelike Bonnet–Myers inequality). *If \mathcal{M} satisfies $\text{wTCD}_p^e(k, N)$, then equality holds throughout (5.5), i.e.*

$$\text{diam}_{k/N} \mathbf{M} = \text{diam}^l \text{spt } \mathbf{m}.$$

Proof. Assume the existence of $\varepsilon > 0$ and $x, y \in \text{spt } \mathbf{m}$ such that

$$\text{diam}_{k/N} \mathbf{M} + 2\varepsilon \leq l(x, y),$$

whence $x \ll y$. By l -geodesy of \mathbf{M} and Corollary 2.17, there are precompact, open sets $A_0, A_1 \subset \mathbf{M}$ with

$$\text{diam}_{k/N} \mathbf{M} + \varepsilon \leq \inf \{l(x', y') : x' \in A_0, y' \in A_1\}. \quad (5.6)$$

Finally, in terms of A_0 and A_1 we define $\mu_0, \mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$ by (5.2). Then the pair (μ_0, μ_1) is strongly timelike p -dualizable by (5.6) and Example 2.59.

Every $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ and every $0 < t < 1$ obeys

$$\sigma_{k_\pi^-/N}^{(1-t)}(\mathbf{c}_\pi) = \sigma_{k_\pi^+/N}^{(t)}(\mathbf{c}_\pi) = \infty \quad (5.7)$$

by Lemma 3.6 and our assumption. On the other hand, our curvature assumption provides us with an ℓ_p -geodesic $(\mu_t)_{t \in [0,1]}$ and $\pi \in \text{OptTGeo}_p^{\text{ac}}(\mu_0, \mu_1, \mathbf{m})$ certifying the displacement semiconvexity inequality of the functional \mathcal{U}_N . Let A'_t be the set from (5.1) in terms of A_0 and A_1 . Combined with Jensen's inequality, (5.7), and the identities (5.3) for $\mathcal{U}_N(\mu_0)$ and $\mathcal{U}_N(\mu_1)$, our discussion yields

$$\mathbf{m}[J(\text{spt } \mu_0, \text{spt } \mu_1)]^{1/N} \geq \mathbf{m}[A'_t]^{1/N} \geq \mathcal{U}_N(\mu_t) \geq \infty.$$

Since \mathbf{m} is finite on compact sets, this is however impossible. \square

5.3. Timelike Schneider inequality.

Theorem 5.6 (Past timelike Schneider inequality). *We assume $\text{wTCD}_p^e(k, N)$ and suppose that some $o \in \text{spt } \mathbf{m}$, $\beta > 0$, and $R > 0$ satisfy*

$$k \geq N \left[\frac{1}{4} + \beta^2 \right] l(\cdot, o)^{-2} \quad \text{on } \{l(\cdot, o) > R\}.$$

Then $\text{diam}^l(I^-(o) \cap \text{spt } \mathbf{m}) \leq R e^{\pi/\beta}$.

Proof. The proof is comparable to the argument for Theorem 5.5. For notational convenience, we set $x_1 := o$. Assume to the contrapositive that some $x_0 \in \text{spt } \mathbf{m}$ and some $0 < \varepsilon < R$ satisfy the inequality

$$l(x_0, x_1) > (R + 4\varepsilon e^{\pi/\beta}) e^{\pi/\beta}. \quad (5.8)$$

As in the proof of Theorem 5.5, we choose precompact, open neighborhoods $A_0 \subset \mathbf{M}$ of x_0 and $A_1 \subset \mathbf{M}$ of x_1 such that for every $x \in \bar{A}_0$ and every $y \in \bar{A}_1$,

$$|l(x, y) - l(x_0, x_1)| \leq \varepsilon e^{\pi/\beta}. \quad (5.9)$$

In particular, for every timelike maximizer $\bar{\gamma}: [0, l(x, x_1)] \rightarrow \mathbf{M}$ with proper-time parametrization connecting x to x_1 we have, whenever $0 < r < \varepsilon$,

$$\begin{aligned} l(\bar{\gamma}_r, x_1) &= l(x, x_1) - r \\ &\geq l(x_0, x_1) - \varepsilon e^{\pi/\beta} - r \\ &> (R + 4\varepsilon e^{\pi/\beta}) e^{\pi/\beta} - \varepsilon e^{\pi/\beta} - r \\ &> R + \varepsilon - r \\ &> R. \end{aligned}$$

For such r , our hypothesis on k together with (5.9) thus yields

$$\begin{aligned} k(\bar{\gamma}_r) &\geq N \left[\frac{1}{4} + \beta^2 \right] l(\bar{\gamma}_r, x_1)^{-2} \\ &= N \left[\frac{1}{4} + \beta^2 \right] [l(x, x_1) - r]^{-2} \\ &\geq N \left[\frac{1}{4} + \beta^2 \right] [\varepsilon e^{\pi/\beta} + l(x_0, x_1) - r]^{-2}. \end{aligned}$$

Next, define $\kappa: [0, \varepsilon] \rightarrow \mathbf{R}$ through

$$\kappa(r) := \left[\frac{1}{4} + \beta^2 \right] [\eta + l(x_0, x_1) - r]^{-2},$$

where $\eta := \varepsilon e^{\pi/\beta}$. In the notation of (3.5), the preceding paragraph provides the bound $k_\gamma^+/N \geq \kappa$ on $[0, \varepsilon]$, which directly transfers to $\sin_{k_\gamma^+/N} \geq \sin_\kappa$ on the same interval by Remark 3.5. We can explicitly express \sin_κ as

$$\sin_\kappa(r) = c \sqrt{\eta + l(x_0, x_1) - r} \sin \left[\beta \log \left[\frac{\eta + l(x_0, x_1) - r}{\eta + l(x_0, x_1)} \right] + \pi \right]$$

for a normalization constant $c > 0$. Of course, the first zero of \sin_κ is $r_1 = 0$, and its second zero is $r_2 = (\eta + l(x_0, x_1)) (1 - e^{-\pi/\beta})$. Consequently, the inequalities $\varepsilon < R$, (5.8), and (5.9) yield, for every $x \in \bar{A}_0$,

$$\begin{aligned} r_2 &= \varepsilon e^{\pi/\beta} (1 - e^{-\pi/\beta}) + l(x_0, x_1) - l(x_0, x_1) e^{-\pi/\beta} \\ &< \varepsilon e^{\pi/\beta} (1 - e^{-\pi/\beta}) + l(x, y) + \varepsilon e^{\pi/\beta} - R - 4\varepsilon e^{\pi/\beta} \\ &= l(x, y) - R - \varepsilon - 2\varepsilon e^{\pi/\beta} \\ &< l(x, y) - 2\varepsilon. \end{aligned} \quad (5.10)$$

To wrap up the argument and to arrive at the desired contradiction, we falsify a weaker version of the timelike Brunn–Minkowski inequality from Theorem 5.1. Consider the set of l -geodesics $G(A_0, A_1)$ connecting A_0 and A_1 as well as the associated t -intermediate point set A'_t , where $0 \leq t \leq 1$, from (5.1). Moreover, in terms of A_0 and A_1 we define $\mu_0, \mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$ as in (5.2), respectively. By construction and Example 2.59, the pair (μ_0, μ_1) is strongly timelike p -dualizable. Furthermore, if $0 < r < \varepsilon$, by (5.10) every $\gamma \in G(A_0, A_1)$ saturates

$$|\dot{\gamma}| - r = l(\gamma_0, \gamma_1) - r > r_2 + \varepsilon.$$

In other words, for sufficiently small $r > 0$, the argument $|\dot{\gamma}| - r$ lies strictly past the first zero of κ for every $\gamma \in G(A_0, A_1)$. The nondegeneracy and monotonicity properties from Lemma 3.6 and the definition of the distortion coefficients under consideration imply for sufficiently small $t > 0$ that

$$\mathfrak{m}[A'_t]^{1/N} \geq \inf \sigma_{\frac{1-t}{k\bar{\gamma}/N}}^{(1-t)}(|\dot{\gamma}|) \mathfrak{m}[A_0]^{1/N} = \infty,$$

where the infimum is taken over all $\gamma \in G(A_0, A_1)$. As in the proof of Theorem 5.5, this contradicts the Radon property of \mathfrak{m} . \square

By Remark 3.15, Theorem 5.6 readily yields the following.

Corollary 5.7 (Future timelike Schneider inequality). *Assume $\text{wTCD}_p^e(k, N)$, and suppose that some $o \in \text{spt } \mathfrak{m}$, $\beta > 0$, and $R > 0$ satisfy*

$$k \geq N \left[\frac{1}{4} + \beta^2 \right] l(o, \cdot)^{-2} \quad \text{on } \{l(o, \cdot) > R\}.$$

Then $\text{diam}^l(I^+(o) \cap \text{spt } \mathfrak{m}) \leq R e^{\pi/\beta}$.

5.4. Timelike Bishop–Gromov inequality. Following [44, Sec. 3.1], let

$$\mathbf{B}^l(x, r) := \{y \in I^+(x) : l(x, y) < r\} \cup \{x\}$$

be the future “ball” anchored at $x \in \mathbf{M}$ with radius $r > 0$. Such a set typically has infinite volume, e.g. it is a hyperboloid in flat Minkowski space. To take the \mathfrak{m} -measure of parts of such balls, one localizes them by star shaped sets. Here a set $E \subset I^+(x) \cup \{x\}$ is called *l-star shaped* with respect to x if every $\gamma \in \text{TGeo}(\mathbf{M})$ with $\gamma_0 = x$ and $\gamma_1 \in E$ does not leave E . Let E be such a set which we assume to be compact in addition. Using the Radon property of \mathfrak{m} , we set

$$\begin{aligned} \mathfrak{v}(r) &:= \mathfrak{m}[\bar{\mathbf{B}}^l(x, r) \cap E], \\ \mathfrak{s}(r) &:= \limsup_{\delta \rightarrow 0} \delta^{-1} \mathfrak{m}[(\bar{\mathbf{B}}^l(x, r + \delta) \setminus \mathbf{B}^l(x, r)) \cap E] \end{aligned} \quad (5.11)$$

where $\bar{\mathbf{B}}^l(x, r)$ denotes the closure of $\mathbf{B}^l(x, r)$.

The following construction of compact and *l*-star shaped sets will be useful in the proof of Theorem 5.9. Here, analogously to (5.1), $G(x, A)$ denotes the set of all $\gamma \in \text{TGeo}(\mathbf{M})$ with $\gamma_0 = x$ and $\gamma_1 \in A$, where $x \in \mathbf{M}$ and $A \subset \mathbf{M}$.

Lemma 5.8 (Filled future cone). *Given any $x \in \mathbf{M}$ and any compact $C \subset \mathbf{M}$ with $\{x\} \times C \subset \{l > 0\}$, the set*

$$E := \bigcup_{t \in [0, 1]} \mathfrak{e}_t[G(x, C)]$$

*is contained in $I^+(x) \cup \{x\}$, compact, and *l*-star shaped.*

Proof. The inclusion $E \subset I^+(x) \cup \{x\}$ is clear by Corollary 2.14.

Now we show compactness of E . Since $E \subset J(x, C)$, by Theorem 2.18 it suffices to prove its sequential compactness. Therefore, let $(z_n)_{n \in \mathbf{N}}$ be a sequence in E and write $z_n := \gamma_{t_n}^n$ for some sequences $(\gamma^n)_{n \in \mathbf{N}}$ in $G(x, C)$ and $(t_n)_{n \in \mathbf{N}}$ in $[0, 1]$, respectively. Since $\inf l(\{x\} \times C) > 0$, Corollary 2.47 entails the existence of a uniform limit $\gamma \in G(x, C)$ of a nonrelabeled subsequence of $(\gamma^n)_{n \in \mathbf{N}}$. By possibly passing to a further nonrelabeled subsequence, the sequence $(t_n)_{n \in \mathbf{N}}$ converges to some $0 \leq t \leq 1$. In particular, $(z_n)_{n \in \mathbf{N}}$ converges to $\gamma_t \in E$.

We are left with *l*-star shapedness of E . Let $\gamma \in \text{TGeo}(\mathbf{M})$ with $\gamma_0 = x$ and $\gamma_1 \in E$. If $\gamma_1 \in C$ we directly get $\gamma_{[0, 1]} \subset E$ by definition of E . Otherwise, there exist $\sigma \in G(x, C)$ and $0 < t < 1$ with $\gamma_1 = \sigma_t$. It follows that

$$l(\gamma_0, \gamma_1) = l(x, \sigma_t) = t l(x, \sigma_1).$$

On the other hand, any $\alpha \in \text{TGeo}(\mathbf{M})$ with $\alpha_0 = \sigma_t$ and $\alpha_1 = \sigma_1$ has length

$$l(\alpha_0, \alpha_1) = l(\sigma_t, \sigma_1) = (1-t)l(x, \sigma_1).$$

As in the proof of Lemma 2.46, an appropriate concatenation of γ and α yields an element of $\text{TGeo}(\mathbf{M})$ starting at x , terminating in C , and containing $\gamma_{[0,1]}$ as a subsegment. By the definition of E , we thus get $\gamma_{[0,1]} \subset E$. \square

For brevity, a *future sphere* is any set of the form

$$S^l(x, r) := \{l(x, \cdot) = r\} = (\bar{\mathbf{B}}^l(x, r) \setminus \mathbf{B}^l(x, r)) \setminus \{l(x, \cdot) = 0\},$$

where $x \in \mathbf{M}$ and $r > 0$. *Past spheres* are defined accordingly with respect to the causal reversal from Example 2.45. If the time orientation does not matter, we simply use the term *sphere*.

Next, we foliate the function k along spheres throughout E as follows. Since the set E is compact, so is $E_r := S^l(x, r) \cap E$ for every $r > 0$ (and contained in $\{l(x, \cdot) = r\}$), and also the intersection of $S^l(x, 0) := \{l(x, \cdot) = 0\}$ with E . Hence k assumes its minimum, say $k_x(r)$, on E_r if the latter is nonempty. Let

$$R_x := \sup\{r \geq 0 : S^l(x, r) \cap E \neq \emptyset\},$$

and let $\underline{k}_x : [0, R_x] \rightarrow \mathbf{R}$ denote the lower semicontinuous envelope of k_x , i.e.

$$\underline{k}_x(r) := \inf \liminf_{n \rightarrow \infty} k_x(r_n), \quad (5.12)$$

where the infimum is taken over all sequences $(r_n)_{n \in \mathbf{N}}$ in $[0, R_x]$ converging to r , where $[0, R_x] := [0, \infty)$ provided $R_x = \infty$. By a slight abuse of notation, we will also consider the function $\underline{k}_x : J^+(x) \rightarrow \mathbf{R}$ defined by

$$\underline{k}_x(y) := \underline{k}_x(l(x, y)). \quad (5.13)$$

Theorem 5.9 (Timelike Bishop–Gromov inequality). *Assume $\text{wTCD}_p^\varepsilon(k, N)$. Furthermore, let $E \subset \mathbf{M}$ be compact and l -star shaped with respect to $x \in \text{spt } \mathbf{m}$. Then for every $0 < r < R < R_x$, the quantities from (5.11) satisfy*

$$\begin{aligned} \frac{\mathbf{v}(r)}{\mathbf{v}(R)} &\geq \frac{\int_0^r \sin_{\underline{k}_x/N}^N(t) dt}{\int_0^R \sin_{\underline{k}_x/N}^N(t) dt}, \\ \frac{\mathbf{s}(r)}{\mathbf{s}(R)} &\geq \frac{\sin_{\underline{k}_x/N}^N(r)}{\sin_{\underline{k}_x/N}^N(R)}. \end{aligned} \quad (5.14)$$

In particular, \mathbf{m} neither gives mass to points in $I^\pm(x)$, nor to spheres at x .

Finally, if \mathbf{m} has an atom at x , in fact we have

$$I^\pm(x) \cap \text{spt } \mathbf{m} = \emptyset. \quad (5.15)$$

Proof. For simplicity, assume $\mathbf{m}[E \cap U] > 0$ for every open neighborhood $U \subset \mathbf{M}$ of x ; we outline the necessary modifications in the more general case later.

We first prove the second estimate from (5.14). Without restriction, we may and will assume $\mathbf{s}(R) > 0$. Lastly, by applying the regularization procedure (3.6) to the function \underline{k}_x from (5.12) relative to the compact space $[0, R + \delta R]$ endowed with the Euclidean metric, we will also first suppose the induced function \underline{k}_x from (5.13) to be bounded and continuous on $\bar{\mathbf{B}}^l(x, R + \delta R) \cap E$ for sufficiently small $\delta > 0$.

Let $\mu_0 := \delta_{A_0}$ and, given $\delta > 0$, let $\mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathbf{m})$ be the uniform distribution of A_1 with respect to \mathbf{m} , where

$$\begin{aligned} A_0 &:= \{x\}, \\ A_1 &:= (\bar{\mathbf{B}}^l(x, R + \delta R) \setminus \mathbf{B}^l(x, R)) \cap E. \end{aligned}$$




FIGURE 2. An illustration of the proof of Theorem 5.9.

When $R + \delta R < R_x$, we first claim the existence of $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ with

$$\begin{aligned} \mathfrak{m}[(\bar{\mathbb{B}}^l(x, r + \delta r) \setminus \mathbb{B}^l(x, r)) \cap E]^{1/N} \\ \geq \int \sigma_{k_{x,\gamma}^+}^{(t)}(|\dot{\gamma}|) d\pi(\gamma) \mathfrak{m}[(\bar{\mathbb{B}}^l(x, R + \delta R) \setminus \mathbb{B}^l(x, R)) \cap E]^{1/N}. \end{aligned} \quad (5.16)$$

Since $\mathfrak{s}(R) > 0$, the measure μ_1 is well-defined, and (5.16) is not void.

Let $t := r/R$, and let $\mu_0^\varepsilon \in \mathcal{P}_c^{\text{ac}}(\mathbb{M}, \mathfrak{m})$ be the uniform distribution of

$$A_0^\varepsilon := \mathbb{B}^d(x, \varepsilon) \cap E.$$

Observe that unlike for A_1 , the definition of A_0^ε contains a metric ball. Here we choose $\varepsilon > 0$ sufficiently small in such a way that $\bar{\mathbb{B}}^d(x, \varepsilon)$ is compact according to Corollary 2.17 and Proposition 2.28, and $A_0^\varepsilon \times A_1 \in \{l > 0\}$, as shown in Figure 2. Finally, define the set $A_t^\varepsilon \subset \mathbb{M}$ according to (5.4), i.e.

$$A_t^\varepsilon := \mathbf{e}_t[G(A_0^\varepsilon, A_1)].$$

By Example 2.59, Remark 5.3, as well as (5.13), following the proof of Theorem 5.1 provides us with a plan $\pi^{\varepsilon, \delta} \in \text{OptTGeo}_p(\mu_0^\varepsilon, \mu_1)$ satisfying

$$\begin{aligned} \mathfrak{m}[A_t^{\varepsilon}]^{1/N} &\geq \sigma_{k_{\pi^\varepsilon}^+}^{(t)}(\mathbf{c}_{\pi^\varepsilon}) \mathcal{U}_N(\mu_1) \\ &\geq \int \sigma_{k_\gamma^+}^{(t)}(|\dot{\gamma}|) d\pi^\varepsilon(\gamma) \mathfrak{m}[A_1]^{1/N} \\ &\geq \int \sigma_{k_{x,\gamma}^+}^{(t)}(|\dot{\gamma}|) d\pi^\varepsilon(\gamma) \mathfrak{m}[A_1]^{1/N}. \end{aligned}$$

Note that $\mu_0 \otimes \mu_1$ constitutes the only coupling of μ_0 and μ_1 ; by construction, it is chronological. Tightness of timelike ℓ_p -optimal dynamical plans as asserted by

Proposition 2.66 gives the existence of a narrow limit $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ of a nonrelabeled subfamily of $(\pi^\varepsilon)_{\varepsilon>0}$. By the lower semicontinuous dependence of the last integrand above on $\gamma \in \text{TGeo}(\mathcal{M})$ provided by Lemma 3.6, [129, Lem. 4.3], and Levi's monotone convergence theorem we obtain

$$\mathfrak{m}[A_t]^{1/N} = \liminf_{\varepsilon \rightarrow 0} \mathfrak{m}[A_t^\varepsilon]^{1/N} \geq \int \sigma_{\underline{k}_x, \gamma/N}^{(t)}(|\dot{\gamma}|) d\pi(\gamma) \mathfrak{m}[A_1]^{1/N}.$$

The definition of A_0 and A_1 , the l -star shapedness of E , and the choice of t easily imply $A_t \subset (\bar{\mathbf{B}}^l(x, r + \delta r) \setminus \mathbf{B}^l(x, r)) \cap E$. This terminates the proof of (5.14).

If $\mathfrak{m}[E \cap U] = 0$ for some open neighborhood $U \subset \mathcal{M}$ of x , we still get (5.16) by the following modification of the previous argument. By using Remark 3.15 and Proposition A.3 below, from $\text{wTCD}_p^e(k, N)$ we obtain the entropic timelike measure contraction property $\text{TMCP}^e(k, N)$ from Definition A.1 for the causal reversal \mathcal{M}^{\leftarrow} in Example 2.45. This suffices to run the above computation with A_0 and μ_0 in place of A_0^ε and μ_0^ε , respectively.

From (5.16), letting $\delta \rightarrow 0$, and following the proof of [80, Thm. 5.9] we get

$$\frac{\mathfrak{s}(r)}{\mathfrak{s}(R)} \geq \sigma_{\underline{k}_x/N}^{(t)}(R)^N = \frac{\sin_{\underline{k}_x/N}^N(r)}{\sin_{\underline{k}_x/N}^N(R)},$$

where the last equality is a consequence of Lemma 3.6, Theorem 5.5, and our choice of t . In particular, as in the proof of [121, Thm. 2.3] we obtain

$$\mathfrak{m}[(\bar{\mathbf{B}}^l(x, r) \setminus \mathbf{B}^l(x, r)) \cap E] = 0.$$

This implies the claim about \mathfrak{m} -negligibility of spheres in the future of x . Indeed, using Polishness of τ and Proposition 2.28 we construct countably many precompact, mutually disjoint sets $(A_i)_{i \in \mathbf{N}}$ covering $\{l(x, \cdot) = r\}$ with $A_i \Subset I^+(x)$ for every $i \in \mathbf{N}$. The claim easily follows by applying Lemma 5.8 and the previous discussion to $C := \bar{A}_i$, where $i \in \mathbf{N}$. The case of past spheres is dealt with analogously, recalling Remark 3.15.

To show $\mathfrak{m}[\{y\}] = 0$ for every $y \in I^+(x)$, apply the previous argument to $r := l(x, y)$ and the compact, l -star shaped set E from Lemma 5.8 defined for $C := \{y\}$; observe that $l(x, y) < R_x$ holds here by nontrivial chronology. By applying the preceding argument to the causal reversal from Example 2.45 while taking Remark 3.15 into account, we obtain the claim about the \mathfrak{m} -measure of truncated spheres in the past of x , points in $I^-(x)$, and truncated past light cones at x .

Lastly, the same proof as for [80, Thm. 5.9] allows us to remove the continuity hypothesis on \underline{k}_x in order to finally obtain the second inequality in (5.14). In turn, this implies the first inequality from (5.14) by [129, Lem. 18.9].

It remains to prove (5.15) in the case \mathfrak{m} gives mass to x . Suppose to the contrary that $y \in I^-(x) \cap \text{spt } \mathfrak{m}$. Then applying the previous argumentation to y replacing x would entail the contradiction $\mathfrak{m}[\{x\}] = 0$. In an analogous way, we falsify the assumption $y \in I^+(x) \cap \text{spt } \mathfrak{m}$. \square

By nontrivial chronology, the following is immediate from Theorem 5.9.

Corollary 5.10 (No atoms). *The reference measure \mathfrak{m} of a $\text{wTCD}_p^e(k, N)$ space has no atoms if it is fully supported.*

The following is a consequence of Theorem 5.9 and argued as for [93, Thm. 5.2]; alternatively, it can be shown by using Lemma 3.21 together with Example 2.43 and [21, Rem. 3.19, Thm. 3.35].

Corollary 5.11 (Dimension bound). *Assume \mathcal{M} is the metric measure spacetime induced by a globally hyperbolic Lipschitz spacetime, where \mathfrak{m} is the top-dimensional*

Lorentzian Hausdorff measure defined in [93, Def. 2.3]. If \mathcal{M} satisfies $\text{wTCD}_p^e(k, N)$, then its geometric Hausdorff dimension in the sense of [93, Def. 3.1] obeys

$$\dim^l \mathbf{M} \leq N + 1. \quad (5.17)$$

6. ONE-DIMENSIONAL LOCALIZATION/NEEDLE DECOMPOSITION

In this chapter, we extend the Lorentzian localization paradigm initiated in [44]. We treat a more general class of potentials and cover spaces that might fail to be timelike nonbranching. A related extension is pursued for different purposes in [1]. Our plan is as follows.

- We introduce the transport sets with and without bad points, and describe transport relations \preceq_u , \succeq_u , and \sim_u thereon.
- The relation \sim_u is an equivalence relation on the transport set without bad points, Theorem 6.14.
- In Corollary 6.16, we show every chain with respect to \preceq_u is parametrized by exactly one l -geodesic.
- Under natural conditions — which can be inferred from curvature and timelike essential nonbranching assumptions using the existence of optimal maps (Theorem 3.23) and local-to-global property (Theorem 3.30) — we show the set of bad points thrown away to get Theorem 6.14 is \mathfrak{m} -negligible, Corollary 6.29.
- Assuming $\text{wTCD}_p^e(k, N)$ and timelike p -essential nonbranching, we show that on the transport set without bad points, \mathfrak{m} admits a disintegration into “absolutely continuous” needles, Theorem 6.31 and Proposition 6.33.
- Finally, Theorem 6.37 ensures a.e. needle constitutes a $\text{CD}(k, N)$ metric measure space in the sense of Ketterer [80].

We recall our assumption of τ being Polish. This is implicitly used various times when we apply results about Suslin measurability [118]; recall Subsection 2.6.1 for relevant basics of Suslin sets.

6.1. Framework. Throughout this chapter, we fix

- a Borel subset $E \subset \mathbf{M}$ which is l -geodesically convex, i.e. for every $\gamma \in \text{TGeo}(\mathbf{M})$ such that $\gamma_0, \gamma_1 \in E$, we have $\gamma_{[0,1]} \subset E$, and
- a Borel function $u: E \rightarrow \mathbf{R}$ obeying the *reverse l -Lipschitz condition*⁸, i.e. every $x, y \in E$ satisfies

$$u(y) - u(x) \geq l(x, y). \quad (6.1)$$

A concrete example from [44] fitting in this setting is given in Example 6.1 below. Observe that a subset $E \subset \mathbf{M}$ is l -geodesically convex with respect to a metric spacetime if and only if it is l^\leftarrow -geodesically convex relative to its causal reversal from Example 2.45. In addition, a function $u: E \rightarrow \mathbf{R}$ obeys (6.1) if and only if $-u$ satisfies (6.1) with l replaced by l^\leftarrow .

In addition, we impose mild assumptions on \mathfrak{m} , namely

- \mathfrak{m} has full topological support, and
- there exists a Borel weight function $w: \mathcal{T}_u^{\text{bp}} \rightarrow (0, \infty)$ with $\inf w(C) > 0$ for every compact set $C \subset E$ and

$$\int_{\mathcal{T}_u^{\text{bp}}} w \, \mathfrak{m} = 1; \quad (6.2)$$

in particular $\mathfrak{m}[\mathcal{T}_u^{\text{bp}}] > 0$. Here $\mathcal{T}_u^{\text{bp}} \subset E$ is the transport set with bad points from Definition 6.6 below.

⁸In [14], we refer to this property as *1-steepness* of u on E .

The hypothesis $\text{spt } \mathfrak{m} = \mathbb{M}$ could be dropped by appropriately restricting all definitions below to $\text{spt } \mathfrak{m}$, but we refrain from that in favor of a simpler presentation. On the other hand, the existence of a weight function w holds in high generality, e.g. on every proper Polish space [43, Lem. 3.3] or simply if $\mathfrak{m}[E]$ is finite (provided $\mathfrak{m}[T^{\text{bP}}] > 0$ in both cases). It is used to normalize \mathfrak{m} in order to apply the disintegration theorem in Section 6.5.

Example 6.1 (Compatibility with [44]). To give a concrete example originating in [44] from which e.g. the synthetic Hawking singularity theorem was obtained, we first recall basic definitions [44, 63] for a set $V \subset \mathbb{M}$ (assuming \mathbb{M} is proper).

We call V *achronal* provided $x \not\ll y$ for every $x, y \in V$; in other words, $I^+(V)$ and $I^-(V)$ have empty intersection. This immediately implies well-definedness of the *signed time* $\tau_V: \mathbb{M} \rightarrow [-\infty, \infty]$ to V given by

$$\tau_V(x) := \sup_{y \in V} l_+(y, x) - \sup_{y \in V} l_+(x, y).$$

Since l_+ is continuous, τ_V is lower semicontinuous on $I^+(V) \cup V$.

Let V be such an achronal Borel set which is in addition *future timelike complete* [63, p. 367], briefly FTC, i.e. for every $y \in I^+(V)$ the intersection $J^-(y) \cap V$ has compact closure relative to V ; observe that this covers especially the case when V is a singleton by Corollary 2.14. Then define

$$\begin{aligned} E &:= I^+(V) \cup V, \\ u &:= \tau_V|_{E^2}. \end{aligned}$$

By achronality of V and Lemma 2.3, for every $\gamma \in \text{TGeo}(\mathbb{M})$ with $\gamma_0, \gamma_1 \in E$ we have $\gamma_t \in I^+(V)$ for every $t \in (0, 1]$, which shows l -geodesic convexity of E . The inequality (6.1) on E^2 is a consequence of the hypothesized FTC property and is shown with the same proof as for [44, Lem. 4.1].

The existence of a function w as above follows from [44, Lems. 4.4, 4.10]. In fact, under the natural condition $\mathfrak{m}[V] = 0$ we have

$$\mathfrak{m} \llcorner I^+(V) = \mathfrak{m} \llcorner \mathcal{T}_u^{\text{bP}}. \quad (6.3)$$

6.2. Transport relations and transport sets. Now we construct

- the transport relation $E_{\sim_u}^{2, \text{bP}}$ and the transport set $\mathcal{T}_u^{\text{bP}}$ relative to E and u with bad points, and
- their counterparts $E_{\sim_u}^2$ and \mathcal{T}_u without bad points.

As shown in Theorem 6.14, $E_{\sim_u}^{2, \text{bP}}$ will be an equivalence relation on \mathcal{T}_u . (In fact, \mathcal{T}_u can be chosen slightly larger.) Compared to [44, Prop. 4.5], timelike nonbranching is not available to us; our approach is instead inspired by the work [37] for metric measure spaces, see also [16, 35, 40, 84].

6.2.1. Transport with bad points. The construction of the indicated sets $E_{\sim_u}^{2, \text{bP}}$ and $\mathcal{T}_u^{\text{bP}}$ follows [44, Sec. 4.1] and is briefly recapitulated here.

We define relations \preceq_u and \succeq_u on E by declaring

$$\begin{aligned} x \preceq_u y &:\iff x = y \quad \text{or} \quad u(y) - u(x) = l(x, y) > 0, \\ x \succeq_u y &:\iff y \preceq_u x. \end{aligned}$$

We consider the associated sets

$$\begin{aligned} E_{\preceq_u}^2 &:= \{(x, y) \in E^2 : x \preceq_u y\}, \\ E_{\succeq_u}^2 &:= \{(x, y) \in E^2 : x \succeq_u y\}. \end{aligned} \quad (6.4)$$

Our assumptions on E and u guarantee $E_{\preceq_u}^2$ and $E_{\succeq_u}^2$ are Borel.

The following straightforward properties should give the reader a better feeling for the definitions of \preceq_u and \succeq_u , respectively.

Lemma 6.2 (Partial ordering). *The relations \preceq_u and \succeq_u are partial orders on E .*

Proof. We only show the claim for \preceq_u , the proof for \succeq_u is completely analogous. Let $x, y, z \in E$. By definition, we have $x \preceq_u x$. If $x \preceq_u y$ and $x \succeq_u y$, then $x = y$ since otherwise $x \ll y$ or $y \ll x$, a contradiction to the respective other half of the assumption. Lastly, assume $x \preceq_u y$ and $y \preceq_u z$. Without restriction, we assume x, y , and z are distinct. Then $x \ll y \ll z$, and in particular $x \ll z$ by the push-up property from Lemma 2.3. By the reverse l -Lipschitz property (6.1) of u and the reverse triangle inequality for l , we easily obtain

$$\begin{aligned} u(z) - u(x) &\geq l(x, z) \\ &\geq l(x, y) + l(y, z) \\ &= [u(y) - u(x)] + [u(z) - u(y)] \\ &= u(z) - u(x), \end{aligned}$$

which forces equality to hold throughout. \square

In a similar way, the following is verified [44, Lem. 4.2].

Lemma 6.3 (Cyclical monotonicity). *$E_{\preceq_u}^2$ is l -cyclically monotone, while $E_{\succeq_u}^2$ is l^\leftarrow -cyclically monotone.*

Next we define

$$\begin{aligned} G &:= \{\gamma \in \text{TGeo}(\mathbf{M}) : \gamma_s \preceq_u \gamma_t \text{ for every } 0 \leq s < t \leq 1\}, \\ G^\leftarrow &:= \{\gamma \in \text{TGeo}^\leftarrow(\mathbf{M}) : \gamma_s \succeq_u \gamma_t \text{ for every } 0 \leq s < t \leq 1\}. \end{aligned} \quad (6.5)$$

As every curve in G and G^\leftarrow is timelike in the respective causal structure, G and G^\leftarrow do not contain constant curves by Corollary 2.14. Moreover, note that if we would a priori replace l by l^\leftarrow , then G^\leftarrow precisely corresponds to the set G , and vice versa — heuristically, $(G^\leftarrow)^\leftarrow = G$.

The key property of G relative to \preceq_u is contained in the following lemma. An evident modification relative to \succeq_u holds for G^\leftarrow ; similarly for Corollary 6.5.

Lemma 6.4 (Propagation of the transport relation). *Let $\gamma \in \text{TGeo}(\mathbf{M})$ such that $\gamma_0 \preceq_u \gamma_1$. Then $\gamma_s \preceq_u \gamma_t$ for every $0 \leq s < t \leq 1$; in other words, $\gamma \in G$.*

Proof. By (6.1), the assumption $\gamma_0 \preceq_u \gamma_1$, and affine parametrization of γ ,

$$\begin{aligned} u(\gamma_t) - u(\gamma_s) &= u(\gamma_1) - u(\gamma_0) - [u(\gamma_s) - u(\gamma_0)] - [u(\gamma_1) - u(\gamma_t)] \\ &\leq l(\gamma_0, \gamma_1) - l(\gamma_0, \gamma_s) - l(\gamma_t, \gamma_1) \\ &= l(\gamma_s, \gamma_t). \end{aligned}$$

Using (6.1) again, the previous inequality is in fact an identity. \square

Corollary 6.5 (Continuity of u along rays). *For every $\gamma \in G$, the function $u \circ \gamma$ is continuous on $[0, 1]$.*

These considerations motivate the definition of a further relation \sim_u on E via

$$x \sim_u y \quad :\iff \quad x \preceq_u y \quad \text{or} \quad x \succeq_u y.$$

Definition 6.6 (Transport with bad points). *We define*

- a. *the transport relation with bad points by*

$$E_{\sim_u}^{2, \text{bp}} := E_{\preceq_u}^2 \cup E_{\succeq_u}^2,$$

- b. *the transport set with bad points by*

$$\mathcal{T}_u^{\text{bp}} := \text{pr}_1[E_{\sim_u}^{2, \text{bp}} \setminus \text{diag}(\mathbf{M}^2)].$$

Clearly, $E_{\sim_u}^{2,\text{bp}}$ is Borel; since the background topology τ is Polish, this implies $\mathcal{T}_u^{\text{bp}}$ is Suslin.

By definition of $\mathcal{T}_u^{\text{bp}}$, for every $x \in \mathcal{T}_u^{\text{bp}}$ there exists $y \in E \setminus \{x\}$ with $x \preceq_u y$ or $x \succeq_u y$, and in the respective case $x \ll y$ or $y \ll x$; in particular, $y \in \mathcal{T}_u^{\text{bp}}$. Yet, this does *not* mean $\mathcal{T}_u^{\text{bp}}$ has no extremal points according to Definition 6.9. Moreover, \sim_u is not generally transitive on the whole of $\mathcal{T}_u^{\text{bp}}$.

6.2.2. *Bad points.* Let us now make precise what we mean by bad points through the following two definitions.

Definition 6.7 (Timelike branching points). *We define*

- a. *the timelike forward branching set by*

$$B := \{x \in \mathcal{T}_u^{\text{bp}} : \text{there exist } y_1, y_2 \in \mathcal{T}_u^{\text{bp}} \\ \text{such that } x \preceq_u y_1 \text{ and } x \preceq_u y_2 \text{ yet } y_1 \not\preceq_u y_2\},$$

- b. *the timelike backward branching set by*

$$B^\leftarrow := \{y \in \mathcal{T}_u^{\text{bp}} : \text{there exist } x_1, x_2 \in \mathcal{T}_u^{\text{bp}} \\ \text{such that } x_1 \preceq_u y \text{ and } x_2 \preceq_u y \text{ yet } x_1 \not\preceq_u x_2\}.$$

Any point in B or B^\leftarrow is called timelike forward branching point or timelike backward branching point, respectively.

Remark 6.8 (About Definition 6.7). Let $x \in B$ be fixed, and let $y_1, y_2 \in \mathcal{T}_u^{\text{bp}}$ be as in the previous definition of B . Then x, y_1 , and y_2 are necessarily distinct; in particular, we have $x \ll y_1$ and $x \ll y_2$. Moreover, if a point $x' \in \mathcal{T}_u^{\text{bp}}$ satisfies $x' \preceq_u x$ in the previous situation, then $x' \in B$.

Analogous statements hold for B^\leftarrow .

Definition 6.9 (Extremal points). *We define*

- a. *the initial point set by*

$$a := \{x \in \mathcal{T}_u^{\text{bp}} : \text{no } x^- \in \mathcal{T}_u^{\text{bp}} \setminus \{x\} \text{ satisfies } x^- \preceq_u x\},$$

- b. *the final point set by*

$$b := \{x \in \mathcal{T}_u^{\text{bp}} : \text{no } y^+ \in \mathcal{T}_u^{\text{bp}} \setminus \{x\} \text{ satisfies } y^+ \succeq_u x\}.$$

Any point in a or b is called initial point or final point, respectively.

In general, there is no relation between the sets B, B^\leftarrow, a , and b .

Any point in the union $B \cup B^\leftarrow \cup a \cup b$ is called *bad point*. This terminology is motivated by the following difficulties encountered in the sequel.

- Timelike branching points prevent \sim_u to be an equivalence relation on $\mathcal{T}_u^{\text{bp}}$, cf. Theorem 6.14.
- On the other hand, the presence of extremal points will have to be waived in order to prove absolute continuity of the conditional densities along a.e. needle in our disintegration Theorem 6.31 later, cf. Proposition 6.33.

For these reasons, these sets will be excluded from $\mathcal{T}_u^{\text{bp}}$ in the next subsection. It is highly nontrivial — and in general false — that $\mathcal{T}_u^{\text{bp}} \setminus (B \cup B^\leftarrow \cup a \cup b)$ remains sufficiently large, yet its difference from $\mathcal{T}_u^{\text{bp}}$ is \mathfrak{m} -negligible under natural conditions, as demonstrated in Section 6.4.

The sets B and B^\leftarrow are Suslin; the tedious proof is similar to the argument for Lemma 6.23 below (which does not create a circular reasoning). On the other hand, Suslin measurability of a and b is easy to infer from the identities

$$a = \mathcal{T}_u^{\text{bp}} \setminus \text{pr}_2[E_{\preceq_u}^2 \setminus \text{diag}(\mathbb{M}^2)], \\ b = \mathcal{T}_u^{\text{bp}} \setminus \text{pr}_1[E_{\succeq_u}^2 \setminus \text{diag}(\mathbb{M}^2)].$$




FIGURE 3. An illustration of the proof of Lemma 6.10.

To get a grip on Definition 6.7, we prove the following lemma. Its statement is only needed again in Lemma 6.23 below, but we will revisit its proof technique already in the proof of Lemma 6.15. Furthermore, it highlights the connection of timelike branching points and branching phenomena for l -geodesics. A further, perhaps more illuminating evidence is deferred to Example 6.12 below, where we relate B and B^\leftarrow to certain cut loci for specific choices of E and u .

We only concentrate on timelike forward branching. With evident changes, all statements hold in the timelike backward branching case as well.

Lemma 6.10 (Branching points vs. branching geodesics I). *Let $x \in \mathcal{T}_u^{\text{bp}}$. Then $x \in B$ if and only if there exist $\gamma^1, \gamma^2 \in G$ from (6.5) with*

$$\min_{t \in [0,1]} d(\gamma_t^1, \gamma_t^2) > 0$$

and such that for every $0 \leq t \leq 1$, we have $x \preceq_u \gamma_t^1$ and $x \preceq_u \gamma_t^2$ yet $\gamma_t^1 \not\sim_u \gamma_t^2$, and

$$u(\gamma_t^1) = u(\gamma_t^2).$$

Proof. The backward implication is clear by definition of B .

To prove the forward implication — which is depicted in Figure 3 — assume $x \in B$, and let $y_1, y_2 \in \mathcal{T}_u^{\text{bp}}$ be as in Definition 6.7. By symmetry of \sim_u we may and will assume $u(y_2) \geq u(y_1)$. By Remark 6.8 and Lemma 6.4 there are $\eta^1, \eta^2 \in G$ with $\eta_0^1 = \eta_0^2 = x$, $\eta_1^1 = y_1$, and $\eta_1^2 = y_2$; in particular, $\eta_t^1, \eta_t^2 \in \mathcal{T}_u^{\text{bp}}$ for every $0 \leq t \leq 1$. Since $x \preceq_u y_1$, we obtain

$$u(y_2) \geq u(y_1) = u(x) + l(x, y_1) > u(x).$$

Corollary 6.5 and the intermediate value theorem imply

$$u(\eta_\alpha^2) = u(y_1) \tag{6.6}$$

for some $0 < \alpha \leq 1$ which will be henceforth fixed. Then $\eta_\alpha^2 \neq y_1$, since otherwise $y_1 \preceq_u y_2$ by Lemma 6.4, in conflict with our assumption $y_1 \not\sim_u y_2$. We even have $\eta_\alpha^2 \not\sim_u y_1$, since otherwise $\eta_\alpha^2 \ll y_1$ or $y_1 \ll \eta_\alpha^2$ by the previous observation. In the first case, by (6.6) the observation $0 = u(y_1) - u(\eta_\alpha^2) = l(\eta_\alpha^2, y_1) > 0$ leads to a contradiction; the case $y_1 \ll \eta_\alpha^2$ is falsified analogously.

For every $0 \leq r \leq 1$, by employing the property $\eta^1, \eta^2 \in G$ and (6.6), a direct computation shows

$$-u(\eta_{\alpha r}^2) = -u(\eta_r^1).$$

As $\eta_\alpha^2 \neq y_1 = \eta_1^1$, continuity implies the existence of $\delta > 0$ such that $\eta_{1-s}^1 \neq \eta_{\alpha(1-s)}^2$ and, following the above argument, that $\eta_{1-s}^1 \not\sim_u \eta_{\alpha(1-s)}^2$ for every $0 \leq s \leq \delta$. A rescaling argument then gives $\gamma^1, \gamma^2 \in G$ with the desired properties. \square

Remark 6.11 (Branching points vs. branching geodesics II). By a slight abuse of terminology, let us call a point $x \in M$ *timelike forward branching* if there exist $\gamma^1, \gamma^2 \in \text{TGeo}(M)$ and $0 < \delta < 1$ with $x \in \gamma^1|_{(0,\delta)} = \gamma^2|_{(0,\delta)}$ yet $\gamma^1 \neq \gamma^2$. Then the proof of Lemma 6.10 shows that if $x \in B \setminus a$, then x is such a timelike forward branching point. Conversely, a timelike forward branching point $x \in \mathcal{J}_u^{\text{bp}}$ belongs to B provided $x \preceq_u \gamma_t^1$ and $x \preceq_u \gamma_t^2$ yet $\gamma_t^1 \not\sim_u \gamma_t^2$ for some $\delta < t \leq 1$.

With similar arguments, the reader may verify the following.

Example 6.12 (Timelike cut loci). Let $o \in M$ be fixed. Let E and u be defined as in Example 6.1 in terms of $V := \{o\}$, i.e.

$$\begin{aligned} E &:= I^+(o) \cup \{o\}, \\ u &:= l(o, \cdot). \end{aligned}$$

Note that in this case, we have $\mathcal{J}_u^{\text{bp}} = I^+(o)$. Then the induced set B^\leftarrow coincides with the *future timelike cut locus* $\text{TCut}^+(o)$ of o . The latter is the set of $y \in I^+(o)$ such that two distinct $\gamma^1, \gamma^2 \in \text{TGeo}(M)$ connect o and y .

Mutatis mutandis, considering instead

$$\begin{aligned} E &:= I^-(o) \cup \{o\}, \\ u &:= -l(\cdot, o) \end{aligned}$$

we observe the set B coincides with the *past timelike cut locus* $\text{TCut}^-(o)$, which is the set of $x \in I^-(o)$ such that two distinct $\gamma^1, \gamma^2 \in \text{TGeo}(M)$ connect x to o .

6.2.3. Transport without bad points.

Definition 6.13 (Transport without bad points). *We define*

- a. *the transport set by*

$$\mathcal{J}_u := \mathcal{J}_u^{\text{bp}} \setminus (B \cup B^\leftarrow \cup a \cup b),$$

- b. *the transport relation by*

$$E_{\sim_u}^2 := E_{\sim_u}^{2,\text{bp}} \cap \mathcal{J}_u^2.$$

In other words, $E_{\sim_u}^2$ corresponds to the restriction of the relation \sim_u to \mathcal{J}_u .

By the discussions from Subsection 6.2.2, the sets \mathcal{J}_u and $E_{\sim_u}^2$ are both Suslin.

Theorem 6.14 (Equivalence relation). *The relation \sim_u constitutes an equivalence relation on $\mathcal{J}_u^{\text{bp}} \setminus (B \cup B^\leftarrow)$, a fortiori on its subset \mathcal{J}_u .*

Proof. We abbreviate $\mathcal{J}'_u := \mathcal{J}_u^{\text{bp}} \setminus (B \cup B^\leftarrow)$. By definition of \sim_u , we have $x \sim_u x$ for every $x \in \mathcal{J}'_u$. Moreover, \sim_u is symmetric by construction. To show transitivity, let $x, y, z \in \mathcal{J}'_u$ be mutually distinct points with $x \sim_u y$ and $y \sim_u z$. By Lemma 6.2, we only have to consider the following two cases.

Case 1. Assume $x \preceq_u y$ and $y \succeq_u z$, hence $z \preceq_u y$. This immediately implies $x \sim_u z$; indeed, otherwise $x \not\sim_u z$ contradicts the assumption $y \notin B^\leftarrow$.

Case 2. The case $x \succeq_u y$ and $y \preceq_u z$ is treated similarly to Case 3 by contradicting the assumption $y \notin B$. \square

The equivalence class of $x \in \mathcal{J}_u$ with respect to \sim_u is denoted \tilde{x}^u .

6.3. Ray map and quotient map. In order to disintegrate \mathbf{m} with respect to u later, we now have to build an \mathbf{m} -measurable quotient map of $E_{\sim_u}^2$.

First, in Lemma 6.15 we prove that every equivalence class of \sim_u is crossed by exactly one l -geodesic. This will make the so-called ray map \mathbf{g} from Corollary 6.16 well-defined. In turn, \mathbf{g} will be used to construct the above mentioned quotient map in Proposition 6.18.

Lemma 6.15 (Nonbranching rays). *Let $x, y_1, y_2 \in \mathcal{T}_u$ satisfy $x \sim_u y_1$ and $x \sim_u y_2$. Then there exists $\gamma \in G$ with $\{x, y_1, y_2\} \subset \gamma_{[0,1]}$.*

Moreover, if any $\gamma' \in G$ satisfies the previous conclusion too, there exists $\gamma'' \in G$ such that $\gamma_{[0,1]} \cup \gamma'_{[0,1]} \subset \gamma''_{[0,1]}$.

Proof. To show the first statement, without restriction we assume x, y_1 , and y_2 are distinct, otherwise we just apply Lemma 6.4. We have to distinguish four cases.

Case 1. Assume $y_1 \preceq_u x$ and $x \preceq_u y_2$. By Remark 6.8, we have $y_1 \ll x \ll y_2$, and by transitivity of \preceq_u we have $y_1 \preceq_u y_2$. Therefore,

$$\begin{aligned} l(y_1, y_2) &= u(y_2) - u(y_1) \\ &= [u(y_2) - u(x)] + [u(x) - u(y_1)] \\ &= l(y_1, x) + l(x, y_2). \end{aligned}$$

As in the proof of Lemma 2.46, the concatenation of two appropriate elements in $\text{TGeo}(\mathbf{M})$ — one going from y_1 to x , one going from x to y_2 — yields an element $\gamma \in \text{TGeo}(\mathbf{M})$ from y_1 to y_2 . By Lemma 6.4, we get $\gamma \in G$ with $\{y_1, x, y_2\} \subset \gamma_{[0,1]}$ and hence the desired curve.

Case 2. Analogously to Case 1, the assumption $y_2 \preceq_u x$ and $x \preceq_u y_1$ yields an element $\gamma \in G$ with $\{y_2, x, y_1\} \subset \gamma_{[0,1]}$.

Case 3. Assume $x \preceq_u y_1$ and $x \preceq_u y_2$. Since \sim_u is an equivalence relation on \mathcal{T}_u by Theorem 6.14, and since $y_1 \neq y_2$, either $y_1 \ll y_2$ or $y_2 \ll y_1$ by Remark 6.8. Without restriction, we assume

$$u(y_2) > u(y_1) > u(x).$$

Indeed, equality in the second place never occurs under our standing hypotheses, and a possible equality in the first place can be easily contradicted as in the proof of Lemma 6.10. Now, assume to the contrapositive that no $\gamma \in G$ from x to y_2 passes through y_1 . Then by geodesy, Lemma 6.4, and a continuity argument there exist $\gamma \in G$ and $0 < t < 1$ with $\gamma_0 = x$, $\gamma_1 = y_2$, $u(\gamma_t) = u(y_1)$, $x \preceq_u \gamma_t$, and $\gamma_t \neq y_1$. Arguing again as in the proof of Lemma 6.10, we get $\gamma_t \not\sim_u y_1$, but this contradicts the assumption $x \notin B$.

Case 4. The case $y_1 \preceq_u x$ and $y_2 \preceq_u x$ is treated as in Case 3 by falsifying the assumption $x \in B^{\leftarrow}$.

The last statement is argued similarly. \square

Now we define the set

$$D_{\mathbf{g}} = \{(\alpha, t, y) \in \mathcal{T}_u \times \mathbf{R} \times \mathcal{T}_u : \alpha \sim_u y, l_+(\alpha, y) - l_+(y, \alpha) = t\}$$

and the multi-valued map $\mathbf{g}: \mathcal{D}(\mathbf{g}) \rightarrow 2^{\mathcal{T}_u}$ in terms of its graph

$$\text{graph } \mathbf{g} := (\mathcal{D}(\mathbf{g}) \times \mathcal{T}_u) \cap D_{\mathbf{g}},$$

where $\mathcal{D}(\mathbf{g}) \subset \mathcal{T}_u \times \mathbf{R}$ is given by

$$\mathcal{D}(\mathbf{g}) := \{(\alpha, t) \in \mathcal{T}_u \times \mathbf{R} : D_{\mathbf{g}} \cap (\{(\alpha, t)\} \times \mathcal{T}_u) \neq \emptyset\}.$$

In other words, to a given pair $(\alpha, t) \in \mathcal{D}(\mathbf{g})$ the multi-valued map \mathbf{g} assigns the nonempty set of points $y \in \mathcal{T}_u$ with “distance” to α being equal to t . It is clear that $(\alpha, 0) \in \mathcal{D}(\mathbf{g})$ for every $\alpha \in \mathcal{T}_u$.

Corollary 6.16 (Ray map). *The map \mathbf{g} defined above is single-valued on $\mathcal{D}(\mathbf{g})$, hence may and will be identified with a nonrelabeled Borel map $\mathbf{g}: \mathcal{D}(\mathbf{g}) \rightarrow \mathcal{T}_u$.*

This map obeys the following properties for every $\alpha \in \text{pr}_1(\mathcal{D}(\mathbf{g}))$.

(i) **Intervals.** *The set*

$$\mathcal{D}(\mathbf{g})(\alpha) := \text{pr}_2[\mathcal{D}(\mathbf{g}) \cap (\{\alpha\} \times \mathbf{R})]$$

is a nonempty convex subset of \mathbf{R} , i.e. an interval.

(ii) **Totality.** *The map $F: \mathcal{D}(\mathbf{g})(\alpha) \rightarrow \mathcal{T}_u$ defined by $F(t) := \mathbf{g}(\alpha, t)$ is injective. Moreover, it maps surjectively onto the ray $\tilde{\alpha}^u$.*

(iii) **Order isometry.** *For every $s, t \in \mathcal{D}(\mathbf{g})(\alpha)$ with the property $s < t$, we have $F(s) \preceq_u F(t)$ and*

$$l(F(s), F(t)) = t - s.$$

Proof. We first show \mathbf{g} is single-valued on $\mathcal{D}(\mathbf{g})$. By definition of the latter, it suffices to prove $\#\mathbf{g}(\alpha, t) \leq 1$ for every $(\alpha, t) \in \mathcal{D}(\mathbf{g})$. We only treat the case $t > 0$; the case $t < 0$ is addressed similarly, while the case $t = 0$ is clear from Corollary 2.14. Thus, let $y_1, y_2 \in \mathbf{g}(\alpha, t)$, and note that y_1 and y_2 are distinct from α . The point α is no timelike forward branching point by definition of \mathcal{T}_u , thus $y_1 \sim_u y_2$. Then Lemma 6.15 implies that x, y_1 , and y_2 lie on a single l -geodesic segment. Since $l(\alpha, y_1) = l(\alpha, y_2) = t > 0$, this forces $y_1 = y_2$.

Borel measurability of \mathbf{g} follows from analyticity of graph \mathbf{g} , compare with the proof of Lemma 6.23 below.

For item (i), observe that $\mathbf{g}(\alpha, 0) = \alpha$ implies $0 \in \mathcal{D}(\mathbf{g})(\alpha)$. Next, we claim $[0, t] \subset \mathcal{D}(\mathbf{g})(\alpha)$ for every $t \in \mathcal{D}(\mathbf{g})(\alpha) \cap (0, \infty)$; the case $t \in \mathcal{D}(\mathbf{g})(\alpha) \cap (-\infty, 0)$ follows by an analogous argument. By definition of \mathbf{g} , we have $\alpha \preceq_u F(t)$ with $l(\alpha, F(t)) = t$. Using Lemma 6.4, let $\gamma \in G$ connect α to $F(t)$. In particular, for $0 < s < t$ we have $\alpha \preceq_u \gamma_{s/t}$ and $l(\alpha, \gamma_{s/t}) = s$, whence $F(s) = \gamma_{s/t}$ as above.

Item (ii) follows from the definition of \mathbf{g} .

Item (iii) in the cases $0 \leq s < t$ or $s < t \leq 0$ follows as in (i) by identifying intermediate points of F and of l -geodesics. If $s < 0 < t$, by definition of \mathbf{g} and transitivity of \sim_u we obtain

$$\begin{aligned} l(F(s), F(t)) &= u(F(t)) - u(F(s)) \\ &= u(F(t)) - u(\alpha) - [u(F(s)) - u(\alpha)] \\ &= l(\alpha, F(t)) - l(\alpha, F(s)) \\ &= t - s. \end{aligned} \quad \square$$

Remark 6.17 (Openness). In fact, since \mathcal{T}_u has empty intersection with the set $a \cup b$ of extremal points, $\mathcal{D}(\mathbf{g})(\alpha)$ is an open subset of \mathbf{R} for every $\alpha \in \text{pr}_1(\mathcal{D}(\mathbf{g}))$.

For every $\alpha \in \text{pr}_1(\mathcal{D}(\mathbf{g}))$, the map $\mathbf{g}(\alpha, \cdot)$ follows the “ l -geodesic ray” $t \mapsto F(t)$ on $\mathcal{D}(\mathbf{g})(\alpha)$. We thus call \mathbf{g} the *ray map*.

Given this map \mathbf{g} , the proof of the following result is now analogous to [40, Lem. 3.9] — with d replaced by l — hence omitted. Observe that the similar argument from [44, Prop. 4.9] does not apply directly since u may attain negative values. Moreover, the σ -compactness in the proof of [40, Lem. 3.9] is unclear in our case and has to be replaced by mere Suslin measurability; this is why the map \mathbf{Q} below is stated as \mathfrak{m} -measurable, and not Borel measurable as in [40, Lem. 3.9].

Recall that a map $\mathbf{Q}: \mathcal{T}_u \rightarrow \mathcal{T}_u$ is called *quotient map* for the equivalence relation \sim_u on \mathcal{T}_u if $\text{graph } \mathbf{Q} \subset E_{\sim_u}^2$, and $\mathbf{Q}(x) = \mathbf{Q}(y)$ whenever $x, y \in \mathcal{T}_u$ obey $x \sim_u y$.

Proposition 6.18 (Quotient map). *There exists an \mathfrak{m} -measurable quotient map $\mathbf{Q}: \mathcal{T}_u \rightarrow \mathcal{T}_u$ for the equivalence relation \sim_u on \mathcal{T}_u .*

Moreover, the image of \mathbf{Q} can be written as a countable union of level sets of u . More precisely, there exist rational numbers $(a_i)_{i \in \mathbf{N}}$ and a sequence of mutually disjoint Suslin subsets $(Q_i)_{i \in \mathbf{N}}$ of \mathcal{T}_u such that $Q_i \subset u^{-1}(a_i)$ for every $i \in \mathbf{N}$ and

$$\mathbf{Q}(\mathcal{T}_u) = \bigcup_{i \in \mathbf{N}} Q_i.$$

Definition 6.19 (Quotient set and transport rays). *The set $Q := \mathbf{Q}(\mathcal{T}_u)$ will be termed the quotient set.*

Any set of the form $M_\alpha := \mathbf{Q}^{-1}(\alpha)$, $\alpha \in Q$, will be called a transport ray.

The following readily follows from basic properties of \mathbf{g} and \mathbf{Q} .

Corollary 6.20 (Right-inverse of \mathbf{g}). *The map $h: \mathcal{T}_u \rightarrow \mathcal{T}_u \times \mathbf{R}$ given by*

$$h(x) := \begin{cases} (x, 0) & \text{if } \mathbf{Q}(x) = x, \\ (\mathbf{Q}(x), l(x, \mathbf{Q}(x))) & \text{if } \mathbf{Q}(x) \neq x \text{ yet } x \preceq_u \mathbf{Q}(x), \\ (\mathbf{Q}(x), -l(\mathbf{Q}(x), x)) & \text{if } \mathbf{Q}(x) \neq x \text{ yet } \mathbf{Q}(x) \preceq_u x \end{cases}$$

is right-inverse to \mathbf{g} on \mathcal{T}_u .

6.4. Negligibility of bad points. As highlighted in Subsection 6.2.2, in general the set $B \cup B^{\leftarrow} \cup a \cup b$ removed from $\mathcal{T}_u^{\text{bp}}$ to define \mathcal{T}_u can be quite large. The goal of this section is to provide a natural condition (Corollary 6.29 below), which combines with Theorems 3.23 and 3.30 to ensure \mathbf{m} -negligibility of the set of bad points relative to $\mathcal{T}_u^{\text{bp}}$.

The hypotheses in all relevant results proven along the way decrease in their generality. Ultimately, they all hold under curvature hypotheses, cf. Theorem 6.31.

In the sequel, we call $\mu \in \mathcal{P}(\mathbf{M})$ an L^∞ -measure if it is \mathbf{m} -absolutely continuous with \mathbf{m} -essentially bounded Radon–Nikodým density.

6.4.1. Negligibility of timelike branching points. The goal now is to prove the following Theorem 6.21. For metric measure spaces, an analogous conclusion has already been observed to be implied by a similar hypothesis [37, 38, 39, 40] sometimes called “good transport behavior” [61]. Note that in our case, the condition on μ_0 and μ_1 below might be empty, e.g. if they do not admit any causal coupling. However, by Example 2.59 and Theorem 3.23 this holds for sufficiently many marginals.

Theorem 6.21 (Negligibility of timelike forward branching points). *Assume the existence of $0 < p < 1$ with the following property. For every L^∞ -measure $\mu_0 \in \mathcal{P}_c(\mathbf{M})$ and every $\mu_1 \in \mathcal{P}_c(\mathbf{M})$ such that*

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\},$$

every ℓ_p -optimal coupling of μ_0 and μ_1 is induced by a map. Then

$$\mathbf{m}[B] = 0.$$

Analogous conditions for the \mathbf{m} -negligibility of B^{\leftarrow} are given in Corollary 6.25.

The proof of Theorem 6.21 requires some preparations. First, we prove the correspondence of Lemma 6.10 between elements of B and suitable pairs of elements in G to hold in a measurable way. To this aim, we need the subsequent version of the von Neumann selection theorem [118, Thm. 5.5.2].

Recall that for any sets X and Y , a map $S: \text{pr}_1(F) \rightarrow Y$ is called *section* of a given set $F \subset X \times Y$ if $\text{graph } S \subset F$.

Theorem 6.22 (Von Neumann selection theorem). *Let X and Y be Polish spaces, $F \subset X \times Y$ be Suslin, and \mathcal{A} be the σ -algebra generated by all Suslin subsets of X . Then there exists an \mathcal{A} -measurable section $S: \text{pr}_1(F) \rightarrow Y$ of F .*

In the sequel, let X be the completion of $\text{TGeo}(\mathbf{M}) \cup \text{TGeo}^{\leftarrow}(\mathbf{M})$ with respect to the uniform distance d_∞ . By a continuity argument, X is Polish with respect to d_∞ . Furthermore, the Borel σ -algebra of X is equal to the Borel σ -algebra of $\mathbf{C}([0, 1]; \mathbf{M})$ relative to X .

Lemma 6.23 (Measurable selection of branching geodesics). *There exists an \mathfrak{m} -measurable map $S: B \rightarrow X^2$ such that for every $(x, \gamma^1, \gamma^2) \in \text{graph } S$ and every $0 \leq t \leq 1$, the triple $(x, \gamma_t^1, \gamma_t^2)$ obeys the conclusions from Lemma 6.10. That is, we have $\gamma^1, \gamma^2 \in G$ such that*

$$\min_{t \in [0, 1]} d(\gamma_t^1, \gamma_t^2) > 0,$$

furthermore $x \preceq_u \gamma_t^1$ and $x \preceq_u \gamma_t^2$ yet $\gamma_t^1 \not\prec_u \gamma_t^2$, and

$$u(\gamma_t^1) = u(\gamma_t^2).$$

Proof. In the notation of Theorem 6.22, let \mathcal{A} denote the σ -algebra generated by all Suslin sets of \mathbf{M} . For later use, we claim that

$$F' := F_1 \cap F_2 \cap F_3 \cap F_4 \cap F_5$$

is a Suslin set, where

$$F_1 := \mathcal{T}_u^{\text{bp}} \times G^2,$$

$$F_2 := \{(x, \gamma^1, \gamma^2) \in \mathbf{M} \times X^2 : x \preceq_u \gamma_0^1\},$$

$$F_3 := \{(x, \gamma^1, \gamma^2) \in \mathbf{M} \times X^2 : x \preceq_u \gamma_0^2\},$$

$$F_4 := \mathbf{M} \times \{(\gamma^1, \gamma^2) \in X^2 : d(\gamma_0^1, \gamma_0^2) > 0\},$$

$$F_5 := \mathbf{M} \times \{(\gamma^1, \gamma^2) \in X^2 : u(\gamma_0^1) = u(\gamma_0^2) \text{ and } u(\gamma_1^1) = u(\gamma_1^2)\}.$$

Clearly, F_4 is Borel measurable. To prove that F_1 is Suslin, it suffices to prove that G is. By Lemma 6.4,

$$\begin{aligned} G &= \{\gamma \in \text{TGeo}(\mathbf{M}) : u(\gamma_1) - u(\gamma_0) = l(\gamma_0, \gamma_1)\} \\ &= \text{TGeo}(\mathbf{M}) \cap \{\gamma \in X : u(\gamma_1) = u(\gamma_0) + l(\gamma_0, \gamma_1)\}. \end{aligned}$$

Arguing as in Subsection 2.7.1, we prove σ -compactness of $\text{TGeo}(\mathbf{M})$ in X , hence $\text{TGeo}(\mathbf{M})$ is a Borel subset of X . The quantities $u(\gamma_1)$ and $u(\gamma_0) + l(\gamma_0, \gamma_1)$ depend Borel measurably on $\gamma \in X$; the claim follows. Similarly, we establish that F_5 is a Suslin set. Finally, by definition of \preceq_u we obtain the decomposition

$$F_2 = A_1 \cup [A_2 \cap A_3 \cap A_4],$$

where

$$A_1 := \{(x, \gamma^1, \gamma^2) \in \mathbf{M} \times X^2 : x = \gamma_0^1\},$$

$$A_2 := U \times [e_0^{-1}(U) \cap X] \times X,$$

$$A_3 := \{(x, \gamma^1) \in \mathbf{M} \times X : u(\gamma_0^1) = u(x) + l(x, \gamma_0^1)\} \times X,$$

$$A_4 := \{(x, \gamma^1) \in \mathbf{M} \times X : l(x, \gamma_0) > 0\}.$$

As above, it is then easily seen that these are all Suslin sets, and so is F_2 ; similarly, we verify that F_3 is Suslin.

Now we construct the selection map S . To exclude possible intersections along the curves (including the endpoints), we consider the Suslin set

$$F := F' \cap [\mathbf{M} \times \{(\gamma^1, \gamma^2) \in X^2 : \min_{t \in [0, 1]} d(\gamma_t^1, \gamma_t^2) > 0\}].$$

By Lemma 6.10, for every $x \in B$ we have $F \cap [\{x\} \times X^2] \neq \emptyset$, in fact $B \subset \text{pr}_1(F)$. The inclusion $B \supset \text{pr}_1(F)$ is satisfied as well, which is derived as in the proof of Lemma 6.10. Therefore, Theorem 6.22 provides us with an \mathcal{A} -measurable selection

map $S: B \rightarrow X^2$. By construction, we have $S(B) \subset G^2$, and the other properties are again argued as in the proof of Lemma 6.10. \square

For the proof of Theorem 6.21, we also need the following result about cyclical monotonicity. The proof is completely analogous to the one of [44, Prop. 4.12].

Lemma 6.24 (A criterion for cyclical monotonicity). *Let $0 < p < 1$ be arbitrary, and let $\Delta \subset E_{\leq u}^2$ satisfy*

$$[u(x_1) - u(x_0)] [u(y_1) - u(y_0)] \geq 0$$

for every $(x_0, y_0), (x_1, y_1) \in \Delta$. Then Δ is l^p -cyclically monotone.

Proof of Theorem 6.21. Assume to the contrary that $\mathfrak{m}[B] > 0$. In this case, we construct an L^∞ -measure $\mu_0 \in \mathcal{P}_c(\mathbf{M})$ and a measure $\mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathfrak{m})$ with

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}$$

and a chronological ℓ_p -optimal coupling $\pi \in \Pi(\mu_0, \mu_1)$ which is not induced by a map, which contradicts our assumption.

We start with some preparations. By inner regularity of \mathfrak{m} , we may and will fix a compact set $C' \subset B$ with $0 < \mathfrak{m}[C'] < \infty$. Let S be the map given by Lemma 6.23, and write $\gamma^j(x) := \text{pr}_j \circ S(x)$, where $j = 1, 2$. An iterated application of Lusin's theorem gives the existence of an \mathfrak{m} -measurable set $C'' \subset C'$ with $\mathfrak{m}[C''] > 0$ such that $S|_{C''}$ and, for every $i = 0, 1$ and every $j = 1, 2$, the restriction of the function $x \mapsto u(\gamma^j(x)_i)$ to C'' is continuous.

Given $x \in C''$, set $\alpha_x := u(\gamma^1(x)_0)$ and $\beta_x := u(\gamma^1(x)_1)$; note that $\alpha_x < \beta_x$ since $\text{pr}_1 \circ S(C'') \subset G$. Using that \mathfrak{m} has full support, by a continuity argument we find an \mathfrak{m} -measurable set $C \subset C''$ with $\mathfrak{m}[C] > 0$, henceforth fixed, such that

$$\sup_{x \in C} \alpha_x < \inf_{x \in C} \beta_x.$$

Now fix $c > 0$ with

$$\sup_{x \in C} \alpha_x < c < \inf_{x \in C} \beta_x. \quad (6.7)$$

Since $u(\gamma^1(x)_t) = u(\gamma^2(x)_t)$ for every $x \in C$ and every $0 \leq t \leq 1$ by Lemma 6.23, there exists a continuous map $\mathfrak{t}: C \rightarrow (0, 1)$ such that for every $x \in C$,

$$u(\gamma^1(x)_{\mathfrak{t}(x)}) = u(\gamma^2(x)_{\mathfrak{t}(x)}) = c. \quad (6.8)$$

For $j = 1, 2$, define the continuous map $T^j: C \rightarrow \mathbf{M}$ by

$$T^j(x) := \gamma^j(x)_{\mathfrak{t}(x)}.$$

Finally, define $\pi \in \mathcal{P}(\mathbf{M}^2)$ by

$$2\pi := (\text{Id}, T^1)_\# [\mathfrak{m}[C]^{-1} \mathfrak{m} \llcorner C] + (\text{Id}, T^2)_\# [\mathfrak{m}[C]^{-1} \mathfrak{m} \llcorner C].$$

Observe that $\mu_0 := (\text{pr}_1)_\# \pi \in \mathcal{P}_c^{\text{ac}}(\mathbf{M}, \mathfrak{m})$ and $\mu_1 := (\text{pr}_2)_\# \pi \in \mathcal{P}_c(\mathbf{M})$. The relations (6.7) imply that π splits mass, i.e. cannot be induced by a map. Furthermore, by construction of T^1 and T^2 , cf. Lemma 6.23, π is chronological; in particular, up to restricting π to a suitable product set we may and will assume

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}.$$

To thus arrive at the desired contradiction, we claim ℓ_p -optimality of π . By construction, π is concentrated on the set

$$\Delta := \{(x, \gamma^1(x)_{\mathfrak{t}(x)}) : x \in C\} \cup \{(x, \gamma^2(x)_{\mathfrak{t}(x)}) : x \in C\}.$$

By (6.8), for every $(x_0, y_0), (x_1, y_1) \in \Delta$ we easily get

$$[u(x_1) - u(x_0)] [u(y_1) - u(y_0)] = 0.$$

Lemma 6.24 entails l^p -cyclical monotonicity of Δ ; since π is chronological, its l_p -optimality follows from [44, Prop. 2.8]. \square

Mutatis mutandis, the following Corollary 6.25 is shown.

Corollary 6.25 (Negligibility of timelike backward branching points). *Assume the existence of $0 < p < 1$ with the following property. For every $\mu_0 \in \mathcal{P}_c(\mathbf{M})$ and every L^∞ -measure $\mu_1 \in \mathcal{P}_c(\mathbf{M})$ such that*

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\},$$

every l_p -optimal coupling of μ_0 and μ_1 is induced by a map. Then

$$\mathbf{m}[B^\leftarrow] = 0.$$

Hence, if both assumptions from Theorem 6.21 and Corollary 6.25 hold, then

$$\mathbf{m}[B \cup B^\leftarrow] = 0.$$

In fact, thanks to Lemma 6.24 the transport exponents in the two previous results do not need to coincide in order to get the same conclusion.

6.4.2. *Negligibility of extremal points.* Now we address the \mathbf{m} -negligibility of the sets a and b from Definition 6.9. We essentially follow [44, Sec. 4.4] up to formulating the results in slightly larger generality as we neither assume curvature bounds nor timelike nonbranching, yet the proofs remain similar.

Given a Suslin set $A \subset \mathcal{T}_u^{\text{bp}}$ and $s \geq 0$, we define the Suslin set $A_s \subset \mathcal{T}_u^{\text{bp}}$ by

$$A_s := \text{pr}_2[\{(\alpha, y) \in (A \times \mathcal{T}_u^{\text{bp}}) \cap E_{\leq u}^2 : l(\alpha, y) = s\}]$$

In other words, A_s is the set of points “translated forward” by the l -distance t according to the relation \leq_u . In particular, the set of all $s \geq 0$ such that $\mathbf{m}[A_s] > 0$ is Lebesgue measurable, which allows us to evaluate its \mathcal{L}^1 -measure. Up to bad points, the set A_s is closely related to the ray map \mathbf{g} from Corollary 6.16.

Lemma 6.26 (Size of sets moved along rays). *Assume the existence of $0 < p < 1$ such that for every L^∞ -measure $\mu_0 \in \mathcal{P}_c(\mathbf{M})$ and every $\mu_1 \in \mathcal{P}_c(\mathbf{M})$ with*

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\},$$

every element $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ is induced by a map and the measure $(\mathbf{e}_t)_\# \pi$ is absolutely continuous with respect to \mathbf{m} for every $0 \leq t < 1$. Then for every Suslin set $A \subset \mathcal{T}_u^{\text{bp}} \setminus b$ with $\mathbf{m}[A] > 0$, there exist $r > 0$ and a compact set $C \subset A$ with

$$\bigcup_{s \in [0, r]} C_s \Subset \mathbf{M},$$

and for every $0 \leq s < r$ we have $C_s \subset \mathcal{T}_u^{\text{bp}} \setminus b$ and $\mathbf{m}[C_s] > 0$.

In particular, we obtain

$$\mathcal{L}^1[\{0 \leq s < r : \mathbf{m}[A_s] > 0\}] > 0. \quad (6.9)$$

Before entering the proof, we relate Lemma 6.26 to Theorem 6.21. This becomes relevant in the proof of Theorem 6.28.

Remark 6.27 (About the hypothesis of Lemma 6.26). The assumption of the above lemma is stronger than the one of Theorem 6.21. Indeed, in the notation of both results, by Proposition 2.66 every chronological l_p -optimal coupling of μ_0 and μ_1 can be written as the push-forward of some $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$ by $(\mathbf{e}_0, \mathbf{e}_1)$. But if the latter is induced by a map, so is the former.

Analogously, if the causal reversal \mathcal{M}^\leftarrow satisfies the assumption of Lemma 6.26, then the hypothesis of Corollary 6.25 holds.

Proof of Lemma 6.26. Arguing as for [44, Prop. 4.12] by using Lemma 6.24, for every $r \geq 0$ we deduce the Suslin set

$$\Lambda_r := \{(\alpha, y) \in (A \times \mathcal{T}_u^{\text{bp}}) \cap E_{\leq u}^2 : l(\alpha, y) = r\}$$

is l^p -cyclically monotone. We claim

$$\bigcup_{r>0} \text{pr}_1(\Lambda_r) = A. \quad (6.10)$$

Indeed, the inclusion “ \subset ” is clear since $\text{pr}_1(\Lambda_r) \subset A$ for every $r > 0$. On the other hand, since A does not contain final points, for every $q \in A$ there exist $r > 0$ and $y \in \mathcal{T}_u^{\text{bp}}$ such that $(\alpha, y) \in \Lambda_r$, which establishes “ \supset ”.

Now observe that if $0 \leq r' \leq r$ then $\text{pr}_1(\Lambda_{r'}) \subset \text{pr}_1(\Lambda_r)$ thanks to Lemma 6.4. Together with (6.10), Levi’s monotone convergence theorem, and our assumption $\mathfrak{m}[A] > 0$, we get $\mathfrak{m}[\text{pr}_1(\Lambda_r)] > 0$ for a sufficiently small parameter $r > 0$ which is henceforth fixed. Owing to inner regularity of \mathfrak{m} , let $C \subset \text{pr}_1(\Lambda_r)$ be compact with $\mathfrak{m}[C] > 0$. Theorem 6.22 (von Neumann’s selection theorem) yields the existence of a Suslin measurable map $T: C \rightarrow \mathcal{T}_u^{\text{bp}}$ with $\text{graph } T \subset \Lambda_r$. By Lusin’s theorem and again inner regularity of \mathfrak{m} , up to shrinking C we may and will assume T to be continuous — in particular, $T(C)$ is a compact subset of $\mathcal{T}_u^{\text{bp}}$.

Now consider the measures $\mu_0 \in \mathcal{P}_c^{\text{ac}}(M, \mathfrak{m})$ and $\mu_1 \in \mathcal{P}_c(M)$ defined by

$$\begin{aligned} \mu_0 &:= \mathfrak{m}[C]^{-1} \mathfrak{m} \llcorner C, \\ \mu_1 &:= T_{\#} \mu_0. \end{aligned}$$

By construction, the coupling $\pi := (\text{Id}, T)_{\#} \mu_0 \in \Pi(\mu_0, \mu_1)$ is concentrated on Λ_r . This yields both its l^p -cyclical monotonicity and its chronology; in particular, π is ℓ_p -optimal by [44, Prop. 2.8]. Since every appropriate restriction of π is ℓ_p -optimal as well [44, Lem. 2.10], we may and will assume

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}.$$

The lifting procedure from Proposition 2.66 and our assumption thus yield the existence of $\boldsymbol{\pi} \in \text{OptTGeo}_p(\mu_0, \mu_1)$ with $(\mathbf{e}_t)_{\#} \boldsymbol{\pi}$ \mathfrak{m} -absolutely continuous for every $0 \leq t < 1$. Lemma 2.63 yields $(\mathbf{e}_t)_{\#} \boldsymbol{\pi} \subset J(\text{spt } \mu_0, \text{spt } \mu_1)$ for every $0 \leq t \leq 1$. In addition, arguing as in Remark 6.27 we deduce $\boldsymbol{\pi} = (\mathbf{e}_0, \mathbf{e}_1)_{\#} \boldsymbol{\pi}$. As T moves mass from $\text{spt } \mu_0$ to $\text{spt } \mu_1$ with constant l -length r , Lemma 2.63 again implies $(\mathbf{e}_t)_{\#} \boldsymbol{\pi}$ is concentrated on C_{rt} for every $0 \leq t < 1$; \mathfrak{m} -absolute continuity of $(\mathbf{e}_t)_{\#} \boldsymbol{\pi}$ yields

$$\mathfrak{m}[C_{rt}] > 0.$$

Collecting everything leads to the desired properties of C_s for every $0 \leq s < r$.

The inequality (6.9) follows straightforwardly from the previous inequality and the inclusion $C_s \subset A_s$ for every $0 \leq s < r$. This terminates the proof. \square

Theorem 6.28 (Negligibility of extremal points). *Let \mathcal{M} and its causal reversal \mathcal{M}^{\leftarrow} satisfy the assumption of Lemma 6.26. Then*

$$\mathfrak{m}[a] = \mathfrak{m}[b] = 0.$$

Proof. We only show $\mathfrak{m}[a] = 0$, the proof of the identity $\mathfrak{m}[b] = 0$ is analogous.

By Remark 6.27 and Corollary 6.25, it suffices to prove

$$\mathfrak{m}[a \setminus B^{\leftarrow}] = 0.$$

Assume to the contrary $\mathfrak{m}[a \setminus B^{\leftarrow}] > 0$. Let $C \subset a \setminus B^{\leftarrow}$ be a compact subset as provided by applying Lemma 6.26 to $A := a \setminus B^{\leftarrow}$. Retaining the notation of this lemma, we first claim $C_s \cap C_{s'} = \emptyset$ for every $0 < s < s' < r$. Indeed, otherwise let $y \in C_s \cap C_{s'}$. By definition, there exist $\alpha, \alpha' \in C$ such that $l(\alpha, y) = s$ and $\alpha \preceq_u y$ as well as $l(\alpha', y) = s'$ and $\alpha' \preceq_u y$. As in Case 1 in the proof of Theorem 6.14 —

here we use that $a \setminus B^\leftarrow$ does not contain timelike backward branching points — this leads to a contradiction.

Again by Lemma 6.26, we get $\mathfrak{m}[C_s \setminus B^\leftarrow] = \mathfrak{m}[C_s] > 0$ for uncountably many $0 \leq s < r$. On the other hand, the set $\bigcup_{s \in [0, r)} (C_s \setminus B^\leftarrow)$ has compact closure, yet is intersected by uncountably many mutually disjoint Suslin sets of positive \mathfrak{m} -measure. This contradicts the finiteness of \mathfrak{m} on compact subsets of M . \square

Combining Theorem 6.28 with Remark 6.27, Theorem 6.21, and Corollary 6.25 leads to the following.

Corollary 6.29 (Negligibility of bad points). *Under the same hypotheses as stated in Theorem 6.28, the Suslin set of bad points is \mathfrak{m} -negligible.*

Remark 6.30 (Comparison with [44]). Retain the notation of Example 6.1. In our case, the transport set \mathcal{T}_V constructed in [44, Sec. 4.1] coincides with $T \cup ((B \cup B^\leftarrow) \setminus (a \cup b))$. In particular, our transport set \mathcal{T}_u differs from the previous set \mathcal{T}_V only by an \mathfrak{m} -negligible set by Theorem 6.28 (or [44, Cor. 4.15, Thm. 4.17]) in the framework considered in [44].

6.5. Disintegration of \mathfrak{m} on the transport set. After these preparations, we are in a position to state and prove our disintegration theorem. This is the place where the weight function w from Section 6.1 enters through the proof.

We recall the quotient map $Q: \mathcal{T}_u \rightarrow \mathcal{T}_u$ from Proposition 6.18, the quotient set $Q := Q(\mathcal{T}_u)$, and the transport rays $M_\alpha := Q^{-1}(\alpha)$, $\alpha \in Q$, from Definition 6.19.

Theorem 6.31 (Needle decomposition). *Suppose \mathcal{M} is a timelike p -essentially nonbranching $\text{TCD}_p^e(k, N)$ space. Then there exist*

- a probability measure $\mathfrak{q} \in \mathcal{P}(Q)$ which is mutually absolutely continuous to $Q_\#[\mathfrak{m} \llcorner \mathcal{T}_u]$, and
- a map $\mathfrak{m}: Q \rightarrow \mathfrak{M}(M)$

satisfying the disintegration formula

$$\mathfrak{m} \llcorner \mathcal{T}_u^{\text{bp}} = \mathfrak{m} \llcorner \mathcal{T}_u = \int_Q \mathfrak{m}_\alpha \, d\mathfrak{q}(\alpha)$$

as well as the subsequent properties.

- (i) **Measurability.** For every \mathfrak{m} -measurable set $B \subset E$, the map $\alpha \mapsto \mathfrak{m}_\alpha[B]$ is \mathfrak{q} -measurable.
- (ii) **Strong consistency.** For \mathfrak{q} -a.e. $\alpha \in Q$, the measure \mathfrak{m}_α is concentrated on the transport ray M_α .
- (iii) **Disintegration.** For every \mathfrak{m} -measurable set $B \subset E$ and every \mathfrak{q} -measurable set $A \subset Q$,

$$\mathfrak{m}[B \cap Q^{-1}(A)] = \int_A \mathfrak{m}_\alpha[B] \, d\mathfrak{q}(\alpha).$$

- (iv) **Local finiteness.** For every compact set $C \subset E$,

$$\mathfrak{q}\text{-esssup}_{\alpha \in Q} \mathfrak{m}_\alpha[C] < \infty.$$

Furthermore, given any $\mathfrak{q} \in \mathcal{P}(Q)$ as above the disintegration \mathfrak{m} is \mathfrak{q} -essentially unique, in the sense that for every map $\mathfrak{m}': Q \rightarrow \mathfrak{M}(M)$ with the same properties,

$$\mathfrak{m} = \mathfrak{m}' \quad \mathfrak{q}\text{-a.e.}$$

Proof. Theorem 3.24 combined with Corollary 6.29 yields

$$\mathfrak{m} \llcorner \mathcal{T}_u^{\text{bp}} = \mathfrak{m} \llcorner \mathcal{T}_u.$$

It thus suffices to find the desired disintegration of $\mathfrak{m} \llcorner \mathcal{T}_u$.

Define $\mathbf{n}, \mathbf{q} \in \mathcal{P}(M)$ by $\mathbf{n} := w \mathbf{m} \llcorner \mathcal{T}_u$, where w is the weight function normalizing $\mathbf{m} \llcorner E$ from Section 6.1, and $\mathbf{q} := \mathbf{Q}_{\sharp} \mathbf{n}$. By construction, \mathbf{q} is concentrated on Q . We claim \mathbf{q} is mutually absolutely continuous to $\mathbf{Q}_{\sharp}[\mathbf{m} \llcorner \mathcal{T}_u]$. Indeed, let $A \subset Q$ be a Suslin set with $\mathbf{q}[A] = 0$. This means

$$\int_{Q^{-1}(A)} w \, d\mathbf{m} \llcorner \mathcal{T}_u = 0, \quad (6.11)$$

and therefore $\mathbf{m} \llcorner \mathcal{T}_u[\mathbf{Q}^{-1}(A)] = 0$ since w is positive on E . On the other hand, the identity $\mathbf{m} \llcorner \mathcal{T}_u[\mathbf{Q}^{-1}(A)] = 0$ clearly implies (6.11), hence $\mathbf{q}[A] = 0$.

The disintegration theorem for probability measures [60, Thm. 452I] applied to \mathbf{n} and \mathbf{q} yields the existence of a map $\mathbf{n} : Q \rightarrow \mathcal{P}(M)$ with the following properties.

- For every \mathbf{n} -measurable set $B \subset E$, the map $\alpha \mapsto \mathbf{n}_{\alpha}[B]$ is \mathbf{q} -measurable.
- For \mathbf{q} -a.e. $\alpha \in Q$, the measure \mathbf{n}_{α} is concentrated on the transport ray M_{α} .
- For every \mathbf{n} -measurable set $B \subset E$ and every \mathbf{q} -measurable set $A \subset Q$,

$$\mathbf{n}[B \cap \mathbf{Q}^{-1}(A)] = \int_A \mathbf{n}_{\alpha}[B] \, d\mathbf{q}(\alpha).$$

Moreover, every map $\mathbf{n}' : Q \rightarrow \mathcal{P}(M)$ with these properties satisfies

$$\mathbf{n} = \mathbf{n}' \quad \mathbf{q}\text{-a.e.}$$

It is now straightforward to verify that setting

$$\mathbf{m} := w^{-1} \mathbf{n},$$

leads to all desired properties; see [44, Sec. 4.2] for details. We only comment on local finiteness, which does not follow from the above disintegration theorem. For a compact subset $C \subset M$, since $c := \inf w(C) > 0$ we obtain

$$\mathbf{q}\text{-esssup}_{\alpha \in Q} \mathbf{m}_{\alpha}[C] \leq c^{-1} \mathbf{q}\text{-esssup}_{\alpha \in Q} \mathbf{n}_{\alpha}[C] \leq c^{-1}. \quad \square$$

Remark 6.32 (Relation of $\mathbf{m} \llcorner E$ and $\mathbf{m} \llcorner \mathcal{T}_u$). If needed for applications, the identity $\mathbf{m} \llcorner E = \mathbf{m} \llcorner \mathcal{T}_u$ has to be checked separately. What we only know a priori thanks to the inclusion $\mathcal{T}_u \subset E$ is the inequality “ \geq ”.

Lastly, we get to the additional important regularity property of absolute continuity of the conditional measures from Theorem 6.31. To this aim, recall the ray map $\mathbf{g} : \mathcal{D}(\mathbf{g}) \rightarrow \mathcal{T}_u$ from Corollary 6.16, where $\mathcal{D}(\mathbf{g}) \subset \mathcal{T}_u \times \mathbf{R}$. For every $\alpha \in Q$, the map $\mathbf{g}(\alpha, \cdot) : \mathcal{D}(\mathbf{g})(\alpha) \rightarrow \mathcal{T}_u$, where $\mathcal{D}(\mathbf{g})(\alpha)$ is an open interval in \mathbf{R} containing 0, parametrizes the transport ray $\tilde{\alpha}^u$ passing through $\alpha = \mathbf{g}(\alpha, 0)$. Since $\mathbf{g}(\alpha, \cdot)$ is an order isometry, we obtain

$$\mathbf{g}(\alpha, \cdot)_{\sharp}[\mathcal{L}^1 \llcorner \mathcal{D}(\mathbf{g})(\alpha)] = \mathcal{H}^1 \llcorner \mathbf{g}(\alpha, \mathcal{D}(\mathbf{g})(\alpha))$$

as an equality between measures on M_{α} . To ease the notation, we write

$$\mathcal{L}_{\alpha}^1 := \mathbf{g}(\alpha, \cdot)_{\sharp}[\mathcal{L}^1 \llcorner \mathcal{D}(\mathbf{g})(\alpha)].$$

The proof of the following result is completely analogous to [44, Prop. 4.16] in the special case of Example 6.1 by using Lemma 6.26 and thus omitted.

Proposition 6.33 (Absolute continuity of conditional measures). *Assume \mathcal{M} is a timelike p -essentially nonbranching $\text{TCD}_p^e(k, N)$ space, and let \mathbf{q} be a disintegration of $\mathbf{m} \llcorner \mathcal{T}_u$ as provided by Theorem 6.31. Then for \mathbf{q} -a.e. $\alpha \in Q$, the conditional measure \mathbf{m}_{α} is absolutely continuous with respect to \mathcal{L}_{α}^1 .*

In the following, for $\alpha \in Q$ we write h_{α} for the (\mathcal{L}_{α}^1 -uniquely determined) Radon–Nikodým derivative of \mathbf{m}_{α} with respect to \mathcal{L}_{α}^1 .

Remark 6.34 (Identification). It will be useful to identify the transport ray $M_\alpha \subset M$, where $\alpha \in Q$, with its parametrization $\mathcal{D}(\mathfrak{g})(\alpha) \subset \mathbf{R}$ by the homeomorphism $\mathfrak{g}(\alpha, \cdot)$; recall Corollary 6.16. We employ barred notations for the real counterparts of all objects associated with M_α as follows:

$$\begin{aligned}\bar{M}_\alpha &:= \mathcal{D}(\mathfrak{g})(\alpha), \\ \bar{\mathfrak{m}}_\alpha &:= \mathfrak{g}(\alpha, \cdot)_\#^{-1} \mathfrak{m}_\alpha, \\ \bar{\mathcal{L}}_\alpha^1 &:= \mathcal{L}^1 \llcorner \bar{M}_\alpha, \\ \bar{h}_\alpha &:= h_\alpha \circ \mathfrak{g}(\alpha, \cdot).\end{aligned}\tag{6.12}$$

From Proposition 6.33, we deduce $\bar{\mathfrak{m}}_\alpha$ is \mathcal{L}^1 -absolutely continuous, and

$$\bar{\mathfrak{m}}_\alpha = \bar{h}_\alpha \bar{\mathcal{L}}_\alpha^1.$$

Arguing as for Theorem 6.31 and Proposition 6.33, we deduce the following.

Theorem 6.35 (Smooth needle decomposition). *Let \mathcal{M} form the metric measure spacetime induced by a smooth, globally hyperbolic Finsler spacetime. Moreover, let E and u be given as in Section 6.1. Then the conclusions of Theorem 6.31 and Proposition 6.33 hold.*

This echoes Klartag’s Riemannian result [84, Thm. 3.23] in the absence of curvature bounds; general uniqueness results of chronological ℓ_p -optimal couplings and ℓ_p -geodesics (see [91, Thm. 5.8, Cor. 5.9] for the Lorentzian case and [26, Thm. 4.17, Cor. 4.18] for general Finsler spacetimes) explain why Theorem 6.35 does not require curvature hypotheses.

6.6. Localization of $\text{TCD}_p^e(k, N)$. The crucial property of our needle decomposition is that the $\text{TCD}_p^e(k, N)$ property passes over to \mathfrak{q} -a.e. needle and makes them one-dimensional variable CD metric measure spaces, Remark 6.39. Its proof follows Cavalletti–Mondino’s counterpart for metric measure spaces [40, Thm. 4.2], recently shown by them for constant k for Lorentzian length spaces [46, Thm. 3.2].

For the following discussion, let us recall the identifications (6.12) from Remark 6.34. Given any $\alpha \in Q$ we define $k_\alpha: \bar{M}_\alpha \rightarrow \mathbf{R}$ by

$$k_\alpha := k \circ \mathfrak{g}(\alpha, \cdot).$$

Thanks to Corollary 6.16, k_α is lower semicontinuous. In particular, the induced distortion coefficients $\sigma_{k_\alpha/N}^{(t)}(\theta)$ are defined for every $0 \leq \theta < \text{diam}^l M_\alpha$ and every $0 \leq t \leq 1$ analogously to Definition 3.4.

Given any $K \in \mathbf{R}$, $\theta > 0$, and $0 \leq t \leq 1$, recalling our simplifying assumption $N > 1$ we need to consider the geometrically averaged distortion coefficients

$$\tau_{K,N}^{(t)}(\theta) := t^{1/N} \sigma_{K/(N-1)}^{(t)}(\theta)^{1-1/N}\tag{6.13}$$

further detailed in Definition 7.4. Although (7.3) below implies

$$\sigma_{K/N}^{(t)}(\theta) \leq \tau_{K,N}^{(t)}(\theta)$$

the following lemma [50, Lem. 2.1] “reverses” this inequality locally at the cost of a slightly worse real constant $K' < K$.

Lemma 6.36 (Local comparison of distortion coefficients). *For every $K, K' \in \mathbf{R}$ with $K' < K$, there exists $\theta > 0$ such that for every $0 < \theta' < \theta$ and every $0 \leq t \leq 1$,*

$$\sigma_{K/N}^{(t)}(\theta') \geq \tau_{K',N}^{(t)}(\theta').$$




FIGURE 4. An illustration of the proof of Theorem 6.37.

Theorem 6.37 (Disintegration into $\text{CD}(k, N)$ needles). *Assume \mathcal{M} is a timelike p -essentially nonbranching $\text{TCD}_p^e(k, N)$ space, and let \mathfrak{q} be a disintegration according to Theorem 6.31. Then $\bar{\mathfrak{m}}_\alpha$ is absolutely continuous with respect to the one-dimensional Lebesgue measure for \mathfrak{q} -a.e. $\alpha \in Q$, and its $\bar{\mathcal{L}}_\alpha^1$ -density \bar{h}_α has a nonrelabeled $\bar{\mathcal{L}}_\alpha^1$ -version which is locally Lipschitz continuous and strictly positive on the open real interval \bar{M}_α , continuous on its closure, and satisfies the following inequality for every $t_0, t_1 \in \bar{M}_\alpha$ with $t_0 < t_1$ and every $0 \leq s \leq 1$:*

$$\begin{aligned} \bar{h}_\alpha((1-s)t_0 + st_1)^{1/(N-1)} &\geq \sigma_{k_\alpha^-/(N-1)}^{(1-s)}(t_1 - t_0) \bar{h}_\alpha(t_0)^{1/(N-1)} \\ &\quad + \sigma_{k_\alpha^+/(N-1)}^{(s)}(t_1 - t_0) \bar{h}_\alpha(t_1)^{1/(N-1)}. \end{aligned}$$

Proof. Roughly speaking, we proceed as follows. We cover E with precompact, l -geodesically convex sets. On each of these sets, we will displacement interpolate mass distributions from $u^{-1}([A_0, A_0 + L_0])$ to $u^{-1}([A_1, A_1 + L_1])$ for suitable real numbers A_0, A_1, L_0 , and L_1 along the l -geodesic needles. The $\text{TCD}_p^e(k, N)$ property along this transport reduces to a pathwise inequality for the respective densities by Theorem 3.27, and some clever choices of the involved constants leads to the desired estimate. The latter both globalizes along the entire transport along the needles and throughout the cover of E ; see also Remark 6.39.

We first suppose E is precompact (and understand all objects relative to E). By (3.6), we may and will then first assume k is continuous. In turn, up to covering the set E with small chronological diamonds, given any $\varepsilon > 0$ we may and will thus assume the existence of $K \in \mathbf{R}$ with $K - \varepsilon \leq k \leq K + \varepsilon$ on E .

By Proposition 6.18, Q is the disjoint union of countably many Suslin subsets $(Q_i)_{i \in \mathbf{N}}$ of \mathcal{T}_u such that $Q_i \subset u^{-1}(a_i)$ for every $i \in \mathbf{N}$, where $(a_i)_{i \in \mathbf{N}}$ are rational numbers. Evidently, it suffices to show the claim for \mathfrak{q} -a.e. $\alpha \in Q_i$, where $i \in \mathbf{N}$ is arbitrary. In the sequel, we identify Q_i with $\mathbf{g}(Q_i, 0)$. Since $\mathcal{D}(\mathbf{g})(\alpha)$ is an open convex interval in \mathbf{R} , up to partitioning Q_i into countably many pieces and arguing

separately we assume the existence of $b_0, b_1 \in \mathbf{R}$ with $(b_0, b_1) \subset \bar{M}_\alpha$ for every $\alpha \in Q_i$. To ease notation, let $a_i = 0$, $b_0 < 0 < b_1$, $\mathfrak{q}[Q_i] > 0$, and $Q_i = Q$.

Let $a_0 < A_0 < A_1 < a_1$ and $L_0, L_1 > 0$ satisfy $A_0 + L_0 < A_1$ and $A_1 + L_1 < a_1$. Define $\mu_0, \mu_1 \in \mathcal{P}(M)$ by

$$\begin{aligned}\mu_0 &= \int_Q \mathfrak{g}(\alpha, \cdot)_{\sharp} [L_0^{-1} \mathcal{L}^1 \llcorner [A_0, A_0 + L_0]] \, d\mathfrak{q}(\alpha), \\ \mu_1 &= \int_Q \mathfrak{g}(\alpha, \cdot)_{\sharp} [L_1^{-1} \mathcal{L}^1 \llcorner [A_1, A_1 + L_1]] \, d\mathfrak{q}(\alpha).\end{aligned}$$

By Proposition 6.33, μ_0 and μ_1 are \mathfrak{m} -absolutely continuous, and for \mathfrak{q} -a.e. $\alpha \in Q$ their respective \mathfrak{m} -densities satisfy the formulas

$$\begin{aligned}\rho_0^{-1} \circ \mathfrak{g}(\alpha, \cdot) &= L_0 \bar{h}_\alpha \quad \text{on } [A_0, A_0 + L_0], \\ \rho_1^{-1} \circ \mathfrak{g}(\alpha, \cdot) &= L_1 \bar{h}_\alpha \quad \text{on } [A_1, A_1 + L_1].\end{aligned}\tag{6.14}$$

Our hypothesis on E makes μ_0 and μ_1 compactly supported. By construction and since $\mathfrak{g}(\alpha, \cdot)$ parametrizes l -geodesics passing through $\alpha = \mathfrak{g}(\alpha, \cdot)$,

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l \geq A_1 - A_0 - L_0\} \subset \{l > 0\}.$$

Given any $0 \leq s \leq 1$, we define $\mu_s \in \mathcal{P}(M)$ by

$$\mu_s := \int_Q \mathfrak{g}(\alpha, \cdot)_{\sharp} [L_s^{-1} \mathcal{L}^1 \llcorner [A_s, A_s + L_s]] \, d\mathfrak{q}(\alpha),$$

where

$$\begin{aligned}A_s &:= (1-s)A_0 + sA_1, \\ L_s &:= (1-s)L_0 + sL_1.\end{aligned}$$

As for μ_0 and μ_1 , μ_s has compact support and is \mathfrak{m} -absolutely continuous with \mathfrak{m} -density given for \mathfrak{q} -a.e. $\alpha \in Q$ by the formula

$$\rho_s^{-1} \circ \mathfrak{g}(\alpha, \cdot) = L_s \bar{h}_\alpha \quad \text{on } [A_s, A_s + L_s].\tag{6.15}$$

Using Lemma 6.24, it is not hard to verify $(\mu_s)_{s \in [0,1]}$ is an ℓ_p -geodesic from μ_0 to μ_1 . By Corollary 3.25, it necessarily coincides with the unique displacement ℓ_p -geodesic from μ_0 to μ_1 — which is also the one provided by Theorem 3.27. In particular, by Remark 3.28 we have local Lipschitz continuity of \bar{h}_α for \mathfrak{q} -a.e. $\alpha \in Q$.

From the latter pathwise version of $\text{TCD}_p^e(k, N)$ and our assumption on k on E , (6.14), and (6.15), we thus get the following estimate for \mathfrak{q} -a.e. $\alpha \in Q$. For every $A_0 \leq t_0 \leq A_0 + L_0$ and every $A_1 \leq t_1 \leq A_1 + L_1$ represented by $t_0 = A_0 + \lambda L_0$ and $t_1 = A_1 + \lambda L_1$ for some $0 \leq \lambda \leq 1$, and for every $0 \leq s \leq 1$,

$$\begin{aligned}L_s^{1/N} \bar{h}_\alpha((1-s)t_0 + st_1)^{1/N} &\geq \sigma_{k_\alpha^-/N}^{(1-s)}(t_1 - t_0) L_0^{1/N} \bar{h}_\alpha(t_0)^{1/N} \\ &\quad + \sigma_{k_\alpha^+/N}^{(s)}(t_1 - t_0) L_1^{1/N} \bar{h}_\alpha(t_1)^{1/N} \\ &\geq \sigma_{(K-\varepsilon)/N}^{(1-s)}(t_1 - t_0) L_0^{1/N} \bar{h}_\alpha(t_0)^{1/N} \\ &\quad + \sigma_{(K-\varepsilon)/N}^{(s)}(t_1 - t_0) L_1^{1/N} \bar{h}_\alpha(t_1)^{1/N}.\end{aligned}$$

Setting $s = 1/2$ in the previous inequality yields

$$\begin{aligned}L_{1/2}^{1/N} \bar{h}_\alpha(A_{1/2} + \lambda L_{1/2})^{1/N} \\ \geq \sigma_{(K-\varepsilon)/N}^{(1/2)}(A_1 - A_0 + \lambda|L_1 - L_0|) [L_0^{1/N} \bar{h}_\alpha(A_0 + \lambda L_0)^{1/N} \\ + L_1^{1/N} \bar{h}_\alpha(A_1 + \lambda L_1)^{1/N}].\end{aligned}$$

By the continuity of \bar{h}_α , we can set $\lambda = 0$ and obtain

$$L_{1/2}^{1/N} \bar{h}_\alpha(t_{1/2})^{1/N} \geq \sigma_{(K-\varepsilon)/N}^{(1/2)}(t_1 - t_0) [L_0^{1/N} \bar{h}_\alpha(t_0)^{1/N} + L_1^{1/N} \bar{h}_\alpha(t_1)^{1/N}], \quad (6.16)$$

where $2t_{1/2} := t_0 + t_1$. The \mathfrak{q} -conegligible set satisfying the previous inequality depends on t_0, t_1, L_0 , and L_1 . However, again by continuity of \bar{h}_α for \mathfrak{q} -a.e. $\alpha \in Q$, we deduce the following. There exists a \mathfrak{q} -conegligible subset of Q on which (6.16) holds for every $a_0 < t_0 < t_1 < a_1$ and every $L_0, L_1 > 0$ with $t_0 + L_0 < t_1$ and $t_1 + L_1 < t_1$. From now on, we fix $\alpha \in Q$ which belongs to this set. Without loss of generality, we suppose $\bar{h}_\alpha(t_0) > 0$ or $\bar{h}_\alpha(t_1) > 0$.

By Lemma 6.36, there exists $\theta > 0$ such that if $t_1 - t_0 < \theta$, then

$$L_{1/2}^{1/N} \bar{h}_\alpha(t_{1/2})^{1/N} \geq \tau_{K-2\varepsilon, N}^{(1/2)}(t_1 - t_0) [L_0^{1/N} \bar{h}_\alpha(t_0)^{1/N} + L_1^{1/N} \bar{h}_\alpha(t_1)^{1/N}].$$

By the definition (6.13), this is equivalent to

$$(L_0 + L_1)^{1/N} \bar{h}_\alpha(t_{1/2})^{1/N} \geq \sigma_{(K-2\varepsilon)/(N-1)}^{(1/2)}(t_1 - t_0)^{1-1/N} [L_0^{1/N} \bar{h}_\alpha(t_0)^{1/N} + L_1^{1/N} \bar{h}_\alpha(t_1)^{1/N}].$$

For sufficiently small $L > 0$, the optimal choices

$$\begin{aligned} L_0 &= L \bar{h}_\alpha(t_0)^{1/(N-1)} [\bar{h}_\alpha(t_0)^{1/(N-1)} + \bar{h}_\alpha(t_1)^{1/(N-1)}]^{-1}, \\ L_1 &= L \bar{h}_\alpha(t_1)^{1/(N-1)} [\bar{h}_\alpha(t_0)^{1/(N-1)} + \bar{h}_\alpha(t_1)^{1/(N-1)}]^{-1} \end{aligned}$$

easily lead to

$$\bar{h}_\alpha(t_{1/2})^{1/(N-1)} \geq \sigma_{(K-2\varepsilon)/(N-1)}^{(1/2)}(t_1 - t_0) [\bar{h}_\alpha(t_0)^{1/(N-1)} + \bar{h}_\alpha(t_1)^{1/(N-1)}].$$

Recall this inequality holds if $t_1 - t_0 < \theta$, where $\theta > 0$ depends on K and ε through Lemma 6.36. Yet, this midpoint inequality globalizes and holds along all interior points of the geodesic which connects fixed yet arbitrary endpoints [47, Lem. 5.1, Thm. 5.2]; consequently, for every $a_0 < t_0 < t_1 < a_1$ and every $0 \leq s \leq 1$,

$$\begin{aligned} \bar{h}_\alpha((1-s)t_0 + st_1)^{1/(N-1)} &\geq \sigma_{(K-2\varepsilon)/(N-1)}^{(1-s)}(t_1 - t_0) \bar{h}_\alpha(t_0)^{1/(N-1)} \\ &\quad + \sigma_{(K-2\varepsilon)/(N-1)}^{(s)}(t_1 - t_0) \bar{h}_\alpha(t_1)^{1/(N-1)} \\ &\geq \sigma_{(k_\alpha - 3\varepsilon)^-/(N-1)}^{(1-s)}(t_1 - t_0) \bar{h}_\alpha(t_0)^{1/(N-1)} \\ &\quad + \sigma_{(k_\alpha - 3\varepsilon)^+/(N-1)}^{(s)}(t_1 - t_0) \bar{h}_\alpha(t_1)^{1/(N-1)}. \end{aligned}$$

By the arbitrariness of ε and Lemma 3.6, we conclude

$$\begin{aligned} \bar{h}_\alpha((1-s)t_0 + st_1)^{1/(N-1)} &\geq \sigma_{k_\alpha^-/(N-1)}^{(1-s)}(t_1 - t_0) \bar{h}_\alpha(t_0)^{1/(N-1)} \\ &\quad + \sigma_{k_\alpha^+/(N-1)}^{(s)}(t_1 - t_0) \bar{h}_\alpha(t_1)^{1/(N-1)}. \end{aligned}$$

By Lemma 3.6 again, we remove the continuity hypothesis on k on E .

Finally, we outline how to drop the initial assumption of precompactness of E . Owing to Polishness of τ , we cover E with countably many small chronological diamonds. The intersection of each of these sets with E is precompact and l -geodesically convex. By the above argumentation, the claimed estimate holds on each of these small sets. We then apply the local-to-global property of $\text{CD}(k_\alpha, N)$ densities — after an evident variable adaptation of [39, Lem. A.3], cf. [24] plus Remark 6.39 below — to extend the statement to all of E . \square

Remark 6.38 (The case $N = 1$). In the limit case $N = 1$, one can follow a similar argument as above and show that for \mathfrak{q} -a.e. $\alpha \in Q$, the \mathcal{L}_1^α -density \bar{h}_α is constant. Compare with the proof of [40, Thm. 4.2].

We summarize Theorem 6.37 by saying \mathfrak{q} is a $\text{CD}(k, N)$ *disintegration*. This is justified by the following remark.

Remark 6.39 (Connection to $\text{CD}(k_\alpha, N)$ metric measure spaces). The conclusion of the previous theorem is equivalent to say \bar{h}_α is a $\text{CD}(k_\alpha, N)$ *density* on \bar{M}_α for \mathfrak{q} -a.e. $\alpha \in Q$. In particular, the triple $(\bar{M}_\alpha, |\cdot - \cdot|, \bar{\mathfrak{m}}_\alpha)$ forms a $\text{CD}(k_\alpha, N)$ metric measure space after Ketterer [80, Def. 4.4]. Both statements follow from evident variable adaptations of [39, Def. A.1, Thm. A.2] further discussed in [24].

The argument for the following consequence of Theorem 6.37 is standard. Recalling (6.12), given any $\alpha \in Q$ and any $x_0, x_1 \in \bar{M}_\alpha$ with $x_0 < x_1$, we consider the (α -dependent) functions $\kappa_{x_0, x_1}^\pm : [0, |x_1 - x_0|] \rightarrow \mathbf{R}$ defined by

$$\begin{aligned}\kappa_{x_0, x_1}^+(s |x_1 - x_0|) &:= \bar{k}_\alpha((1-s)x_0 + s x_1), \\ \kappa_{x_0, x_1}^-(s |x_1 - x_0|) &:= \bar{k}_\alpha((1-s)x_1 + s x_0).\end{aligned}\tag{6.17}$$

Corollary 6.40 (Comparison inequality). *Let \mathfrak{q} be a $\text{CD}(k, N)$ disintegration. Let $\alpha \in Q$ such that \bar{h}_α satisfies the conclusion of Theorem 6.37. If $a < t_0 < t_1 < b$ all belong to \bar{M}_α , we have*

$$\left[\frac{\sin_{\kappa_{t_0, b}^-}(b - t_1)}{\sin_{\kappa_{t_0, b}^-}(b - t_0)} \right]^{N-1} \leq \frac{\bar{h}_\alpha(t_1)}{\bar{h}_\alpha(t_0)} \leq \left[\frac{\sin_{\kappa_{a, t_1}^+}(t_1 - a)}{\sin_{\kappa_{a, t_1}^+}(t_0 - a)} \right]^{N-1}.$$

7. APPLICATIONS

Now we apply the localization machinery from the previous chapter to sharpen the geometric inequalities obtained in Chapter 5 in the following setting.

Hypothesis 7.1 (Third standing assumption). *\mathcal{M} is a timelike p -essentially non-branching $\text{TCD}_p^e(k, N)$ metric measure spacetime according to Definition 2.50 with $\text{spt } \mathfrak{m} = \mathcal{M}$.*

We need *not* assume the existence of a weight function w according to Section 6.1 a priori. For an appropriate choice of E and u , such a weight is always provided by default thanks to (the proof of) Theorem 7.2. In particular, the Polish background topology τ is not assumed to be proper.

7.1. Localization by optimal transport. We apply the disintegration technique developed in Chapter 6 to the localization with respect to a mean zero function f (roughly speaking, by localizing with respect to a Kantorovich potential u for the ℓ_1 -optimal transport problem from f_+ to f_-). The additional benefit is that the mean zero property of f passes over to a.e. transport ray; see also [1]. Inter alia, this will be used to deduce the sharp Brunn–Minkowski inequality in Theorem 7.6 from its one-dimensional counterpart. A related result for constant k has recently been shown in [46, Thm. 5.2].

Theorem 7.2 (Localization by mean zero functions). *Assume Hypothesis 7.1. Let a nontrivial function $f \in L^1(\mathcal{M}, \mathfrak{m})$ vanish outside a compact subset of \mathcal{M} , suppose*

$$\int_{\mathcal{M}} f \, \text{d}\mathfrak{m} = 0,$$

and assume the probability measures $\mu_0, \mu_1 \in \mathcal{P}_c^{\text{ac}}(\mathcal{M}, \mathfrak{m})$ defined by

$$\begin{aligned}\mu_0 &:= \left\| f_+ \right\|_{L^1(\mathcal{M}, \mathfrak{m})}^{-1} f_+ \, \mathfrak{m}, \\ \mu_1 &:= \left\| f_- \right\|_{L^1(\mathcal{M}, \mathfrak{m})}^{-1} f_- \, \mathfrak{m}\end{aligned}$$

satisfy the chronology condition

$$\text{spt } \mu_0 \times \text{spt } \mu_1 \subset \{l > 0\}.\tag{7.1}$$

Then there exists a Suslin set $\mathcal{T}_u \subset \mathbf{M}$ such that

- (i) $\mathfrak{m}[\mathcal{T}_u] > 0$,
- (ii) $\mathfrak{m} \llcorner \mathcal{T}_u$ admits a $\text{CD}(k, N)$ disintegration \mathfrak{q} according to Section 6.6,
- (iii) f vanishes identically \mathfrak{m} -a.e. on $\mathcal{T}_u^c := \mathbf{M} \setminus \mathcal{T}_u$, and
- (iv) f has mean zero with respect to \mathfrak{m}_α for \mathfrak{q} -a.e. $\alpha \in Q$, i.e.

$$\int_{\mathbf{M}_\alpha} f \, d\mathfrak{m}_\alpha = 0.$$

Proof. In order to localize we construct $E \subset \mathbf{M}$, a function $u: E \rightarrow \mathbf{R}$, and a weight function $w: \mathcal{T}_u^{\text{bp}} \rightarrow (0, \infty)$ which match the framework of Section 6.1. The hypothesized chronology condition implies the pair (μ_0, μ_1) obeys strong l_+ -Kantorovich duality according to Definition C.1 by Example C.2. The Kantorovich duality formula for ℓ_1 from Theorem C.3 implies the existence of

- a compact, causally convex set $E \subset \mathbf{M}$ containing $\text{spt } \mu_0 \cup \text{spt } \mu_1$, and
- a Borel Kantorovich potential $u: E \rightarrow \mathbf{R}$ for the ℓ_1 -optimal transport from μ_0 to μ_1 satisfying the reverse l -Lipschitz condition (6.1).

Clearly, E is l -geodesically convex. It remains to prove the existence of a weight function according to (6.2). We claim $\mathfrak{m}[\mathcal{T}_u^{\text{bp}}] > 0$ — since $\mathcal{T}_u^{\text{bp}} \subset E$ is precompact, by the Radon property of \mathfrak{m} one can then simply take

$$w := \mathfrak{m}[\mathcal{T}_u^{\text{bp}}]^{-1} 1_{\mathcal{T}_u^{\text{bp}}}.$$

To this aim, fix an ℓ_1 -optimal coupling $\pi \in \Pi(\mu_0, \mu_1)$ of μ_0 and μ_1 , cf. Lemma 2.56. By the chronology condition (7.1) and Theorem C.3, π is concentrated on

$$E_{\preceq_u}^2 := E_{\preceq_u}^2 \setminus \text{diag}(\mathbf{M}^2),$$

where $E_{\preceq_u}^2$ is from (6.4). In particular $\pi[(\mathcal{T}_u^{\text{bp}})^2] = 1$, and therefore

$$\mu_0[\mathcal{T}_u^{\text{bp}}] = \mu_1[\mathcal{T}_u^{\text{bp}}] = 1. \quad (7.2)$$

Since μ_0 is \mathfrak{m} -absolutely continuous, this forces $\mathfrak{m}[\mathcal{T}_u^{\text{bp}}] > 0$.

By Corollary 6.29, we thus obtain $\mathfrak{m}[\mathcal{T}_u] > 0$. This proves (i).

Thus, the localization procedure from Chapter 6 relative to E and u applies and provides us with an associated $\text{CD}(k, N)$ disintegration \mathfrak{q} of $\mathfrak{m} \llcorner \mathcal{T}_u$, yielding (ii).

Next, we show (iii). Assume to the contrary the existence of an \mathfrak{m} -measurable set $B \subset \mathcal{T}_u^c$ with $\mathfrak{m}[B] > 0$ and $f \neq 0$ on B . We may and will assume $\mu_0[B] > 0$; the proof in the case $\mu_1[B] > 0$ is analogous. Moreover, by Corollary 6.29, up to removing an \mathfrak{m} -negligible set we may and will assume $B \subset (\mathcal{T}_u^{\text{bp}})^c$. Yet, by using (7.2) we arrive at the contradiction

$$1 = \mu_0[\mathbf{M}] \geq \mu_0[\mathcal{T}_u^{\text{bp}}] + \mu_0[B] > 1.$$

Finally, we turn to (iv). We first claim

$$\mathbf{Q}_\# \mu_0[A] = \mathbf{Q}_\# \mu_1[A]$$

for every \mathfrak{q} -measurable set $A \subset Q$. Indeed, for a fixed coupling π as in the first paragraph, by using Corollary 6.29 and

$$\mu_0[\mathcal{T}_u] = \mu_1[\mathcal{T}_u] = 1,$$

implied by (7.2) and \mathfrak{m} -absolute continuity of μ_0 and μ_1 , as above we obtain

$$\pi[E_{\preceq_u}^2 \cap \mathcal{T}_u^2] = 1.$$

Moreover, if $x, y \in \mathcal{T}_u$ satisfy $x \preceq_u y$, by Theorem 6.14 they necessarily belong to the same transport ray, symbolically $\mathbf{Q}(x) = \mathbf{Q}(y)$. This easily yields

$$(\mathbf{Q}^{-1}(A) \times E) \cap E_{\preceq_u}^2 \cap \mathcal{T}_u^2 = (E \times \mathbf{Q}^{-1}(A)) \cap E_{\preceq_u}^2 \cap \mathcal{T}_u^2.$$

These observations entail

$$\begin{aligned}\mu_0[\mathbb{Q}^{-1}(A)] &= \pi[(\mathbb{Q}^{-1}(A) \times E) \cap E_{\prec_u}^2 \cap \mathcal{T}_u^2] \\ &= \pi[(E \times \mathbb{Q}^{-1}(A)) \cap E_{\prec_u}^2 \cap \mathcal{T}_u^2] \\ &= \mu_1[\mathbb{Q}^{-1}(A)],\end{aligned}$$

which is the claimed identity. From this, we deduce

$$\begin{aligned}\int_A \int_{M_\alpha} f_+(x) \, d\mathbf{m}_\alpha(x) \, d\mathbf{q}(\alpha) &= \|f_+\|_{L^1(M, \mathbf{m})} \mathbb{Q}_\# \mu_0[A] \\ &= \|f_-\|_{L^1(M, \mathbf{m})} \mathbb{Q}_\# \mu_1[A] \\ &= \int_A \int_{M_\alpha} f_-(x) \, d\mathbf{m}_\alpha(x) \, d\mathbf{q}(\alpha).\end{aligned}$$

The arbitrariness of A terminates the proof of (iv). \square

Remark 7.3 (Smooth localization by mean zero functions). In the smooth setting of Theorem 6.35, the conclusions (i), (iii), and (iv) from Theorem 7.2 remain valid irrespective of curvature assumptions.

7.2. Sharp geometric inequalities. Using Theorem 7.2, we now sharpen various results obtained in Chapter 5. We consider the following variable generalization of the geometrically averaged distortion coefficients from (6.13), cf. [80, Def. 4.1].

Definition 7.4 (Averaged distortion coefficients along l -geodesics). *Given any $\gamma \in \text{TGeo}(M)$, $0 \leq \theta \leq |\dot{\gamma}|$, and $0 \leq t \leq 1$, we set*

$$\tau_{k_\gamma^\pm, N}^{(t)}(\theta) := t^{1/N} \sigma_{k_\gamma^\pm / (N-1)}^{(t)}(\theta)^{1-1/N}.$$

Remark 7.5 (The case $N = 1$). The interpretation of the previous definition in the case $N = 1$ is as follows. If k_γ^\pm is nonpositive on $[0, 1]$, then $\tau_{k_\gamma^\pm, 1}^{(t)}(\theta) = t$ for every $0 \leq \theta \leq |\dot{\gamma}|$ and every $0 \leq t \leq 1$; otherwise $\tau_{k_\gamma^\pm, 1}^{(t)}(\theta) < \infty$ if and only if $\theta = 0$.

By [80, Cor. 4.2], in the context of Definition 7.4 we have

$$\tau_{k_\gamma^\pm, N}^{(t)}(\theta) \geq \sigma_{k_\gamma^\pm / N}^{(t)}(\theta). \quad (7.3)$$

For the next result, we recall the definition

$$A_t := \mathbf{e}_t[G(A_0, A_1)]$$

of t -intermediate points of l -geodesics starting in A_0 and ending in A_1 from (5.4).

Theorem 7.6 (Sharp timelike Brunn–Minkowski inequality I). *Let Hypothesis 7.1 hold. Let $A_0, A_1 \subset M$ be Borel sets with positive and finite \mathbf{m} -measure such that the pair $(\mu_0, \mu_1) \in \mathcal{P}^{\text{ac}}(M, \mathbf{m})^2$ is timelike p -dualizable, where*

$$\begin{aligned}\mu_0 &:= \mathbf{m}[A_0]^{-1} \mathbf{m} \llcorner A_0, \\ \mu_1 &:= \mathbf{m}[A_1]^{-1} \mathbf{m} \llcorner A_1.\end{aligned}$$

Let π denote the only element of $\text{OptTGeo}_p(\mu_0, \mu_1)$. Then for every $0 \leq t \leq 1$,

$$\begin{aligned}\mathbf{m}[A_t]^{1/N} &\geq \mathbf{m}[\text{spt}(\mathbf{e}_t)_\# \pi]^{1/N} \\ &\geq \inf \tau_{k_\gamma^-, N}^{(1-t)}(|\dot{\gamma}|) \mathbf{m}[A_0]^{1/N} + \inf \tau_{k_\gamma^+, N}^{(t)}(|\dot{\gamma}|) \mathbf{m}[A_1]^{1/N},\end{aligned}$$

where both infima are taken over all $\gamma \in \text{spt} \pi$.

The inequality holds a fortiori if both infima are taken over all $\gamma \in G(A_0, A_1)$.

Proof. By inner regularity of \mathfrak{m} , we may and will assume A_0 and A_1 are compact. By replacing M with the compact emerald $J(A_0, A_1)$ and by the Radon property of \mathfrak{m} , up to normalization we may and will assume \mathfrak{m} is a probability measure. Moreover, we will first consider the reinforced chronology condition $A_0 \times A_1 \subset \{l > 0\}$, cf. Example 2.59. Then the function $f \in L^1(M, \mathfrak{m})$ defined by

$$f := \mathfrak{m}[A_0]^{-1} 1_{A_0} - \mathfrak{m}[A_1]^{-1} 1_{A_1}$$

satisfies the hypotheses of Theorem 7.2. Let \mathfrak{q} denote a disintegration of $\mathfrak{m} \llcorner \mathcal{T}_u$ provided by this result; in particular,

$$\mathfrak{m} \llcorner \mathcal{T}_u = \int_Q \mathfrak{m}_\alpha \, d\mathfrak{q}(\alpha).$$

We trivially disintegrate the restriction of \mathfrak{m} to \mathcal{T}_u^c according to

$$\mathfrak{m} \llcorner \mathcal{T}_u^c = \int_{\mathcal{T}_u^c} \delta_\alpha \, d\mathfrak{m}(\alpha).$$

Since $Q \cap \mathcal{T}_u^c = \emptyset$, the definitions $\mathfrak{q}' := \mathfrak{q} + \mathfrak{m} \llcorner \mathcal{T}_u^c$ and $\mathfrak{m}' := 1_Q \mathfrak{m} + 1_{\mathcal{T}_u^c} \delta$ lead to the well-defined disintegration formula

$$\mathfrak{m} = \int_{Q \cup \mathcal{T}_u^c} \mathfrak{m}'_\alpha \, d\mathfrak{q}'(\alpha).$$

Given any $\alpha \in Q$, we use the notations $A_{0,\alpha} := A_0 \cap M_\alpha$ and $A_{1,\alpha} := A_1 \cap M_\alpha$. Let $A_{t,\alpha}$ be the set of t -intermediate points for the one-dimensional nondecreasing (i.e. W_2 -optimal) transport from $\mathfrak{m}[A_{0,\alpha}]^{-1} \mathfrak{m}_\alpha \llcorner A_{0,\alpha}$ to $\mathfrak{m}[A_{1,\alpha}]^{-1} \mathfrak{m}_\alpha \llcorner A_{1,\alpha}$ (regarded as a subset of M_α by order isometry). Note that π is necessarily concentrated on affinely parametrized subsegments of the considered transport rays by Theorem 3.24. As \mathfrak{q} is a $\text{CD}(k, N)$ disintegration and hence obeys the conclusion of Theorem 6.37, by an evident variable adaptation of [41, Lem. 3.2] \mathfrak{q}' -a.e. $\alpha \in Q$ thus satisfies

$$\mathfrak{m}_\alpha[A_{t,\alpha}] \geq \left[\inf \tau_{k_\gamma^-, N}^{(1-t)}(|\dot{\gamma}|) \mathfrak{m}_\alpha[A_{0,\alpha}]^{1/N} + \inf \tau_{k_\gamma^+, N}^{(t)}(|\dot{\gamma}|) \mathfrak{m}_\alpha[A_{1,\alpha}]^{1/N} \right]^N,$$

where both infima are taken over all $\gamma \in \text{spt } \pi$, which makes these coefficients notably independent of α — we agree on the latter interpretation of all appearing infima for the rest of the proof. Finally, observe that the mean zero property of f with respect to \mathfrak{m}_α for \mathfrak{q}' -a.e. $\alpha \in Q$ translates into

$$\frac{\mathfrak{m}_\alpha[A_{0,\alpha}]}{\mathfrak{m}[A_0]} = \frac{\mathfrak{m}_\alpha[A_{1,\alpha}]}{\mathfrak{m}[A_1]}.$$

From this, we deduce

$$\mathfrak{m}_\alpha[A_{t,\alpha}] \geq \frac{\mathfrak{m}_\alpha[A_{0,\alpha}]}{\mathfrak{m}[A_0]} \left[\inf \tau_{k_\gamma^-, N}^{(1-t)}(|\dot{\gamma}|) \mathfrak{m}[A_0]^{1/N} + \inf \tau_{k_\gamma^+, N}^{(t)}(|\dot{\gamma}|) \mathfrak{m}[A_1]^{1/N} \right]^N. \quad (7.4)$$

Next, we claim (7.4) to hold for \mathfrak{q}' -a.e. $\alpha \in \mathcal{T}_u^c$, properly interpreted with \mathfrak{m}_α replaced by δ_α , and $A_{0,\alpha}$, $A_{1,\alpha}$, and $A_{t,\alpha}$ replaced by A_0 , A_1 , and A_t , respectively. This is trivially satisfied if $\mathfrak{q}'[\mathcal{T}_u^c] = \mathfrak{m}[\mathcal{T}_u^c] = 0$, thus we suppose $\mathfrak{m}[\mathcal{T}_u^c] > 0$. Since $\mathfrak{m}[A_0]^{-1} 1_{A_0} + \mathfrak{m}[A_1]^{-1} 1_{A_1} = f = 0$ \mathfrak{m} -a.e. on \mathcal{T}_u^c by Theorem 7.2, we get

$$\mathfrak{m}[\mathcal{T}_u^c \cap (A_0 \triangle A_1)] = 0.$$

Since $A_0 \cap A_1 = \emptyset$ by our chronology assumption and Corollary 2.14,

$$\mathfrak{q}' \llcorner \mathcal{T}_u^c = \mathfrak{m} \llcorner \mathcal{T}_u^c = \mathfrak{m} \llcorner (\mathcal{T}_u^c \cap A_0^c \cap A_1^c).$$

This implies $\delta_\alpha[A_0] = 0$ for \mathfrak{q}' -a.e. $\alpha \in \mathcal{T}_u^c$, which trivially verifies (7.4).

Taking these two observations together, we obtain

$$\mathfrak{m}[\text{spt}(\mathbf{e}_t)_\# \pi] = \int_{Q \cup \mathcal{T}_u^c} \mathfrak{m}'_\alpha[\text{spt}(\mathbf{e}_t)_\# \pi \cap M_\alpha] \, d\mathfrak{q}'(\alpha)$$

$$\begin{aligned}
&\geq \int_{Q \cup \mathcal{J}_u^c} \mathbf{m}'_\alpha[A_{t,\alpha}] \, d\mathbf{q}'(\alpha) \\
&\geq \left[\inf \tau_{k\bar{\gamma}, N}^{(1-t)}(|\dot{\gamma}|) \mathbf{m}[A_0]^{1/N} \right. \\
&\quad \left. + \inf \tau_{k\bar{\gamma}, N}^{(t)}(|\dot{\gamma}|) \mathbf{m}[A_1]^{1/N} \right]^N \int_{Q \cup \mathcal{J}_u^c} \frac{\mathbf{m}'_\alpha[A_{0,\alpha}]}{\mathbf{m}[A_0]} \, d\mathbf{q}'(\alpha) \\
&= \left[\inf \tau_{k\bar{\gamma}, N}^{(1-t)}(|\dot{\gamma}|) \mathbf{m}[A_0]^{1/N} + \inf \tau_{k\bar{\gamma}, N}^{(t)}(|\dot{\gamma}|) \mathbf{m}[A_1]^{1/N} \right]^N.
\end{aligned}$$

This establishes the claim under the hypothesis $A_0 \times A_1 \subset \{l > 0\}$.

If μ_0 and μ_1 are merely timelike p -dualizable, we proceed as outlined in the last part of the proof of Theorem 3.27 by reducing the claim to the previous case by appropriately decomposing $\{l > 0\}$ relative to a timelike p -dualizing coupling of μ_0 and μ_1 . The respective t -intermediate point sets have \mathbf{m} -negligible intersection by Lemma 2.71. The timelike Brunn–Minkowski inequality for A_0 and A_1 then follows by concavity of the function $r \mapsto r^{1/N}$ on $[0, \infty)$. We leave out the details. \square

Since the proofs of Theorem 5.5 and Theorem 5.6 are both conducted by contradicting the timelike Brunn–Minkowski inequality, from Theorem 7.6 we obtain the following improved versions under Hypothesis 7.1. In particular, we point out our simplifying assumptions $N > 1$ and $\text{spt } \mathbf{m} = \mathbf{M}$.

Corollary 7.7 (Sharp timelike Bonnet–Myers inequality). *We have*

$$\text{diam}_{k/(N-1)} \mathbf{M} = \text{diam}^l \mathbf{M}.$$

Corollary 7.8 (Sharp timelike Schneider inequality). *Assume that some $o \in \mathbf{M}$, $\beta > 0$, and $R > 0$ satisfy*

$$k \geq (N-1) \left[\frac{1}{4} + \beta^2 \right] l(\cdot, o)^{-2} \quad \text{on } \{l(\cdot, o) > R\}.$$

Then $\text{diam}^l I^-(o) \leq R e^{\pi/\beta}$.

Analogously, if instead the asymptotic condition

$$k \geq (N-1) \left[\frac{1}{4} + \beta^2 \right] l(o, \cdot)^{-2} \quad \text{on } \{l(o, \cdot) > R\} \tag{7.5}$$

holds, then $\text{diam}^l I^+(o) \leq R e^{\pi/\beta}$.

Example 7.9 (Sharpness of Theorem 7.6 and Corollary 7.8). Here, we outline an example of a smooth metric measure spacetime satisfying (7.5) for $\beta = 0$ with infinite l -diameter. In particular, this shows that the factor β^2 in the hypotheses of Corollary 7.8 cannot be replaced by a nonpositive number.

Fix $N \in \mathbf{N}_{\geq 2}$. Let $I := (0, \infty)$, and define $f: I \rightarrow (0, \infty)$ by $f(r) := \sqrt{r}$. Let h denote the standard Euclidean metric on \mathbf{R}^{N-1} . Consider the warped product $\mathbf{M} := I \times_f \mathbf{R}^{N-1}$, i.e. we endow $I \times \mathbf{R}^{N-1}$ with the Lorentzian metric

$$g := dr^2 - f(r)^2 h = dr^2 - r h.$$

By [13, Thm. 3.66], (\mathbf{M}, g) is a smooth, globally hyperbolic Lorentzian spacetime. Moreover, in any direction $v \in T_x \mathbf{M}$ of the form $v = \partial_r + w$, where w is a smooth vector field on \mathbf{R}^{N-1} , the induced Ricci tensor is computed in a standard way [13, p. 114], yielding

$$\begin{aligned}
\text{Ric}(v, v) &= -(N-1) \frac{f''(r)}{f(r)} g(v, v) + (N-2) \left[|f'(r)|^2 - f(r) f''(r) \right] h(w, w) \\
&\geq \frac{N-1}{4} r^{-2} g(v, v).
\end{aligned}$$

By Theorem 4.1, the induced metric measure spacetime \mathcal{M} obeys $\text{TCD}_p(k, N)$ for a continuous function $k: \mathbb{M} \rightarrow (0, \infty)$ such that for every $o \in \mathbb{M}$, it behaves like $(N-1)l(o, \cdot)^{-2}/4$ for $l(o, \cdot)$ sufficiently large. Thus (7.5) with $\beta = 0$ is saturated for sufficiently large $R > 0$, yet $\text{diam}^l I^+(o) = \infty$; of course, in this example, by the precise asymptotics of the Ricci tensor, (7.5) cannot hold for any $\beta > 0$.

In particular, this example shows the sharpness of Theorem 7.6.

Moreover, using Theorem 7.6 we sharpen Theorem 5.9 under Hypothesis 7.1 as follows. We retain the inherent notation fixed in Section 5.4. The sharpness of the subsequent theorem is seen as in Example 7.9, see [44, Rem. 5.11] for constant k .

Corollary 7.10 (Sharp timelike Bishop–Gromov inequality). *Let $E \subset \mathbb{M}$ be compact and l -star shaped with respect to $x \in \mathbb{M}$. Then for every $0 < r < R < R_x$, we have*

$$\begin{aligned} \frac{\mathbf{v}(r)}{\mathbf{v}(R)} &\geq \frac{\int_0^r \sin_{k_x/(N-1)}^{N-1}(t) dt}{\int_0^R \sin_{k_x/(N-1)}^{N-1}(t) dt}, \\ \frac{\mathbf{s}(r)}{\mathbf{s}(R)} &\geq \frac{\sin_{k_x/(N-1)}^{N-1}(r)}{\sin_{k_x/(N-1)}^{N-1}(R)}. \end{aligned}$$

Corollary 7.11 (Sharp dimension bound). *Retain the assumptions and notation from Corollary 5.11. The same argument for this corollary now yields the following improved bound on geometric dimension (5.17):*

$$\dim^l \mathbb{M} \leq N.$$

This is shown to be sharp by Minkowski spacetime [93, Thm. 4.8].

7.3. Hawking-type singularity theorem and integral curvature bounds. In order to formulate our main result of this section, Theorem 7.14, we introduce some machinery, following [44, Sec. 5.1] and [68, Ch. 2].

7.3.1. Coarea formula and mean curvature bounds. We retain Hypothesis 7.1. Moreover, we assume we are situated in the setting of Example 6.1. This means \mathbb{M} is proper, u is the future distance function τ_V from a given achronal FTC Borel set $V \subset \mathbb{M}$ with $\mathbf{m}[V] = 0$, and $E := I^+(V) \cup V$. Given any $t \geq 0$, we consider the map f_t assigning to any $\alpha \in Q$ the unique point $f_t(\alpha) \in (\text{cl } \mathbb{M}_\alpha \cap \{u = t\}) \setminus b$ provided the latter set is nonempty, symbolically $\alpha \in \mathcal{D}(f_t)$. By definition, f_t is injective and $f_0(\alpha) \in V$ for every $\alpha \in \mathcal{D}(f_0)$. The Borel measurable map $f_0 \circ \mathbb{Q}: \mathcal{T}_u \rightarrow V$ should be interpreted as a “footpoint projection”. We define $\mathbf{s}: V \rightarrow [0, \infty]$ by

$$\mathbf{s}(x) := \sup\{t \geq 0 : \mathbb{Q}(x) \in \mathcal{D}(f_t)\}.$$

Lastly, for every \mathbf{m} -measurable set $A \subset \mathcal{T}_u$ with $\mathbf{m}[A] < \infty$, Theorem 6.31 implies

$$\mathbf{m}[A] = \int_0^\infty \mathfrak{h}_t[A] dt = \int_0^\infty \mathfrak{h}_t[A \cap \{u = t\}] dt, \quad (7.6)$$

where for every $t \geq 0$,

$$\mathfrak{h}_t := (f_t)_\# [\bar{h}_\cdot(t) \mathfrak{q}]$$

can be interpreted as the surface measure of the level set $\{u = t\}$; here \bar{h}_α , where $\alpha \in Q$, is from (6.12). Compare with [44, Prop. 5.1].

The following notion of upper mean curvature bounds for V has been proposed in [44, Def. 5.2] following [81]. It is consistent with the smooth case [44, Rem. 5.4], and is motivated by the classical realization of mean curvature as the first variation of surface area. Analogously, lower mean curvature bounds for V can be synthesized.

Definition 7.12 (Upper mean curvature bounds). *We say V has forward mean curvature bounded from above by $-\beta$, where $\beta \in \mathbf{R}$, if*

- a. \mathfrak{h}_0 is a Radon measure satisfying $(f_0)_\# \mathfrak{q} \ll \mathfrak{h}_0$, and
- b. for every compactly supported, bounded Borel function $\phi: V \rightarrow [0, \infty)$ and every induced normal variation

$$V_{t,\phi} := \{u \leq t\phi \circ f_0 \circ \mathbf{Q}\} \cap \mathcal{T}_u, \quad (7.7)$$

where $t > 0$, we have

$$-\beta \int_V \phi^2 d\mathfrak{h}_0 \geq \liminf_{t \rightarrow 0} \frac{2}{t^2} \left[\mathfrak{m}[V_{t,\phi}] - t \int_V \phi d\mathfrak{h}_0 \right].$$

7.3.2. Segment inequality. Next, given any Borel set $A \subset V$, following [68, Def. 2.9] the *future development* of A up to time $T > 0$ is defined by

$$A^{0,T} := \{f_t \circ \mathbf{Q}(x) \in \mathcal{T}_u : x \in A, 0 \leq t \leq \min\{T, \mathfrak{s}(x)\}\}.$$

In other words, $A^{0,T}$ is the set of all points lying on a transport ray which starts in A and has l -distance at most T to V , including earlier termination of the respective needle. The set $\text{Reach}_L := \{x \in V : \mathfrak{s}(x) \geq L\}$ of points of reach L denotes the set of all $x \in V$ such that the unique transport ray \mathbf{M}_α , where $\alpha \in \mathbf{Q}$, whose closure contains x has length at least $L > 0$.

The next result is a version of the Riemannian segment inequality by Cheeger–Colding [48] (see also [117]), adapted to the synthetic setting from [68, Prop. 3.9]. For a Borel function $\psi: \mathbf{M} \rightarrow [0, \infty)$ and $T > 0$, define $F_{\psi,T}: V \rightarrow [0, \infty)$ by

$$F_{\psi,T}(x) := \int_0^{\min\{T, \mathfrak{s}(x)\}} \psi \circ f_t \circ \mathbf{Q}(x) dt.$$

Additionally, given any $\alpha \in \mathbf{Q}$ and any $x_0 < x_1$ both belonging to $\bar{\mathbf{M}}_\alpha$, we consider the nonnegative (α -dependent) function $\varkappa_{x_0, x_1}^-: [0, |x_1 - x_0|] \rightarrow [0, \infty)$ according to (6.17) yet defined in terms of k_- instead of k , i.e.

$$\varkappa_{x_0, x_1}^-(s|x_1 - x_0|) := (\bar{k}_-)_\alpha((1-s)x_1 + sx_0). \quad (7.8)$$

Proposition 7.13 (Segment inequality). *Retain the previous notation. Given any $\delta > 0$, let $A \subset \text{Reach}_{T+\delta}$ be a Borel set with $0 < \mathfrak{h}_0[A] < \infty$. Then*

$$\mathfrak{h}_0\text{-essinf}_{x \in A} F_{\psi,T}(x) \leq c'^{-1} \mathfrak{h}_0[A]^{-1} \int_{A^{0,T}} \psi d\mathfrak{m},$$

where

$$c' := \mathfrak{q}\text{-essinf}_{\alpha \in \mathbf{Q}(A)} \left[\frac{\sin_{-\varkappa_{0, T+\delta}^-/(N-1)}(\delta)}{\sin_{-\varkappa_{0, T+\delta}^-/(N-1)}(T+\delta)} \right]^{N-1}.$$

In particular, if $k \geq K$ on $I^+(A) \cup A$ for some constant $K < 0$,

$$\mathfrak{h}_0\text{-essinf}_{x \in A} F_{\psi,T}(x) \leq \left[\frac{\sinh(\sqrt{-K/(N-1)}(T+\delta))}{\sinh(\sqrt{-K/(N-1)}\delta)} \right]^{N-1} \mathfrak{h}_0[A]^{-1} \int_{A^{0,T}} \psi d\mathfrak{m}.$$

Proof. The second statement follows from Example 3.2 and the monotonicity property stated in Lemma 3.6. For the first statement, we may and will assume $c' > 0$. For simplicity, we also suppose k is continuous. The lower semicontinuous case can be handled by approximation along \mathfrak{q} -a.e. needle, recalling (3.6).

For \mathfrak{q} -a.e. $\alpha \in \mathbf{Q}(A)$, combining Corollary 6.40 with Lemma 3.6 and the trivial inequality $k \geq -k_-$ on \mathbf{M} implies the following inequality for every $t_0 < t_1 < b$ belonging to $\bar{\mathbf{M}}_\alpha$:

$$\bar{h}_\alpha(t_1) \geq \left[\frac{\sin_{-\varkappa_{t_0, b}^-/(N-1)}(b-t_1)}{\sin_{-\varkappa_{t_0, b}^-/(N-1)}(b-t_0)} \right]^{N-1} \bar{h}_\alpha(t_0)$$

Replacing b by $T + \delta$, t_1 by any number $0 < t < T$, and letting $t_0 \rightarrow 0$ yields

$$\bar{h}_\alpha(t) \geq \left[\frac{\sin_{-\kappa_{0,T+\delta}^-/(N-1)}(T + \delta - t)}{\sin_{-\kappa_{0,T+\delta}^-/(N-1)}(T + \delta)} \right]^{N-1} \bar{h}_\alpha(0) \geq c' \bar{h}_\alpha(0).$$

In the last step, we have used nondecreasingness of $\sin_{-\kappa_{0,T+\delta}^-/(N-1)}$.

The coarea formula (7.6) implies

$$\begin{aligned} \int_{A^{0,T}} \psi \, \mathrm{d}\mathbf{m} &= \int_0^T \int_{f_t \circ \mathbf{Q}(A)} \psi \, \mathrm{d}\mathbf{h}_t \, \mathrm{d}t \\ &= \int_0^T \int_{\mathbf{Q}(A)} \psi \circ f_t(\alpha) \bar{h}_\alpha(t) \, \mathrm{d}\mathbf{q}(\alpha) \, \mathrm{d}t \\ &\geq c' \int_0^T \int_{\mathbf{Q}(A)} \psi \circ f_t(\alpha) \bar{h}_\alpha(0) \, \mathrm{d}\mathbf{q}(\alpha) \, \mathrm{d}t \\ &= c' \int_0^T \int_A \psi \circ f_t \circ \mathbf{Q} \, \mathrm{d}\mathbf{h}_0 \, \mathrm{d}t \\ &\geq c' \mathfrak{h}_0[A] \mathfrak{h}_0\text{-essinf}_{x \in A} F_{\psi,T}(x); \end{aligned}$$

here we have used the identities $\alpha = \mathbf{Q} \circ f_0(\alpha)$ for every $\alpha \in Q$ and $f_0 \circ \mathbf{Q}(A) = A$ in the second last identity, and the last equality follows by Fubini's theorem. \square

7.3.3. Main result. We are ready to state our main result of this section. It is a version of Hawking's singularity theorem [74], assuming an integral curvature bound instead of a "pointwise" inequality. It is proven along similar lines as the genuine synthetic Hawking singularity theorem [44, Thm. 5.6]. For smooth Lorentzian spacetimes, it has been shown in [68, Thm. 4.1] assuming a uniform lower bound on the timelike Ricci curvature; we point out that our result is new even in the weighted smooth case. A past version holds with evident changes.

For $T > 0$, $\eta \geq 0$, and $\alpha \in Q$ we recall the definition (7.8) of the nonpositive (α -dependent) potential $\kappa_{0,T+\eta}^-: [0, T + \eta] \rightarrow [0, \infty)$ defined by

$$\kappa_{0,T+\eta}^-(s(T + \eta)) := (\bar{k}_-)_\alpha((1 - s)(T + \eta)).$$

Theorem 7.14 (Synthetic Hawking-type singularity theorem I). *Let Hypothesis 7.1 hold, and assume \mathbf{M} is proper. Let $V \subset \mathbf{M}$ be an achronal FTC Borel set such that $\mathbf{m}[V] = 0$ and $\mathfrak{h}_0[V] < \infty$. We assume the forward mean curvature of V is bounded from above by $-\beta$, where $\beta > 0$. Let $T > 0$ and $\delta > 0$ obey*

$$\mathfrak{h}_0[V]^{-1} \int_{V^{0,T}} k_- \, \mathrm{d}\mathbf{m} < c \left[\beta - \frac{N-1}{T} \right], \quad (7.9)$$

where

$$c := \mathfrak{q}\text{-essinf}_{\alpha \in \mathbf{Q}(V)} \left[\frac{\sin_{-\kappa_{0,T+\delta}^-/(N-1)}(\delta)}{\sin_{-\kappa_{0,T+\delta}^-/(N-1)}(T + \delta)} \right]^{N-1}.$$

Then we have

$$\mathfrak{h}_0[V \setminus \text{Reach}_{T+\delta}] > 0.$$

In other words, there exists a transport ray starting in V whose maximal domain of definition in $I^+(V) \cup V$ is strictly contained in $[0, T + \delta)$.

It is implicit in the Definition 7.12 of upper mean curvature boundedness that $\mathfrak{h}_0[V] > 0$, so that the statement makes sense. Moreover, if $\mathfrak{h}_0[V] = \infty$, an evident variant of this theorem still holds when replacing V by an appropriate subset V' thereof such that $0 < \mathfrak{h}_0[V'] < \infty$.

Proof of Theorem 7.14. As for Proposition 7.13, we assume continuity of k .

Our hypotheses imply $c > 0$ and $\beta > T^{-1}$. We abbreviate

$$L := c^{-1} \mathfrak{h}_0[V]^{-1} \int_{V^{0,T}} k_- \, d\mathbf{m}. \quad (7.10)$$

We assume to the contrary that

$$\mathfrak{h}_0[V \setminus \text{Reach}_{T+\delta}] = 0.$$

Therefore Proposition 7.13 is applicable to $A := \text{Reach}_{T+\delta} \cap V$, up to removing an \mathfrak{h}_0 -negligible set; moreover, $\mathfrak{h}_0[A] = \mathfrak{h}_0[V]$. Combining this with (7.9),

$$\mathfrak{h}_0\text{-essinf}_{x \in A} F_{k_-,T}(x) \leq c^{-1} \mathfrak{h}_0[A]^{-1} \int_{A^{0,T}} k_- \, d\mathbf{m} \leq L.$$

Let $\varepsilon > 0$ be a small number to be chosen later. Then there is a Borel set $A' \subset A$ with $\mathfrak{h}_0[A'] > 0$ such that

$$\mathfrak{h}_0\text{-esssup}_{x \in A'} F_{k_-,T}(x) \leq L + \varepsilon. \quad (7.11)$$

Let $\phi: V \rightarrow [0, \infty)$ be a nontrivial bounded Borel function with compact support in A' ; by scaling, we may and will henceforth assume $\sup \phi(V) < T$. By definition of A' , \mathfrak{q} -a.e. transport ray \mathbf{M}_α has length at least $T + \delta$. Recalling the definition (7.7) of the normal variation $V_{t,\phi}$ induced by ϕ , for every $0 < t < 1$ we obtain

$$\begin{aligned} \mathfrak{m}[V_{t,\phi}] - t \int_V \phi \, d\mathfrak{h}_0 &= \int_Q \int_0^t \bar{h}_\alpha(s \phi \circ f_0(\alpha)) \phi \circ f_0(\alpha) \, ds \, d\mathfrak{q}(\alpha) \\ &\quad - t \int_Q \phi \circ f_0(\alpha) \bar{h}_\alpha(0) \, d\mathfrak{q}(\alpha) \\ &= \int_Q \int_0^t [\bar{h}_\alpha(s \phi \circ f_0(\alpha)) - \bar{h}_\alpha(0)] \, ds \phi \circ f_0(\alpha) \, d\mathfrak{q}(\alpha). \end{aligned} \quad (7.12)$$

Let the nonnegative function $\varkappa_{0,T}^-: [0, T] \rightarrow [0, \infty)$ be defined as in (7.8). As for Proposition 7.13, we prove that for \mathfrak{q} -a.e. $\alpha \in Q$, every $0 < s < t$ certifies

$$\bar{h}_\alpha(s \phi \circ f_0(\alpha)) \geq \left[\frac{\sin_{-\varkappa_{0,T}^-/(N-1)}(T - s \phi \circ f_0(\alpha))}{\sin_{-\varkappa_{0,T}^-/(N-1)}(T)} \right]^{N-1} \bar{h}_\alpha(0).$$

Consider the twice continuously differentiable function $v: [0, T] \rightarrow \mathbf{R}$ with

$$v(\theta) := \sin_{-\varkappa_{0,T}^-/(N-1)}(\theta).$$

Note that v still depends on $\alpha \in Q$; this will not affect our estimates, since they are performed along \mathfrak{q} -a.e. transport ray. By boundedness of $-k_-$ along \mathfrak{q} -a.e. needle starting at A' , a Taylor expansion thus gives

$$\begin{aligned} \mathfrak{m}[V_{t,\phi}] - t \int_V \phi \, d\mathfrak{h}_0 &\geq \int_Q \int_0^t \left[\left[\frac{v(T - s \phi \circ f_0(\alpha))}{v(T)} \right]^{N-1} - 1 \right] \, ds \bar{h}_\alpha(0) \phi \circ f_0(\alpha) \, d\mathfrak{q}(\alpha) \\ &= \int_Q \int_0^t \left[- (N-1) \frac{v'(T)}{v(T)} s \phi \circ f_0(\alpha) + o(s) \right] \, ds \bar{h}_\alpha(0) \phi \circ f_0(\alpha) \, d\mathfrak{q}(\alpha). \end{aligned}$$

On the other hand, along \mathfrak{q} -a.e. ray starting in A' we have

$$\begin{aligned} (N-1) v'(T) &= N-1 + \int_0^T \varkappa_{0,T}^-(r) v(r) \, dr \\ &\leq N-1 + \left[\int_0^T (\bar{k}_-)_\alpha(T-r) \, dr \right] v(T) \end{aligned}$$

$$\leq N - 1 + (L + \varepsilon)v(T);$$

here we have successively used the ODE solved by v , the nonnegativity of $\varkappa_{0,T}^-$, the nondecreasingness of v , and (7.11). Moreover, Sturm's comparison theorem implies $v(T) \geq T$, cf. e.g. [80, Thm. 3.1]. Combining these observations yields

$$\begin{aligned} \mathfrak{m}[V_{t,\phi}] - t \int_V \phi \, d\mathfrak{h}_0 &\geq - \int_Q \int_0^t \left[\frac{N-1}{v(T)} + L + \varepsilon \right] s \, ds \bar{h}_\alpha(0) \phi^2 \circ f_0(\alpha) \, d\mathfrak{q}(\alpha) \\ &\quad + \int_Q \int_0^t o(s) \, ds \bar{h}_\alpha(0) \phi \circ f_0(\alpha) \, d\mathfrak{q}(\alpha) \\ &\geq - \left[\frac{N-1}{T} + L + \varepsilon \right] \frac{t^2}{2} \int_Q \bar{h}_\alpha(0) \phi^2 \circ f_0(\alpha) \, d\mathfrak{q}(\alpha) \\ &\quad + \int_Q o(t^2) \bar{h}_\alpha(0) \phi \circ f_0(\alpha) \, d\mathfrak{q}(\alpha). \end{aligned}$$

The upper boundedness of the forward mean curvature of V according to Definition 7.12 by $-\beta$ and Fatou's lemma yield

$$-\beta \int_V \phi^2 \, d\mathfrak{h}_0 \geq \liminf_{t \rightarrow 0} \frac{2}{t^2} \left[\mathfrak{m}[V_{t,\phi}] - t \int_V \phi \, d\mathfrak{h}_0 \right] \geq - \left[\frac{N-1}{T} + L + \varepsilon \right] \int_V \phi^2 \, d\mathfrak{h}_0.$$

Canceling the integrals on both sides of this inequality and recalling (7.10) yields

$$c^{-1} \mathfrak{h}_0[V]^{-1} \int_{V_{0,T}} k_- \, d\mathfrak{m} \geq \beta - \frac{N-1}{T} - \varepsilon.$$

Choosing ε small enough contradicts our hypothesis (7.9). \square

Remark 7.15 (The case $N = 1$). If $N = 1$, achronal FTC Borel sets with vanishing \mathfrak{m} -measure cannot have a negative forward upper mean curvature bound. Indeed, assume such a set $V \subset \mathbb{M}$ exists. Since $N = 1$, by Remark 6.38 \bar{h}_α is constant for \mathfrak{q} -a.e. $\alpha \in Q$. On the other hand, letting $t \rightarrow 0$ in (7.12) — with an appropriate function ϕ — and using the upper boundedness of the forward mean curvature of V by a negative number easily leads to a contradiction.

APPENDIX A. VARIABLE TIMELIKE MEASURE CONTRACTION PROPERTY

A.1. Definition and basic properties. Now we discuss a weaker variant of the entropic timelike curvature-dimension condition from Definition 3.13.

Definition A.1 (Variable timelike measure contraction property). *We term \mathcal{M} to satisfy the future entropic timelike measure contraction property $\text{TMCP}^{e,+}(k, N)$ if for every $\mu_0 \in \mathcal{P}_c^{\text{ac}}(\mathbb{M}, \mathfrak{m})$ and every $x_1 \in \bigcap_{x \in \text{spt } \mu_0} I^+(x) \cap \text{spt } \mathfrak{m}$, there exist*

- an $\ell_{1/2}$ -geodesic $(\mu_t)_{t \in [0,1]}$ ⁹ connecting μ_0 to $\mu_1 := \delta_{x_1}$, and
- a plan $\pi \in \text{OptTGeo}_{1/2}(\mu_0, \mu_1)$

such that for every $0 \leq t \leq 1$,

$$\mathcal{U}_N(\mu_t) \geq \sigma_{\frac{k\pi}{N}}^{(1-t)}(\mathfrak{c}\pi) \mathcal{U}_N(\mu_0).$$

Moreover, we say that \mathcal{M} satisfies the past entropic timelike measure contraction property $\text{TMCP}^{e,-}(k, N)$ if its causal reversal \mathcal{M}^\leftarrow obeys $\text{TMCP}^{e,+}(k, N)$.

Lastly, if \mathcal{M} certifies $\text{TMCP}^{e,+}(k, N)$ and $\text{TMCP}^{e,-}(k, N)$ simultaneously, it is termed to have the entropic timelike measure contraction property $\text{TMCP}^e(k, N)$.

⁹By the assumptions on μ_0 and μ_1 as well as Corollary 2.64, we could equivalently start with rough or displacement $\ell_{1/2}$ -geodesics here.

In other words, the past version $\text{TMCP}^{e,-}(k, N)$ asks for convexity properties of \mathcal{U}_N along the transport from a point mass to an \mathfrak{m} -absolutely continuous, compactly supported mass distribution in its “chronological” *future*. Compared to [44], we have decided to include both time directions in our timelike measure contraction property to make it invariant under causal reversal (recall Remark 3.15).

Remark A.2 (Independence of the transport exponent). This remark should be compared to [45, Rem. 2.4]. The two objects $(\mu_t)_{t \in [0,1]}$ and π from Definition A.1 do not depend on the chosen exponent $1/2$. Indeed, for μ_0 to μ_1 as above, for every $0 < p < 1$ we have

$$\text{OptTGeo}_p(\mu_0, \mu_1) = \text{OptTGeo}_{1/2}(\mu_0, \mu_1) \quad (\text{A.1})$$

since $\Pi_{\ll}(\mu_0, \mu_1) = \Pi(\mu_0, \mu_1)$ is a singleton. Furthermore, $(\mu_t)_{t \in [0,1]}$ forms in fact a displacement ℓ_p -geodesic. Indeed, by Corollary 2.64 it is represented by a plan $\pi \in \text{OptTGeo}_{1/2}(\mu_0, \mu_1)$. Given any $0 \leq s < t \leq 1$, from (A.1) we thus get

$$\begin{aligned} \ell_p(\mu_s, \mu_t) &\geq \|l \circ (\mathbf{e}_s, \mathbf{e}_t)\|_{L^p(\text{TGeo}(\mathbf{M}), \pi)} \\ &= (t - s) \|l \circ (\mathbf{e}_0, \mathbf{e}_1)\|_{L^p(\text{TGeo}(\mathbf{M}), \pi)} \\ &= (t - s) \ell_p(\mu_0, \mu_1). \end{aligned}$$

The claim follows as in Remark 5.3.

Several basic statements valid under $\text{wTCD}_p^e(k, N)$ have evident counterparts for $\text{TMCP}^e(k, N)$. In particular, this pertains to

- consistency in k and N (Proposition 3.18),
- behavior under measure perturbations (Proposition 3.20),
- maps and uniqueness for chronological ℓ_p -optimal couplings (Theorem 3.23), and
- maps, uniqueness, and regularity of displacement ℓ_p -geodesics (Theorem 3.24) together with its Corollary 3.25.

Thanks to Remark A.2, these results automatically hold for every $0 < p < 1$. We omit the details, but point out that in some of their proofs, the following implication has been implicitly used. No circular reasoning occurs, though.

Proposition A.3 (Consistency with Definition 3.13). *The condition $\text{wTCD}_p^e(k, N)$ implies $\text{TMCP}^e(k, N)$.*

Proof. Let μ_0 and μ_1 be as in Definition A.1. Since

$$\text{spt } \mu_0 \times \{x_1\} \subset \{l > 0\},$$

for sufficiently large $n \in \mathbf{N}$ Proposition 2.28 implies

$$\text{spt } \mu_0 \times \bar{\mathbf{B}}^d(x_1, 2^{-n}) \subset \{l > 0\}. \quad (\text{A.2})$$

As $x_1 \in \text{spt } \mathfrak{m}$, the uniform distribution $\mu_1^n \in \mathcal{P}^{\text{ac}}(\mathbf{M}, \mathfrak{m})$ of $\mathbf{B}^d(x_1, 2^{-n})$ with respect to \mathfrak{m} is well-defined. Moreover, by Corollary 2.17 and shifting n if necessary we may and will assume it to have compact support. In particular, as everything will take place only on the compact causal diamond between $\text{spt } \mu_0$ and $\bar{\mathbf{B}}^d(x_1, 1/2)$, by (3.6) we may and will assume k to be continuous on this set. The general case will easily follow by Levi’s monotone convergence theorem and Lemma 3.6.

For every $n \in \mathbf{N}$, by (A.2) and Example 2.59, $\text{wTCD}_p^e(k, N)$ endows us with

- an ℓ_p -geodesic $(\mu_t^n)_{t \in [0,1]}$ from $\mu_0^n := \mu_0$ to μ_1^n , and
- a plan $\pi^n \in \text{OptTGeo}_p(\mu_0, \mu_1^n)$

such that for every $0 \leq t \leq 1$,

$$\begin{aligned} \mathcal{U}_N(\mu_t^n) &\geq \sigma_{k_{\pi^n}^-/N}^{(1-t)}(\mathbf{c}_{\pi^n}) \mathcal{U}_N(\mu_0) + \sigma_{k_{\pi^n}^+/N}^{(t)}(\mathbf{c}_{\pi^n}) \mathcal{U}_N(\mu_1^n) \\ &\geq \sigma_{k_{\pi^n}^-/N}^{(1-t)}(\mathbf{c}_{\pi^n}) \mathcal{U}_N(\mu_0). \end{aligned} \quad (\text{A.3})$$

Now we pass to the limit as $n \rightarrow \infty$. Clearly, $(\mu_1^n)_{n \in \mathbf{N}}$ converges narrowly to μ_1 and is uniformly compactly supported according to Definition 2.51. Combining (A.2), Corollary 2.47, and Prokhorov's theorem as in the proof of Proposition 2.66, the sequence $(\pi^n)_{n \in \mathbf{N}}$ has a narrow limit $\pi \in \mathcal{P}(\mathcal{C}([0, 1]; \mathbf{M}))$ concentrated on the set $G_r \subset \text{TGeo}(\mathbf{M})$ from Corollary 2.47 along a nonrelabeled subsequence, where $r := \inf l(\text{spt } \mu_0 \times \{x_1\}) > 0$. Since $(\mathbf{e}_0, \mathbf{e}_1)_{\sharp} \pi$ constitutes the unique (chronological) coupling of μ_0 and μ_1 , we obtain $\pi \in \text{OptTGeo}_p(\mu_0, \mu_1)$. This yields

$$\liminf_{n \rightarrow \infty} \sigma_{k_{\pi^n}^-/N}^{(1-t)}(\mathbf{c}_{\pi^n}) \geq \sigma_{k_{\pi}^-/N}^{(1-t)}(\mathbf{c}_{\pi}) \quad (\text{A.4})$$

for every $0 \leq t \leq 1$ by Lemma 3.11, as well as

$$\begin{aligned} \ell_p(\mu_0, \mu_1) &= \|l \circ (\mathbf{e}_0, \mathbf{e}_1)\|_{L^p(\text{TGeo}(\mathbf{M}), \pi)} \\ &\leq \liminf_{n \rightarrow \infty} \|l \circ (\mathbf{e}_0, \mathbf{e}_1)\|_{L^p(\text{TGeo}(\mathbf{M}), \pi^n)} \\ &= \liminf_{n \rightarrow \infty} \ell_p(\mu_0, \mu_1^n). \end{aligned}$$

The inequality follows from continuity of l_+ and [129, Lem. 4.3]. In turn, the latter inequality combined with Theorem 2.68 entails the existence of a rough ℓ_p -geodesic $(\mu_t)_{t \in [0, 1]}$ from μ_0 to μ_1 such that, up to a further nonrelabeled subsequence, $(\mu_t^n)_{n \in \mathbf{N}}$ converges narrowly to μ_t every rational $0 \leq t \leq 1$. For such t , invoking Lemma 2.49, upper semicontinuity of \mathcal{U}_N , (A.3), and (A.4) yields

$$\mathcal{U}_N(\mu_t) \geq \limsup_{n \rightarrow \infty} \mathcal{U}_N(\mu_t^n) \geq \sigma_{k_{\pi}^-/N}^{(1-t)}(\mathbf{c}_{\pi}) \mathcal{U}_N(\mu_0).$$

But by Corollary 2.64, $(\mu_t)_{t \in [0, 1]}$ constitutes in fact a displacement ℓ_p -geodesic and in particular narrowly continuous. Employing the upper semicontinuity properties of \mathcal{U}_N and time continuity of the involved distortion coefficients, ensured by continuity of k , the previous estimate extends to every $0 \leq t \leq 1$ too. Recalling Remark A.2, the proof is terminated. \square

A.2. Geometric consequences. With appropriate modifications, all geometric inequalities from Chapter 5 hold under our variable timelike measure contraction property. Instead of stating them in their nonsharp form, in the sequel we use localization to directly show their sharp versions. This is analogous to Chapters 6 and 7, whence we only sketch the arguments and leave their details to the reader. Similarly to Hypothesis 7.1, in this section we work in the following setting.

Hypothesis A.4 (Modification of Hypothesis 7.1). *\mathcal{M} is a timelike p -essentially nonbranching $\text{TMCP}^e(k, N)$ metric measure spacetime according to Definition 2.50 with $\text{spt } \mathbf{m} = \mathbf{M}$.*

The key is the following analog of Theorem 6.37. Recall the definitions (6.12) of the barred quantities appearing therein.

Theorem A.5 (Disintegration into $\text{MCP}(k, N)$ needles). *Assume Hypothesis A.4. Let \mathfrak{q} be a disintegration according to Theorem 6.31. Then the conditional measure $\bar{\mathfrak{m}}_{\alpha}$ is absolutely continuous with respect to the one-dimensional Lebesgue measure for \mathfrak{q} -a.e. $\alpha \in Q$, and its $\bar{\mathcal{L}}_{\alpha}^1$ -density \bar{h}_{α} has a nonrelabeled $\bar{\mathcal{L}}_{\alpha}^1$ -version which is locally Lipschitz continuous and strictly positive on the open real interval \bar{M}_{α} , continuous*

on its closure, and satisfies the following inequality for every $t_0, t_1 \in \bar{M}_\alpha$ with $t_0 < t_1$ and every $0 \leq s \leq 1$:

$$\bar{h}_\alpha((1-s)t_0 + st_1)^{1/(N-1)} \geq \sigma_{k_\alpha^-/(N-1)}^{(1-s)}(t_1 - t_0) \bar{h}_\alpha(t_0)^{1/(N-1)}. \quad (\text{A.5})$$

Proof. As Theorem 6.31 assumes the stronger $\text{TCD}_p^e(k, N)$ condition, we first have to prove the existence of a disintegration \mathfrak{q} satisfying its conclusions. However, the curvature assumption only enters its proof only through Corollary 6.29 (negligibility of bad points). In turn, the latter result only assumes the setting of Lemma 6.26, which is in place under Hypothesis A.4 as noted in the previous section.

The proof of the claimed absolute continuity of the conditional measures is then analogous to [44, Prop. 4.16]. The inequality (A.5) is shown along the same lines as for Theorem 6.37; see [44, Thm. 4.18] for the argument for constant k . \square

Following the lines of Theorem 7.6, the subsequent result is then shown. Recall (5.4) for the definition of the set

$$A_t := \mathbf{e}_t[G(A_0, A_1)]$$

of t -intermediate points of l -geodesics from A_0 to A_1 , and Definition 7.4 for the definition of the averaged distortion coefficients appearing below.

Theorem A.6 (Sharp timelike Brunn–Minkowski inequality II). *Let Hypothesis A.4 hold. Let $A_0 \subset M$ be a Borel set with positive and finite \mathbf{m} -measure. Assume $x \ll x_1$ for every x in the closure of A_0 . Define $\mu_0, \mu_1 \in \mathcal{P}(M)$ by*

$$\begin{aligned} \mu_0 &:= \mathbf{m}[A_0]^{-1} \mathbf{m} \llcorner A_0, \\ \mu_1 &:= \delta_{x_1}. \end{aligned}$$

Let π denote the only element of $\text{OptTGeo}_p(\mu_0, \mu_1)$. Then for every $0 \leq t \leq 1$,

$$\mathbf{m}[A_t]^{1/N} \geq \mathbf{m}[\text{spt}(\mathbf{e}_t)_\# \pi]^{1/N} \geq \inf \tau_{k_\gamma^-/N}^{(1-t)}(|\dot{\gamma}|) \mathbf{m}[A_0]^{1/N},$$

where the infimum is taken over all $\gamma \in \text{spt} \pi$.

The inequality holds a fortiori if the infimum is taken over all $\gamma \in G(A_0, A_1)$.

In turn, the conclusions of the

- sharp timelike Bonnet–Myers inequality (Corollary 7.7),
- sharp timelike Schneider inequality (Corollary 7.8),
- sharp timelike Bishop–Gromov inequality (Corollary 7.10), and
- sharp dimension bound (Corollary 7.11)

remain unchanged under the weaker Hypothesis A.4: their proofs all employ the timelike Brunn–Minkowski inequality in its version from Theorem A.6.

Lastly, Theorem 7.14 translates into the following analog. Indeed, the estimate from Corollary 6.40 used in its proof remains valid under (A.5).

Theorem A.7 (Synthetic Hawking-type singularity theorem II). *Let Hypothesis A.4 hold, and assume M is proper. Let $V \subset M$ be an achronal FTC Borel set such that $\mathbf{m}[V] = 0$ and $\mathfrak{h}_0[V] < \infty$. We assume the forward mean curvature of V is bounded from above by $-\beta$, where $\beta > 0$. Let $T > 0$ and $\delta > 0$ obey*

$$\mathfrak{h}_0[V]^{-1} \int_{V^{0,T}} k_- \, \text{d}\mathbf{m} < c \left[\beta - \frac{N-1}{T} \right],$$

where

$$c := \mathfrak{q}\text{-essinf}_{\alpha \in \mathcal{Q}(V)} \left[\frac{\sin_{-\kappa_{0,T+\delta}^-/(N-1)}(\delta)}{\sin_{-\kappa_{0,T+\delta}^-/(N-1)}(T+\delta)} \right]^{N-1}.$$

Then we have

$$\mathfrak{h}_0[V \setminus \text{Reach}_{T+\delta}] > 0.$$

In other words, there exists a transport ray starting in V whose maximal domain of definition in $I^+(V) \cup V$ is strictly contained in $[0, T + \delta)$.

APPENDIX B. LIMIT CURVES; LIFTING CAUSALITY FROM EVENTS TO MEASURES

In this appendix, we prove Theorem 2.18 (enhanced topological properties of M) and the limit curve Theorem 2.32 for general strictly causal curves. Both proofs rely on linking our approach from Chapter 2 to the one of *bounded Lorentzian metric spaces* from Minguzzi–Suhr [98], to which we refer for precise definitions as we will only borrow several results from [98]. Moreover, inspired by [51] we study causality — most notably, global hyperbolicity — for the space $\mathcal{P}(M)$.

In the following, only our first standing Hypothesis 2.8 is needed. The second-countability of τ will not be required.

B.1. From metric spacetimes to bounded Lorentzian metric spaces. The bounded Lorentzian metric spaces we construct are quotients [97, Sec. 1.3] of causal diamonds in M . To do so, fix $a, b \in M$ with $a \ll b$. Then

- $X := J(a, b)$ is a space with relative topology inherited from τ , and
- the restriction $\tau := l_+|_{X^2}$ to X^2 is continuous, finite by Corollary 2.16, has compact superlevel sets, and satisfies the reverse triangle inequality (2.2).

Hence, we are in the setting of [98, Sec. 1.3]. Following their procedure, we define an equivalence relation \sim on X by $x \sim y$ if $\tau(x, z) = \tau(y, z)$ and $\tau(z, x) = \tau(z, y)$ for every $z \in X$. Let \tilde{X} denote the quotient space with quotient topology $\tilde{\tau}$, and let $\pi: X \rightarrow \tilde{X}$ denote the canonical projection given by $\pi(x) := \tilde{x}$. The function $\tilde{\tau}: \tilde{X}^2 \rightarrow [0, \infty)$ given by $\tilde{\tau}(\tilde{x}, \tilde{y}) := \tau(x, y)$ is well-defined and continuous.

The first part of the following lemma is then easy to deduce. The crucial point for us lies in the second statement.

Lemma B.1 (Quotienting out undistinguished points). *The tuple $(\tilde{X}, \tilde{\tau})$ constitutes a bounded Lorentzian metric space in the sense of [98, Def. 1.1]. Moreover, the restriction of π to $I(a, b)$ is a homeomorphism onto $\pi(I(a, b))$.*

Proof. The first statement follows from [98, Prop. 1.18].

We now show the second claim. Clearly, π is surjective and continuous on X , hence its restriction to $I(a, b)$ is.

We now claim injectivity of π on $I(a, b)$. Let $x, y \in I(a, b)$ with $\tilde{x} = \tilde{y}$. Owing to the nontrivial chronology in Definition 2.7, let $\gamma: [0, 1] \rightarrow M$ be a timelike curve with midpoint x . Let $z \in M$ [sic] obey $x \ll z$. Since $I(a, b)$ is open and the curve γ is continuous, we have $\gamma_t \in X$ for $0 \leq t \leq 1$ sufficiently close to $1/2$. Up to moving t closer to $1/2$, by continuity of l_+ we have $l(\gamma_t, z) > 0$. Then the reverse triangle inequality (2.2) and our assumption $\tilde{x} = \tilde{y}$ imply

$$l(y, z) \geq l(y, \gamma_t) + l(\gamma_t, z) = l(x, \gamma_t) + l(\gamma_t, z) > 0,$$

and thus $y \ll z$. Reversing the roles of x and y and working with a timelike curve passing through y thus leads to $I^+(x) = I^+(y)$, where both chronological futures are understood globally with respect to l . Corollary 2.20 then implies $x = y$, hence the claimed injectivity of the restriction of π to $I(a, b)$.

Finally, as X is compact, \tilde{X} is Hausdorff as a Polish space, and π is continuous, the latter is a closed map. Thus the restriction of π to $\pi^{-1}(\pi(I(a, b))) = I(a, b)$ is relatively closed, which implies homeomorphy. \square

B.2. Proof of Theorem 2.18. Before entering the proof of the mentioned result, we restate it for convenience.

Theorem B.2 (Polish compact sets). *Hypothesis 2.8 implies the relative topology of τ on every compact subset $C \subset M$ is Polish.*

Proof. We first claim that the relative topology on $I(a, b)$ is metrizable for every $a, b \in M$. Let $(\tilde{X}, \tilde{\tau})$ as in Lemma B.1. By [98, Thm. 1.9], the latter is a Polish space, which immediately implies metrizability of I by Lemma B.1 (e.g. by taking the pullback metric of π under a complete metric \tilde{d} inducing the target topology).

Nontrivial chronology shows M is locally metrizable — every point in M has a neighborhood which is metrizable with respect to the subspace topology. Since the latter is Hausdorff by Corollary 2.17, and C is clearly paracompact, Smirnov’s metrization theorem [53, Thm. 4.4.19] implies metrizability of the relative topology of τ on the set C .

Now the argument is as for [98, Thm. 1.9]. Since C is compact, it is Lindelöf. Together with metrizability, this implies second-countability of C . As every second countable space which is also Hausdorff and locally compact is Polish, the claim follows from Corollary 2.17 again. \square

As for Corollary 2.16, we stress that the previous proof only requires l_+ to be continuous near the diagonal of M^2 .

B.3. Proof of Theorem 2.32. Now we turn to our general limit curve theorem. For convenience, we restate it here as well.

Theorem B.3 (Limit curve theorem). *Assume Hypothesis 2.8. Let $C_0, C_1 \subset M$ be compact subsets, and let $(\gamma^n)_{n \in \mathbf{N}}$ be a sequence of strictly causal curves on $[0, 1]$ such that $\gamma_0^n \in C_0$ and $\gamma_1^n \in C_1$ for every $n \in \mathbf{N}$. Then $(\gamma^n)_{n \in \mathbf{N}}$ has a subsequence which — after reparametrization — converges uniformly, in the sense of Definition 2.30,*

- (i) *to a strictly causal curve defined on a compact interval $[\alpha, \beta] \subset \mathbf{R}$ of positive length unless the endpoints of the extracted subsequence converge to the same limit, and*
- (ii) *to a constant causal curve defined on $[0, 1]$ otherwise.*

Proof. By compactness of C_0 and C_1 , the sequences $(\gamma_0^n)_{n \in \mathbf{N}}$ and $(\gamma_1^n)_{n \in \mathbf{N}}$ have limits $a'' \in C_0$ and $b'' \in C_1$ along subsequences that we do not relabel. By nontrivial chronology in Definition 2.7, let $a, a', b, b' \in M$ with $a \ll a' \ll a''$ and $b'' \ll b' \ll b$. Recall that a, a' , and a'' are mutually distinct by Corollary 2.14, and the same holds for b, b' , and b'' . Furthermore, by Lemma 2.19, global hyperbolicity, and continuity of l_+ , $J(a', b')$ is compactly contained in $I(a, b)$. Lastly, up to removing finitely many members of either sequence, we assume $\gamma_0^n \in I^+(a')$ and $\gamma_1^n \in I^-(b')$ for every $n \in \mathbf{N}$; in other words, by causality and Lemma 2.3 the image of γ^n is entirely contained in $I(a', b')$ for every $a, b \in M$. See Figure 5.

We now consider the bounded Lorentzian metric space $(\tilde{X}, \tilde{\tau})$ from Lemma B.1 constructed from $X := J(a, b)$. We retain all inherent notation from above, and we will show uniform convergence with respect to a metric d as in Remark 2.31. Without restriction, to make \tilde{X} compact we add its spacelike boundary to it [98, Rem. 1.2, Cor. 1.6, Prop. 1.7]. Lastly, by [98, Thm. 5.7], \tilde{X} admits a bounded and uniformly continuous time function $t: \tilde{X} \rightarrow \mathbf{R}$ in the terminology of [98, Def. 5.6].

By using the reverse triangle inequality (2.2) and injectivity of π implied by Lemma B.1, it is easy to see that $\tilde{\gamma}^n: [0, 1] \rightarrow \tilde{X}$ is isocausal in the sense of [98, Def. 5.1, Sec. 5.2]¹⁰ for every $n \in \mathbf{N}$, where

$$\tilde{\gamma}^n := \pi \circ \gamma^n;$$

¹⁰This means strictly causal in our terminology.




FIGURE 5. An illustration of the proof of Theorem 2.32.

since $\tilde{\gamma}_{[0,1]}^n \subset I(a, b)$, none of these curves passes through the spacelike boundary.

First assume the endpoint limits a'' and b'' extracted above to be distinct. As $\tilde{\tau}([a''], [b'']) = \tau(a'', b'') = l(a'', b'') \geq 0$ by definition of the quotient time separation $\tilde{\tau}$, we obtain $\mathfrak{t}([b'']) > \mathfrak{t}([a''])$, and by continuity of \mathfrak{t} we may and will assume the difference $\mathfrak{t}(\tilde{\gamma}_1^n) - \mathfrak{t}(\tilde{\gamma}_0^n)$ to be bounded away from zero uniformly in $n \in \mathbf{N}$. Thanks to [98, Prop. 5.8], the curve $\tilde{\gamma}^n$ can be reparametrized with respect to \mathfrak{t} for every $n \in \mathbf{N}$. More precisely, an explicit possibility to do so is to consider the reparametrization $\tilde{\sigma}^n: [\alpha_n, \beta_n] \rightarrow \tilde{X}$ defined by

$$\begin{aligned}\tilde{\sigma}^n &:= \tilde{\gamma}^n \circ (\mathfrak{t} \circ \tilde{\gamma}^n)^{-1}, \\ \alpha_n &:= \mathfrak{t}(\tilde{\gamma}_0^n), \\ \beta_n &:= \mathfrak{t}(\tilde{\gamma}_1^n).\end{aligned}$$

Then $(\alpha_n)_{n \in \mathbf{N}}$ and $(\beta_n)_{n \in \mathbf{N}}$ have no common accumulation point. Therefore, the limit curve theorem [98, Thm. 5.12] entails uniform convergence of $(\tilde{\sigma}^n)_{n \in \mathbf{N}}$ to a strictly causal curve $\tilde{\sigma}: [\alpha, \beta] \rightarrow \tilde{X}$, where $\alpha, \beta \in \mathbf{R}$ with $\alpha < \beta$; more precisely,

$$\begin{aligned}\alpha &:= \mathfrak{t}(a''), \\ \beta &:= \mathfrak{t}(b'').\end{aligned}$$

Since $J(a', b')$ is compactly contained in $I(a, b)$, by Lemma B.1 the map π is a homeomorphism of $J(a', b')$ onto its compact image $\pi(J(a', b'))$. In particular, the inverse of this map is uniformly continuous. As such maps preserve uniform convergence, we get uniform convergence of $(\pi^{-1} \circ \tilde{\sigma}^n)_{n \in \mathbf{N}}$ to $\pi^{-1} \circ \tilde{\sigma}$. The former

are reparametrizations of $(\gamma^n)_{n \in \mathbf{N}}$, and we obtain our claim. Note again that the limit curve (or any reparametrization thereof) thus obtained is strictly causal.

Lastly, we investigate the case $a'' = b''$. By Proposition 2.28, there exists a neighborhood $V \subset \mathbf{B}^d(a'', \varepsilon)$ such that every strictly causal curve whose endpoints lie in V does not leave $\mathbf{B}^d(a'', \varepsilon)$. For every sufficiently large $n \in \mathbf{N}$, our hypothesis yields $\gamma_0^n, \gamma_1^n \in V$. Consequently,

$$\limsup_{n \rightarrow \infty} \sup_{t \in [0,1]} d(a'', \gamma_t^n) \leq \varepsilon.$$

Since ε was arbitrary, the claim is established. \square

B.4. Causality for probability measures. In [51] a natural causal structure of the space $\mathcal{P}(\mathbf{M})$ was introduced and studied for smooth spacetimes \mathbf{M} . As outlined here, the definition easily extends to our nonsmooth setting. In fact, we succeed in showing global hyperbolicity of \mathbf{M} (according to Theorem 2.12) “lifts” to $\mathcal{P}(\mathbf{M})$.

Define a relation \preceq on $\mathcal{P}(\mathbf{M})$ by $\mu \preceq \nu$ provided $\Pi_{\preceq}(\mu, \nu) \neq \emptyset$. This can be restated by saying $\ell_p(\mu, \nu) \geq 0$ for some — equivalently, every — $0 < p \leq 1$. As for the causal relation \leq on \mathbf{M} , given $\mathcal{X}, \mathcal{Y} \subset \mathcal{P}(\mathbf{M})$ we then define

- the *causal future* of \mathcal{X} by $\mathcal{J}^+(\mathcal{X}) := \{\ell_p(\mu, \cdot) \geq 0\}$,
- the *causal past* of \mathcal{Y} by $\mathcal{J}^-(\mathcal{Y}) := \{\ell_p(\cdot, \nu) \geq 0\}$, and
- the *causal emerald* between \mathcal{X} and \mathcal{Y} by $\mathcal{J}(\mathcal{X}, \mathcal{Y}) := \mathcal{J}^+(\mathcal{X}) \cap \mathcal{J}^-(\mathcal{Y})$.

Proposition B.4 (Order property of \preceq). *The relation \preceq is an order, i.e. it is reflexive, antisymmetric, and transitive.*

Proof. Transitivity follows from the gluing lemma. Nonemptiness of $\Pi_{\preceq}(\mu, \mu)$ for every $\mu \in \mathcal{P}(\mathbf{M})$ follows from causality of the diagonal coupling of μ . In order to prove antisymmetry, we first claim this coupling is the *only* element in $\Pi_{\preceq}(\mu, \mu)$. This, however, follows from [51, Thm. 12] whose proof extends with no changes to our setting by using the generalized time function from Theorem 2.21.

Now let $\mu, \nu \in \mathcal{P}(\mathbf{M})$ obey $\mu \preceq \nu \preceq \mu$. By the previous part, gluing of elements of $\Pi_{\preceq}(\mu, \nu)$ and $\Pi_{\preceq}(\nu, \mu)$ induces a measure $\omega \in \mathcal{P}(\mathbf{M}^3)$ with marginals μ, ν , and μ . By the previous paragraph, $(\text{pr}_1, \text{pr}_3)_\# \omega$ is the diagonal coupling of μ , and hence ω -a.e. $(x, y, z) \in \mathbf{M}^3$ satisfies $x \leq y \leq z = x$, thus $x = y$ by causality. Consequently, $(\text{pr}_1, \text{pr}_2)_\# \omega$ is concentrated on the diagonal of \mathbf{M}^2 , and hence $\mu = \nu$. \square

A nice consequence of independent interest is the following.

Theorem B.5 (Global hyperbolicity). *Hypothesis 2.8 implies the space $\mathcal{P}(\mathbf{M})$ is globally hyperbolic in the following sense.*

- (i) *The relation \preceq is a narrowly closed order.*
- (ii) *If $\mathcal{C} \subset \mathcal{P}(\mathbf{M})$ is narrowly compact, so is $\mathcal{J}(\mathcal{C}, \mathcal{C})$.*

Proof. To show (i), by Proposition B.4 it suffices to prove narrow closedness. For sequences $(\mu_n)_{n \in \mathbf{N}}$ and $(\nu_n)_{n \in \mathbf{N}}$ in $\mathcal{P}(\mathbf{M})$ satisfying $\mu_n \preceq \nu_n$ for every $n \in \mathbf{N}$ and converging to $\mu, \nu \in \mathcal{P}(\mathbf{M})$, respectively, let $(\pi_n)_{n \in \mathbf{N}}$ a sequence of elements $\pi_n \in \Pi_{\preceq}(\mu_n, \nu_n)$. Since the marginal sequences are tight, so is $(\pi_n)_{n \in \mathbf{N}}$, and a nonrelabeled subsequence thus converges narrowly to some $\pi \in \Pi(\mu, \nu)$ by Prokhorov (Theorem 2.54). Alexandrov (Theorem 2.53) shows π is causal.

We turn to the proof of (ii). We first argue that if $\mu, \nu \in \mathcal{P}(\mathbf{M})$ satisfy $\mu \preceq \nu$, then every closed set $C \subset \mathbf{M}$ satisfies

$$\begin{aligned} \mu[J^+(C)] &\leq \nu[J^+(C)], \\ \nu[J^-(C)] &\leq \mu[J^-(C)]; \end{aligned} \tag{B.1}$$

compare with [51, Thm. 8]. This statement makes sense as $J^\pm(C)$ is closed, hence Borel, by closedness of C . Now, by the push-up property the indicator function

$1_{J^+(C)}$ is a causal function on M , in the sense that if $x, y \in M$ satisfy $x \leq y$, then $1_{J^+(C)}(x) \leq 1_{J^+(C)}(y)$. Given $\pi \in \Pi_{\leq}(\mu, \nu)$, this easily yields

$$\mu[J^+(C)] = \int_{M^2} 1_{J^+(C)} \circ \text{pr}_1 \, d\pi \leq \int_{M^2} 1_{J^+(C)} \circ \text{pr}_2 \, d\pi = \nu[J^+(C)].$$

The second inequality in (B.1) follows analogously. Finally, let \mathcal{C} as hypothesized. Given $\varepsilon > 0$, Prokhorov's theorem implies the existence of a compact set $C \subset M$ such that every $\mu \in \mathcal{C}$ satisfies

$$\mu[C^c] \leq \varepsilon.$$

Given any $\sigma \in \mathcal{J}(\mathcal{C}, \mathcal{C})$ and $\mu, \nu \in \mathcal{C}$ with $\mu \preceq \sigma \preceq \nu$, (B.1) then implies

$$\begin{aligned} \sigma[J(C, C)^c] &\leq \sigma[J^+(C)^c] + \sigma[J^-(C)^c] \\ &= 2 - \sigma[J^+(C)] - \sigma[J^-(C)] \\ &\leq 2 - \mu[J^+(C)] - \nu[J^-(C)] \\ &\leq 2 - \mu[C] - \nu[C] \\ &\leq 2\varepsilon. \end{aligned}$$

Together with closedness of \preceq , this shows the desired compactness by Prokhorov's theorem. The proof is terminated. \square

Note that in general, the conclusion of Theorem B.5 “projects down” to M , in the sense that it implies \leq is a closed order, and the induced causal emeralds are compact. This can be seen by considering Dirac masses, respectively.

Similarly to Lemma 2.49, we can now shortly discuss (pre)compactness of intermediate point sets of measures. The corresponding Lemma B.6 is only used in the proof of Theorem 3.26, but might be of independent interest.

For $\mu, \nu \in \mathcal{P}(M)$ and $0 \leq t \leq 1$, we define the (possibly empty, p -dependent) t -intermediate point set

$$\mathcal{Z}_t^p(\mu, \nu) := \{\sigma \in \mathcal{J}(\mu, \nu) : \ell_p(\mu, \sigma) = t \ell_p(\mu, \nu), \ell_p(\sigma, \nu) = (1-t) \ell_p(\mu, \nu)\}.$$

Moreover, for $\mathcal{C} \subset \mathcal{P}(M)^2$ we define

$$\begin{aligned} \mathcal{Z}_t^p(\mathcal{C}) &:= \bigcup_{(\mu, \nu) \in \mathcal{C}} \mathcal{Z}_t^p(\mu, \nu), \\ \mathcal{Z}^p(\mathcal{C}) &:= \bigcup_{t \in [0, 1]} \mathcal{Z}_t^p(\mathcal{C}). \end{aligned}$$

The following result is not straightforward due to the anticipated lack of semi-continuity of ℓ_p [44, Rem. 2.17]. Still, the identities defining intermediate point sets and the reverse triangle inequality (2.14) for ℓ_p guarantee enough rigidity.

Lemma B.6 (Intermediate point sets inherit compactness II). *Let $0 \leq t \leq 1$. If $\mathcal{C} \subset \mathcal{P}(M)^2$ is precompact, so are $\mathcal{Z}_t^p(\mathcal{C})$ and $\mathcal{Z}^p(\mathcal{C})$. Moreover, if $\mathcal{C} = \{(\mu, \nu)\}$ for $\mu, \nu \in \mathcal{P}_c(M)$, then $\mathcal{Z}_t^p(\mathcal{C})$ and $\mathcal{Z}^p(\mathcal{C})$ are compact.*

Proof. The first claim follows from Theorem B.5 as for Lemma 2.49.

The last statement is trivial if $\mu \not\preceq \nu$, hence we assume $\ell_p(\mu, \nu) \geq 0$. It suffices to show $\mathcal{Z}^p(\mu, \nu)$ is narrowly closed under the given assumption $\mathcal{C} = \{(\mu, \nu)\}$. Then $C = J(\text{spt } \mu, \text{spt } \nu)$ is compact. Let $(\sigma_n)_{n \in \mathbf{N}}$ be a sequence in $\mathcal{Z}^p(\mathcal{C})$ narrowly converging to $\sigma \in \mathcal{P}(M)$. We claim $\sigma \in \mathcal{Z}^p(\mathcal{C})$. Indeed, let $(\mu_n, \nu_n)_{n \in \mathbf{N}}$ be a sequence in $\mathcal{Z}(\mathcal{C})$ and let $(t_n)_{n \in \mathbf{N}}$ be a sequence in $[0, 1]$ — without loss of generality converging to $0 \leq t \leq 1$ — such that, for every $n \in \mathbf{N}$,

$$\begin{aligned} \ell_p(\mu, \sigma_n) &= t_n \ell_p(\mu, \nu), \\ \ell_p(\sigma_n, \nu) &= (1-t_n) \ell_p(\mu, \nu). \end{aligned}$$

By [91, Prop. 2.9], we have $\text{spt } \sigma_n \subset J(\text{spt } \mu, \text{spt } \nu)$ for every $n \in \mathbf{N}$. Lemma 2.56 implies the existence of ℓ_p -optimal couplings $\pi_1^n \in \Pi_{\leq}(\mu, \sigma_n)$ and $\pi_2^n \in \Pi_{\leq}(\sigma_n, \nu)$. By Prokhorov's theorem, Alexandrov's theorem, and [129, Lem. 4.3] there exist couplings $\pi_1 \in \Pi_{\leq}(\mu, \sigma)$ and $\pi_2 \in \Pi_{\leq}(\sigma, \nu)$ such that

$$\begin{aligned} \ell_p(\mu, \nu) &\geq \ell_p(\mu, \sigma) + \ell_p(\sigma, \nu) \\ &\geq \|l\|_{L^p(\mathbf{M}^2, \pi_1)} + \|l\|_{L^p(\mathbf{M}^2, \pi_2)} \\ &\geq \limsup_{n \rightarrow \infty} \|l\|_{L^p(\mathbf{M}^2, \pi_1^n)} + \limsup_{n \rightarrow \infty} \|l\|_{L^p(\mathbf{M}^2, \pi_2^n)} \\ &= \limsup_{n \rightarrow \infty} t_n \ell_p(\mu, \nu) + \limsup_{n \rightarrow \infty} (1 - t_n) \ell_p(\mu, \nu) \\ &= \ell_p(\mu, \nu). \end{aligned}$$

This forces the implicit inequalities

$$\begin{aligned} \ell_p(\mu, \sigma) &\geq t \ell_p(\mu, \nu), \\ \ell_p(\sigma, \nu) &\geq (1 - t) \ell_p(\mu, \nu) \end{aligned}$$

to be equalities. The desired inclusion $\sigma \in \mathcal{Z}_t^p(\mathcal{C})$ is shown. \square

APPENDIX C. KANTOROVICH DUALITY FORMULA FOR ℓ_1

The disintegration result from Theorem 7.2 requires the existence of appropriate data E and u according to Section 6.1 in order to perform the disintegration. In this chapter, we describe this procedure by using basic Kantorovich duality theory for ℓ_1 . As in [1], we obtain an analog of the Kantorovich duality formula [77] which mirrors the famous duality formula for the 1-Wasserstein distance on a metric space in terms of 1-Lipschitz functions [129, p. 72]; cf. Theorem C.3.

Following [44, Sec. 2.4], we start with basic definitions from Kantorovich duality theory for ℓ_1 (see also [91, Ch. 4]). Given two sets $A, B \subset \mathbf{M}$, a function $\phi: A \rightarrow \mathbf{R}$ is *l-concave* relative to (A, B) if there exists a function $\psi: B \rightarrow \mathbf{R}$ with

$$\phi(x) = \inf_{y \in B} [\psi(y) - l(x, y)].$$

The *l-transform* of ϕ is the function $\phi^{(l)}: B \rightarrow \mathbf{R} \cup \{-\infty, \infty\}$ given by

$$\phi^{(l)}(y) := \sup_{x \in A} [\phi(x) + l(x, y)]. \quad (\text{C.1})$$

Analogous notions are set up by replacing l with l_+ .

Definition C.1 (Strong Kantorovich duality [44, Def. 2.23]). *A pair $(\mu, \nu) \in \mathcal{P}(\mathbf{M})^2$ is said to satisfy strong l-Kantorovich duality if*

- $0 < \ell_1(\mu, \nu) < \infty$, and
- there exist closed subsets $A, B \subset \mathbf{M}$ with $\mu[A] = \nu[B] = 1$, a μ -integrable upper semicontinuous function $\phi: A \rightarrow \mathbf{R}$ which is *l-concave* relative to (A, B) such that its *l-transform* $\phi^{(l)}: B \rightarrow \mathbf{R}$ is ν -integrable and obeys

$$\ell_1(\mu, \nu) = \int_{\mathbf{M}} \phi^{(l)} d\nu - \int_{\mathbf{M}} \phi d\mu. \quad (\text{C.2})$$

The notion of strong l_+ -Kantorovich duality of (μ, ν) is defined analogously by replacing l with l_+ .

Example C.2 (Criterion for strong Kantorovich duality). By the proofs of [44, Thm. 2.26, Cor. 2.29], every pair $(\mu, \nu) \in \mathcal{P}_c(\mathbf{M})^2$ satisfying

$$\text{spt } \mu \times \text{spt } \nu \subset \{l > 0\}$$

satisfies strong l_+ -Kantorovich duality (hence strong *l*-Kantorovich duality).

Theorem C.3 (Kantorovich duality for ℓ_1). *Assume that the pair $(\mu, \nu) \in \mathcal{P}_c(\mathbb{M})^2$ satisfies strong l -Kantorovich duality. Then the following hold.*

(i) **Kantorovich duality formula.** *We have*

$$\ell_1(\mu, \nu) = \inf \left[\int_{\mathbb{M}} u \, d\nu - \int_{\mathbb{M}} u \, d\mu \right].$$

The infimum is taken over all Borel functions $u: E \rightarrow \mathbf{R}$ defined on some Borel set $E \subset \mathbb{M}$ containing $\text{spt } \mu \cup \text{spt } \nu$ with $u \in L^1(\mathbb{M}, \mu) \cap L^1(\mathbb{M}, \nu)$ and for every $x, y \in E$, we have

$$u(y) - u(x) \geq l(x, y). \quad (\text{C.3})$$

(ii) **Existence of minimizers.** *The infimum is attained at a function $u: E \rightarrow \mathbf{R}$ with the additional property that E constitutes a compact, causally convex set containing $\text{spt } \mu \cup \text{spt } \nu$.*

(iii) **Optimal couplings.** *For every minimizer u in the above Kantorovich–Rubinstein formula and every ℓ_1 -optimal coupling $\pi \in \Pi(\mu, \nu)$,*

$$u \circ \text{pr}_2 - u \circ \text{pr}_1 = l \quad \pi\text{-a.e.}$$

Proof. To some extent, the argument is standard, cf. e.g. [129, p. 68]; let us only outline the key steps. Retain the notation from Definition C.1. Without loss of generality, we assume A and B are compact. On $E := J(A, B)$, define $u: E \rightarrow \mathbf{R}$ by the formula (C.1), that is,

$$u(y) := \sup_{x \in A} [\phi(x) + l(x, y)].$$

E is causally convex and, by global hyperbolicity, compact. Our specific choice of E ensures u does not attain the value $-\infty$. It is straightforward to verify (C.3) for every $x, y \in E$. Furthermore, we have $u = \phi^{(l)}$ on B , in particular $u \in L^1(\mathbb{M}, \nu)$. On the other hand, we have $u \geq \phi$ on A , while u is bounded from above on the compact set A by upper semicontinuity of ϕ and l , whence $u \in L^1(\mathbb{M}, \mu)$.

Since $0 < \ell_1(\mu, \nu) < \infty$ by the hypothesized strong l -Kantorovich duality, given any $\varepsilon > 0$ there is an ε -almost maximizer $\pi \in \Pi_{\leq}(\mu, \nu)$ in the definition of $\ell_1(\mu, \nu)$. By the previous part, this entails

$$\begin{aligned} \ell_1(\mu, \nu) &\leq \int_{\mathbb{M}^2} l \, d\pi + \varepsilon \\ &\leq \int_{\mathbb{M}} u \, d\nu - \int_{\mathbb{M}} u \, d\mu + \varepsilon \\ &\leq \int_{\mathbb{M}} \phi^{(l)} \, d\nu - \int_{\mathbb{M}} \phi \, d\mu + \varepsilon. \end{aligned}$$

Combining the arbitrariness of ε and (C.2) proves u attains the infimum in the Kantorovich duality formula for ℓ_1 . This yields the first two claims.

The last follows from the simple observation that the Kantorovich duality formula phrased in terms of an ℓ_1 -optimal coupling $\pi \in \Pi(\mu, \nu)$ forces equality in (C.3) for π -a.e. $(x, y) \in \mathbb{M}^2$. \square

REFERENCES

- [1] A. AKDEMIR. In preparation.
- [2] A. AKDEMIR, A. COLINET, R. J. MCCANN, F. CAVALLETTI, F. SANTARCANGELO. Independence of synthetic curvature dimension conditions on transport distance exponent. *Trans. Amer. Math. Soc.* **374** (2021), no. 8, 5877–5923.
- [3] L. AKÉ HAU, A. J. CABRERA PACHECO, D. A. SOLIS. On the causal hierarchy of Lorentzian length spaces. *Classical Quantum Gravity* **37** (2020), no. 21, 215013, 22 pp.
- [4] S. B. ALEXANDER, R. L. BISHOP. Lorentz and semi-Riemannian spaces with Alexandrov curvature bounds. *Comm. Anal. Geom.* **16** (2008), no. 2, 251–282.

- [5] A. D. ALEXANDROV. Additive set-functions in abstract spaces. *Rec. Math. [Mat. Sbornik] N.S.* **13(55)** (1943), 169–238.
- [6] L. AMBROSIO. *Lecture notes on optimal transport problems*. Mathematical aspects of evolving interfaces (Funchal, 2000), 1–52, Lecture Notes in Math., 1812. Springer-Verlag, Berlin, 2003
- [7] L. AMBROSIO, N. GIGLI. *A user's guide to optimal transport*. Modelling and optimisation of flows on networks, 1–155, Lecture Notes in Math., 2062, Fond. CIME/CIME Found. Subser., Springer, Heidelberg, 2013.
- [8] L. AMBROSIO, N. GIGLI, G. SAVARÉ. *Gradient flows in metric spaces and in the space of probability measures*. Second edition. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 2008. x+334 pp.
- [9] L. ANDERSSON, R. HOWARD. Comparison and rigidity theorems in semi-Riemannian geometry. *Comm. Anal. Geom.* **6** (1998), no. 4, 819–877.
- [10] E. AUBRY. Finiteness of π_1 and geometric inequalities in almost positive Ricci curvature. *Ann. Sci. École Norm. Sup. (4)* **40** (2007), no. 4, 675–695.
- [11] K. BACHER, K.-T. STURM. Localization and tensorization properties of the curvature-dimension condition for metric measure spaces. *J. Funct. Anal.* **259** (2010), no. 1, 28–56.
- [12] J. K. BEEM, P. E. EHRLICH. Cut points, conjugate points and Lorentzian comparison theorems. *Math. Proc. Cambridge Philos. Soc.* **86** (1979), no. 2, 365–384.
- [13] J. K. BEEM, P. E. EHRLICH, K. L. EASLEY. *Global Lorentzian geometry*. Second edition. Monographs and Textbooks in Pure and Applied Mathematics, 202. Marcel Dekker, Inc., New York, 1996. xiv+635 pp.
- [14] T. BERAN, M. BRAUN, M. CALISTI, N. GIGLI, R. J. MCCANN, A. OHANYAN, F. ROTT, C. SÄMANN. In preparation.
- [15] A. N. BERNAL, M. SÁNCHEZ. Globally hyperbolic spacetimes can be defined as ‘causal’ instead of ‘strongly causal’. *Classical Quantum Gravity* **24** (2007), no. 3, 745–749.
- [16] S. BIANCHINI, F. CAVALLETTI. The Monge problem for distance cost in geodesic spaces. *Comm. Math. Phys.* **318** (2013), no. 3, 615–673.
- [17] P. BILLINGSLEY. *Convergence of probability measures*. Second edition. Wiley Series in Probability and Statistics: Probability and Statistics. A Wiley-Interscience Publication. John Wiley & Sons, Inc., New York, 1999. x+277 pp.
- [18] V. I. BOGACHEV. *Measure theory. II*. Springer-Verlag, Berlin, 2007. xiv+575 pp.
- [19] L. BOMBELLI, J. NOLDUS. The moduli space of isometry classes of globally hyperbolic spacetimes. *Classical Quantum Gravity* **21** (2004), no. 18, 4429–4453.
- [20] A. BORDE. Geodesic focusing, energy conditions and singularities. *Classical Quantum Gravity* **4** (1987), no. 2, 343–356.
- [21] M. BRAUN. Rényi’s entropy on Lorentzian spaces. Timelike curvature-dimension conditions. *J. Math. Pures Appl. (9)* **177** (2023), 46–128.
- [22] ———. Good geodesics satisfying the timelike curvature-dimension condition. *Nonlinear Anal.* **229** (2023), Paper No. 113205, 30 pp.
- [23] M. BRAUN, M. CALISTI. Timelike Ricci bounds for low regularity spacetimes by optimal transport. To appear in *Commun. Contemp. Math.*
- [24] ———. In preparation.
- [25] M. BRAUN, K. HABERMANN, K.-T. STURM. Optimal transport, gradient estimates, and pathwise Brownian coupling on spaces with variable Ricci bounds. *J. Math. Pures Appl. (9)* **147** (2021), 60–97.
- [26] M. BRAUN, S. OHTA. Optimal transport and timelike lower Ricci curvature bounds on Finsler spacetimes. Preprint, [arXiv:2305.04389](https://arxiv.org/abs/2305.04389), 2023.
- [27] M. BRAUN, C. SÄMANN. In preparation.
- [28] C. J. FEWSTER, E.-A. KONTOU. A semiclassical singularity theorem. *Classical Quantum Gravity* **39** (2022), no. 7, Paper No. 075028, 33 pp.
- [29] P. J. BROWN, C. J. FEWSTER, E.-A. KONTOU. A singularity theorem for Einstein-Klein-Gordon theory. *Gen. Relativity Gravitation* **50** (2018), no. 10, Paper No. 121, 24 pp.
- [30] YU. BURAGO, M. GROMOV, G. PEREL’MAN. A. D. Aleksandrov spaces with curvatures bounded below. *Uspekhi Mat. Nauk* **47** (1992), no. 2, 3–51, 222; translation in *Russian Math. Surveys* **47** (1992), no. 2, 1–58.
- [31] D. BURAGO, YU. BURAGO, S. IVANOV. *A course in metric geometry*. Graduate Studies in Mathematics, 33. American Mathematical Society, Providence, RI, 2001. xiv+415 pp.
- [32] A. BURTSCHER, L. GARCÍA-HEVELING. Time functions on Lorentzian length spaces. Preprint, [arXiv:2108.02693](https://arxiv.org/abs/2108.02693), 2021.
- [33] A. BURTSCHER, C. KETTERER, R. J. MCCANN, E. WOOLGAR. Inscribed radius bounds for lower Ricci bounded metric measure spaces with mean convex boundary. *SIGMA Symmetry Integrability Geom. Methods Appl.* **16** (2020), Paper No. 131, 29 pp.
- [34] H. BUSEMANN. Timelike spaces. *Dissertationes Math. (Rozprawy Mat.)* **53** (1967), 52 pp.

- [35] L. A. CAFFARELLI, M. FELDMAN, R. J. MCCANN. Constructing optimal maps for Monge's transport problem as a limit of strictly convex costs. *J. Amer. Math. Soc.* **15** (2002), no. 1, 1–26.
- [36] S. CARROLL. *Spacetime and geometry*. An introduction to general relativity. Addison Wesley, San Francisco, CA, 2004. xiv+513 pp.
- [37] F. CAVALLETTI. Monge problem in metric measure spaces with Riemannian curvature-dimension condition. *Nonlinear Anal.* **99** (2014), 136–151.
- [38] F. CAVALLETTI, N. GIGLI, F. SANTARCANGELO. Displacement convexity of entropy and the distance cost optimal transportation. *Ann. Fac. Sci. Toulouse Math. (6)* **30** (2021), no. 2, 411–427.
- [39] F. CAVALLETTI, E. MILMAN. The globalization theorem for the curvature-dimension condition. *Invent. Math.* **226** (2021), no. 1, 1–137.
- [40] F. CAVALLETTI, A. MONDINO. Sharp and rigid isoperimetric inequalities in metric-measure spaces with lower Ricci curvature bounds. *Invent. Math.* **208** (2017), no. 3, 803–849.
- [41] ———. Sharp geometric and functional inequalities in metric measure spaces with lower Ricci curvature bounds. *Geom. Topol.* **21** (2017), no. 1, 603–645.
- [42] ———. Optimal maps in essentially non-branching spaces. *Commun. Contemp. Math.* **19** (2017), no. 6, 1750007, 27 pp.
- [43] ———. New formulas for the Laplacian of distance functions and applications. *Anal. PDE* **13** (2020), no. 7, 2091–2147.
- [44] ———. Optimal transport in Lorentzian synthetic spaces, synthetic timelike Ricci curvature lower bounds and applications. To appear in *Camb. J. Math.*
- [45] ———. A review of Lorentzian synthetic theory of timelike Ricci curvature bounds. *Gen. Relativity Gravitation* **54** (2022), no. 11, Paper No. 137, 39 pp.
- [46] ———. *A sharp isoperimetric-type inequality for Lorentzian spaces satisfying timelike Ricci lower bounds*. Preprint, [arXiv:2401.03949](https://arxiv.org/abs/2401.03949), 2024.
- [47] F. CAVALLETTI, K.-T. STURM. Local curvature-dimension condition implies measure-contraction property. *J. Funct. Anal.* **262** (2012), no. 12, 5110–5127.
- [48] J. CHEEGER, T. H. COLDING. Lower bounds on Ricci curvature and the almost rigidity of warped products. *Ann. of Math. (2)* **144** (1996), no. 1, 189–237.
- [49] P. T. CHRUSCIEL, J. D. E. GRANT. On Lorentzian causality with continuous metrics. *Classical Quantum Gravity* **29** (2012), no. 14, 145001, 32 pp.
- [50] Q. DENG, K.-T. STURM. Localization and tensorization properties of the curvature-dimension condition for metric measure spaces. II. *J. Funct. Anal.* **260** (2011), no. 12, 3718–3725.
- [51] M. ECKSTEIN, T. MILLER. Causality for nonlocal phenomena. *Ann. Henri Poincaré* **18** (2017), no. 9, 3049–3096.
- [52] M. EICHMAIR, G. J. GALLOWAY, D. POLLACK. Topological censorship from the initial data point of view. *J. Differential Geom.* **95** (2013), no. 3, 389–405.
- [53] R. ENGELKING. *General topology*. Translated from the Polish by the author. Second edition. Sigma Series in Pure Mathematics, 6. Heldermann Verlag, Berlin, 1989. viii+529 pp.
- [54] M. ERBAR, K. KUWADA, K.-T. STURM. On the equivalence of the entropic curvature-dimension condition and Bochner's inequality on metric measure spaces. *Invent. Math.* **201** (2015), no. 3, 993–1071.
- [55] M. FELDMAN, R. J. MCCANN. Monge's transport problem on a Riemannian manifold. *Trans. Amer. Math. Soc.* **354** (2002), no. 4, 1667–1697.
- [56] J. R. FLISS, B. FREIVOGEL, E.-A. KONTOU. The double smeared null energy condition. *SciPost Phys.* **14** (2023), no. 2, Paper No. 024, 33 pp.
- [57] L. H. FORD. Quantum coherence effects and the second law of thermodynamics. *Proc. R. Soc. Lond. A* **364** (1978), 227–236.
- [58] T. FRANKEL, G. GALLOWAY. Energy density and spatial curvature in general relativity. *J. Math. Phys.* **22** (1981), no. 4, 813–817.
- [59] B. FREIVOGEL, E.-A. KONTOU, D. KROMMYDAS. The return of the singularities: applications of the smeared null energy condition. *SciPost Phys.* **13** (2022), no. 1, Paper No. 001, 27 pp.
- [60] D. H. FREMLIN. *Measure theory. Vol. 4*. Topological measure spaces. Part I, II. Corrected second printing of the 2003 original Torres Fremlin, Colchester, 2006. Part I: 528 pp.; Part II: 439+19 pp. (errata).
- [61] F. GALAZ-GARCÍA, M. KELL, A. MONDINO, G. SOSA. On quotients of spaces with Ricci curvature bounded below. *J. Funct. Anal.* **275** (2018), no. 6, 1368–1446.
- [62] G. J. GALLOWAY. Compactness criteria for Riemannian manifolds. *Proc. Amer. Math. Soc.* **84** (1982), no. 1, 106–110.
- [63] ———. Curvature, causality and completeness in space-times with causally complete space-like slices. *Math. Proc. Cambridge Philos. Soc.* **99** (1986), no. 2, 367–375.

- [64] ———. Maximum principles for null hypersurfaces and null splitting theorems. *Ann. Henri Poincaré* **1** (2000), no. 3, 543–567.
- [65] L. GARCÍA-HEVELING. Causality theory of spacetimes with continuous Lorentzian metrics revisited. *Classical Quantum Gravity* **38** (2021), no.14, Paper No. 145028, 13 pp.
- [66] N. GIGLI, A. MONDINO, G. SAVARÉ. Convergence of pointed non-compact metric measure spaces and stability of Ricci curvature bounds and heat flows. *Proc. Lond. Math. Soc.* (3) **111** (2015), no. 5, 1071–1129.
- [67] M. GRAF. Singularity theorems for C^1 -Lorentzian metrics. *Comm. Math. Phys.* **378** (2020), no. 2, 1417–1450.
- [68] M. GRAF, E.-A. KONTOU, A. OHANYAN, B. SCHINNERL. Hawking-type singularity theorems for worldvolume energy inequalities. Preprint, [arXiv:2209.04347](https://arxiv.org/abs/2209.04347), 2022.
- [69] M. GRAF, E. LING. Maximizers in Lipschitz spacetimes are either timelike or null. *Classical Quantum Gravity* **35** (2018), no. 8, 087001, 6 pp.
- [70] A. GREVEN, P. PFAFFELHUBER, A. WINTER. Convergence in distribution of random metric measure spaces (Λ -coalescent measure trees). *Probab. Theory Related Fields* **145** (2009), no. 1–2, 285–322.
- [71] M. GROMOV. *Metric structures for Riemannian and non-Riemannian spaces*. Based on the 1981 French original. With appendices by M. Katz, P. Pansu and S. Semmes. Translated from the French by Sean Michael Bates. Progress in Mathematics, 152. Birkhäuser Boston, Inc., Boston, MA, 1999. xx+585 pp.
- [72] M. GROMOV, V. D. MILMAN. Generalization of the spherical isoperimetric inequality to uniformly convex Banach spaces. *Compositio Math.* **62** (1987), no. 3, 263–282.
- [73] S. W. HAWKING. Black holes in general relativity. *Comm. Math. Phys.* **25** (1972), 152–166.
- [74] S. W. HAWKING, G. F. R. ELLIS. *The large scale structure of space-time*. Cambridge Monogr. Math. Phys., No. 1. Cambridge University Press, London-New York, 1973. xi+391 pp.
- [75] R. A. HOUNNONKPE, E. MINGUZZI. Globally hyperbolic spacetimes can be defined without the ‘causal’ condition. *Classical Quantum Gravity* **36** (2019), no. 19, 197001, 10 pp.
- [76] R. KANNAN, L. LOVÁSZ, M. SIMONOVITS. Isoperimetric problems for convex bodies and a localization lemma. *Discrete Comput. Geom.* **13** (1995), no. 3–4, 541–559.
- [77] L. KANTOROVITCH. On the translocation of masses. *C. R. (Doklady) Acad. Sci. URSS (N.S.)* **37** (1942), 199–201.
- [78] M. KELL, S. SUHR. On the existence of dual solutions for Lorentzian cost functions. *Ann. Inst. H. Poincaré C Anal. Non Linéaire* **37** (2020), no. 2, 343–372.
- [79] C. KETTERER. Evolution variational inequality and Wasserstein control in variable curvature context. Preprint, [arXiv:1509.02178](https://arxiv.org/abs/1509.02178), 2015.
- [80] ———. On the geometry of metric measure spaces with variable curvature bounds. *J. Geom. Anal.* **27** (2017), no. 3, 1951–1994.
- [81] ———. The Heintze-Karcher inequality for metric measure spaces. *Proc. Amer. Math. Soc.* **148** (2020), no. 9, 4041–4056.
- [82] ———. Stability of metric measure spaces with integral Ricci curvature bounds. *J. Funct. Anal.* **281** (2021), no. 8, Paper No. 109142, 48 pp.
- [83] ———. Characterization of the null energy condition via displacement convexity of entropy. To appear in *J. London Math. Soc.*
- [84] B. KLARTAG. Needle decompositions in Riemannian geometry. *Mem. Amer. Math. Soc.* **249** (2017), no. 1180, v+77 pp.
- [85] E. H. KRONHEIMER, R. PENROSE. On the structure of causal spaces. *Proc. Cambridge Philos. Soc.* **63** (1967), 481–501.
- [86] M. KUNZINGER, M. OBERGUGGENBERGER, J. A. VICKERS. Synthetic versus distributional lower Ricci curvature bounds. To appear in *Proc. R. Soc. Edinb. A: Math.*
- [87] M. KUNZINGER, C. SÄMANN. Lorentzian length spaces. *Ann. Global Anal. Geom.* **54** (2018), no. 3, 399–447.
- [88] M. KUNZINGER, R. STEINBAUER. Null distance and convergence of Lorentzian length spaces. *Ann. Henri Poincaré* **23** (2022), no. 12, 4319–4342.
- [89] J. LOTT, C. VILLANI. Ricci curvature for metric-measure spaces via optimal transport. *Ann. of Math.* (2) **169** (2009), no. 3, 903–991.
- [90] L. LOVÁSZ, M. SIMONOVITS. Random walks in a convex body and an improved volume algorithm. *Random Structures Algorithms* **4** (1993), no. 4, 359–412.
- [91] R. J. MCCANN. Displacement convexity of Boltzmann’s entropy characterizes the strong energy condition from general relativity. *Camb. J. Math.* **8** (2020), no. 3, 609–681.
- [92] ———. A synthetic null energy condition. To appear in *Comm. Math. Phys.*
- [93] R. J. MCCANN, C. SÄMANN. A Lorentzian analog for Hausdorff dimension and measure. *Pure Appl. Anal.* **4** (2022), no. 2, 367–400.

- [94] E. MINGUZZI. On the causal properties of warped product spacetimes. *Classical Quantum Gravity* **24** (2007), no. 17, 4457–4474.
- [95] ———. Convexity and quasi-uniformizability of closed preordered spaces. *Topology Appl.* **160** (2013), no. 8, 965–978.
- [96] ———. Lorentzian causality theory. *Living Reviews in Relativity* **3** (2019), no. 3, 202 pp.
- [97] ———. Further observations on the definition of global hyperbolicity under low regularity. *Classical Quantum Gravity* **40** (2023), no. 18, Paper No. 185001, 9 pp.
- [98] E. MINGUZZI, S. SUHR. Lorentzian metric spaces and their Gromov-Hausdorff convergence. Preprint, [arXiv:2209.14384](https://arxiv.org/abs/2209.14384), 2022.
- [99] A. MONDINO, S. SUHR. An optimal transport formulation of the Einstein equations of general relativity. *J. Eur. Math. Soc.* **25** (2023), no. 3, pp. 933–994.
- [100] O. MÜLLER. Gromov-Hausdorff metrics and dimensions of Lorentzian length spaces. Preprint, [arXiv:2209.12736](https://arxiv.org/abs/2209.12736), 2022.
- [101] L. NACHBIN. *Topology and order*. Translated from the Portuguese by Lulu Bechtolsheim Van Nostrand Mathematical Studies, No. 4 D. Van Nostrand Co., Inc., Princeton, N.J.-Toronto, Ont.-London 1965 vi+122 pp.
- [102] J. NOLDUS. A Lorentzian Gromov-Hausdorff notion of distance. *Classical Quantum Gravity* **21** (2004), no. 4, 839–850.
- [103] S. OHTA. (K, N) -convexity and the curvature-dimension condition for negative N . *J. Geom. Anal.* **26** (2016), no. 3, 2067–2096.
- [104] ———. *Comparison Finsler geometry*. Springer Monogr. Math. Springer, Cham, [2021], ©2021. xxii+316 pp.
- [105] B. O’NEILL. *Semi-Riemannian geometry. With applications to relativity*. Pure and Applied Mathematics, 103. Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York, 1983. xiii+468 pp.
- [106] L. E. PAYNE, H. F. WEINBERGER. An optimal Poincaré inequality for convex domains. *Arch. Rational Mech. Anal.* **5** (1960), 286–292.
- [107] R. PENROSE. Gravitational collapse and space-time singularities. *Phys. Rev. Lett.* **14** (1965), 57–59.
- [108] P. PETERSEN, G. WEI. Relative volume comparison with integral curvature bounds. *Geom. Funct. Anal.* **7** (1997), no. 6, 1031–1045.
- [109] ———. Analysis and geometry on manifolds with integral Ricci curvature bounds. II. *Trans. Amer. Math. Soc.* **353** (2001), no. 2, 457–478.
- [110] T. RAJALA. Interpolated measures with bounded density in metric spaces satisfying the curvature-dimension conditions of Sturm. *J. Funct. Anal.* **263** (2012), no. 4, 896–924.
- [111] T. RAJALA, K.-T. STURM. Non-branching geodesics and optimal maps in strong $CD(K, \infty)$ -spaces. *Calc. Var. Partial Differential Equations* **50** (2014), no. 3–4, 831–846.
- [112] T. A. ROMAN. On the “averaged weak energy condition” and Penrose’s singularity theorem. *Phys. Rev. D* (3) **37** (1988), no. 2, 546–548.
- [113] C. SÄMANN. Global hyperbolicity for spacetimes with continuous metrics. *Ann. Henri Poincaré* **17** (2016), no. 6, 1429–1455.
- [114] R. SCHNEIDER. Konvexe Flächen mit langsam abnehmender Krümmung. *Arch. Math. (Basel)* **23** (1972), 650–654.
- [115] R. D. SORKIN, E. WOOLGAR. A causal order for spacetimes with C^0 Lorentzian metrics: proof of compactness of the space of causal curves. *Classical Quantum Gravity* **13** (1996), no. 7, 1971–1993.
- [116] C. SORMANI, C. VEGA. Null distance on a spacetime. *Classical Quantum Gravity* **33** (2016), no. 8, 085001, 29 pp.
- [117] C. SPROUSE. Integral curvature bounds and bounded diameter. *Comm. Anal. Geom.* **8** (2000), no. 3, 531–543.
- [118] S. M. SRIVASTAVA. *A course on Borel sets*. Graduate Texts in Mathematics, 180. Springer-Verlag, New York, 1998. xvi+261 pp.
- [119] R. STEINBAUER. The singularity theorems of general relativity and their low regularity extensions. *Jahresber. Dtsch. Math.-Ver.* **125** (2023), no. 2, 73–119.
- [120] K.-T. STURM. On the geometry of metric measure spaces. I. *Acta Math.* **196** (2006), no. 1, 65–131.
- [121] ———. On the geometry of metric measure spaces. II. *Acta Math.* **196** (2006), no. 1, 133–177.
- [122] ———. *Metric measure spaces with variable Ricci bounds and couplings of Brownian motions*. Festschrift Masatoshi Fukushima, 553–575, Interdiscip. Math. Sci., 17, World Sci. Publ., Hackensack, NJ, 2015.
- [123] ———. Distribution-valued Ricci bounds for metric measure spaces, singular time changes, and gradient estimates for Neumann heat flows. *Geom. Funct. Anal.* **30** (2020), no. 6, 1648–1711.

- [124] V.N. SUDAKOV. Geometric problems in the theory of infinite-dimensional probability distributions. *Proc. Steklov Inst. Math.* **141** (1979), 1–178.
- [125] S. SUHR. Theory of optimal transport for Lorentzian cost functions. *Münster J. Math.* **11** (2018), no. 1, 13–47.
- [126] F. J. TIPLER. Energy conditions and space-time singularities. *Phys. Rev. D* (3) **17** (1978), no. 10, 2521–2528.
- [127] J.-H. TREUDE, J. D. E. GRANT. Volume comparison for hypersurfaces in Lorentzian manifolds and singularity theorems. *Ann. Global Anal. Geom.* **43** (2013), no. 3, 233–251.
- [128] N. S. TRUDINGER, X.-J. WANG. On the Monge mass transfer problem. *Calc. Var. Partial Differential Equations* **13** (2001), no. 1, 19–31.
- [129] C. VILLANI. *Optimal transport. Old and new.* Grundlehren der mathematischen Wissenschaften, 338. Springer-Verlag, Berlin, 2009. xxii+973 pp.
- [130] R. F. WALD. *General relativity.* University of Chicago Press, Chicago, IL, 1984. xiii+491 pp.
- [131] S. WILLARD. *General topology.* Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont. 1970 xii+369 pp.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TORONTO, 40 ST. GEORGE STREET ROOM 6290, TORONTO, ONTARIO M5S 2E4, CANADA

Email address: braun@math.toronto.edu

Email address: mccann@math.toronto.edu