

Clocks and Chronogeometry: Rotating Spacetimes and the Relativistic Null Hypothesis

Tushar Menon, Niels Linnemann, and James Read

Abstract

Recent work in the physics literature demonstrates that, in particular classes of rotating spacetimes, physical light rays in general do not traverse null geodesics. Having presented this result, we discuss its philosophical significance, both for the clock hypothesis (and, in particular, a recent purported proof thereof for light clocks), and for the operational meaning of the metric field.

1 Introduction

2 Fletcher's Theorem

2.1 Maudlin on the clock hypothesis in special relativity

2.2 Fletcher's result in special relativity

2.3 Fletcher's theorem in general relativity

3 Electromagnetism and the Geometrical-Optical Limit

3.1 Maxwell's equations in curved spacetime

3.2 The geometrical-optical limit

3.3 Rotating spacetimes

3.4 Aren't Gödel spacetimes unphysical?

4 The Clock Hypothesis and Chronogeometry

4.1 Natural and mathematical observations

4.2 Clock registry discord

4.3 Chronogeometry

5 Conclusion

1 Introduction

In a recent paper, Fletcher ([2013], p. 1370) proves a remarkable theorem, which he interprets as demonstrating that ‘for any timelike curve in any spacetime, there is a light clock that measures the curve’s length as accurately and regularly as one wishes’. We take ‘measurement of a curve’s length’ to mean that a measuring device records intervals of proper time along the curve as given by the metric field of that spacetime. Fletcher takes a kinematical stance on the nature of light in relativistic spacetime theories, insofar as he takes it to be a defining characteristic of light rays that they always traverse null geodesics of the metric field. This turns out to be a crucial premise in the construction of his central theorem. However, one might ask: is this strictly true, of physical light rays, constructed of Maxwell fields? If the answer to this question is ‘no’, then there remains room to doubt the practical import of Fletcher’s result.

In this paper, we discuss a recent result by Asenjo and Hojman ([2017b]), which demonstrates that, at least in minimally coupled Maxwell theory in curved spacetime, the answer to the above question is indeed negative: in certain spacetimes, light rays, to the extent that they can be defined, do not traverse null geodesics; rather, their velocity is spacetime-dependent (that is, is dependent upon the position in the spacetime manifold of the light ray under consideration). In light of this dynamical (rather than kinematical) underpinning of the behaviour of light rays, we argue that it is not the case that physical light clocks (even if idealized) can be used, in general, to measure a spacetime curve’s length arbitrarily accurately.

The structure of this article is as follows. Having presented Fletcher’s theorem in Section 2, and the central result from Asenjo and Hojman in Section 3, the remainder of the article constitutes an in-depth exploration of the philosophical ramifications of the union of these

works.

To be more specific, in Section 4 we discuss the import of these results for the clock hypothesis—that is, for the foundational principle in general relativity that there exist ideal clocks that can measure the proper time along their worldlines, regardless of whether or not those worldlines be geodesic. Here, our central point is that the results of Asenjo and Hojman shake our confidence that certain clocks (in particular, light clocks) do indeed satisfy the clock hypothesis—*pace* Fletcher.

This done, we proceed to discuss the operational meaning, or ‘chronogeometric significance’, of the metric field in general relativity—that is, the metric field’s being surveyed by rods and clocks built from matter fields, in the sense of the latter reading off proper time intervals given by the metric field. Since the results of Asenjo and Hojman imply that the clock hypothesis need not be satisfied by what are traditionally understood to be ‘good’ clocks (namely, light clocks), a broader problem arises for the operational meaning of the metric field in general relativity—for it is harder for the metric field to be surveyed by matter fields, and thereby to acquire chronogeometric significance, than has hitherto been appreciated. This has consequences in particular for advocates of the ‘geometrical approach’ to spacetime theories,¹ according to whom the metric field (in some sense) compels configurations of matter fields to survey its structure—for the results presented in this paper provide evidence that, even in general relativity, such is not the case.² Moreover, if, following Knox ([2013], 2017)), one takes a field’s having chronogeometric significance to be a necessary condition for its qualifying as spatiotemporal, then these observations raise broader concerns regarding the status of spacetime in general relativity.

2 Fletcher’s Theorem

Fletcher situates his result in the context of the argument in (Maudlin [2012], Chapter 5) regarding the clock hypothesis in special relativity:

Maudlin [...] has recently argued that, given some additional assumptions, one

¹For works advocating this view, see, for example, (Friedman [1983]; Maudlin [2012]).

²For further recent discussion of the geometrical approach to spacetime theories, see, for example, (Brown and Read [2018]; Read [unpublished]; Read et al. [2018]).

can prove that the quantity an inertially moving light clock measures in Minkowski spacetime is the proper time along its worldline [...] The present paper generalizes [Maudlin's] result, indicating a direction in which one can extend Maudlin's argument to light clocks undergoing arbitrary acceleration in arbitrary spacetimes. (Fletcher [2013], p. 1370)

Anticipating our objection to Fletcher's interpretation of his theorem, we use in this paper the term 'null clock', rather than the more tendentious term 'light clock', to refer to a Langevin clock in which the oscillating material traverses null geodesics.³

The structure of this section is as follows. We begin in Section 2.1 by introducing the clock hypothesis (cf. Section 4) and examining Maudlin's argument for its approximate validity in special relativity; this relies on a particular idiosyncrasy of Minkowski spacetime—namely, its globally flat metric and affine structure. We discuss, briefly, the representation of a null clock and the manner in which it can be taken to read off intervals of proper time along its worldline. We then consider Maudlin's conditions for an ideal clock to measure such intervals even when subjected to certain forces. In Section 2.2, we introduce Fletcher's notation and recast Maudlin's argument in his terms. In Section 2.3, we remove the restrictions that Minkowski geometry imposes, and present Fletcher's result in its original context—general relativity.

2.1 Maudlin on the clock hypothesis in special relativity

Maudlin ([2012], p. 75) states the clock hypothesis as follows⁴:

The amount of time that an accurate clock shows to have elapsed between two events is proportional to the [i]nterval along the clock's trajectory between the two events.

³By 'Langevin clock', we mean (following Brown and Read [2018], Section 2) a clock consisting of two mirrors and an oscillating medium. The light clock is the paradigm example of a Langevin clock.

⁴As Maudlin himself points out, 'hypothesis' is something of a misnomer for this statement. It is more accurately seen as leading to a definition of a clock, as a physical body that measures the proper time along its worldline. For a contrary attitude, see (Brown and Read [2016], Section III.C), in which it is argued that no physical clock can be regarded as satisfying exactly the clock hypothesis—for all clocks will break eventually, when subject to sufficiently great accelerations. For these authors, what is of relevance is not exact satisfaction of the clock hypothesis, but rather approximate satisfaction, to some requisite degree of accuracy—this is what they dub the 'clock condition'. We return to these matters in Section 4.

A significant point to note about this hypothesis is that it makes a claim regarding clocks in all frames—not merely inertial frames. Maudlin argues for the validity of the clock hypothesis in special relativity for a particular class of clocks—null clocks. We begin by reviewing his construction.

Consider a smooth, connected, paracompact, Hausdorff, Lorentzian metric manifold $\langle M, g_{ab} \rangle$.⁵ In this subsection, we set g_{ab} to be η_{ab} , the flat Minkowski metric field of special relativity. Define two timelike (with respect to η_{ab}) curves $\gamma : I \rightarrow M$ and ${}^\alpha\gamma : J \rightarrow M$, where I and J are some open intervals on the real line, \mathbb{R} .

The Minkowski spacetime representation of a null clock consists of two material ‘mirror’ worldlines, represented by $\gamma[I]$ and ${}^\alpha\gamma[J]$, and a massless particle, the trajectory of which is represented by a series of null geodesics, bouncing between $\gamma[I]$ and ${}^\alpha\gamma[J]$.⁶ Call the trajectory of such a particle between two successive bounces a ‘null ray’. In a Lorentz coordinate frame,⁷ a configuration in which the null clock is at rest can be represented as in Figure 1.




Figure 1: A null clock in an inertial frame of special relativity.

Introducing some useful notation, we consider the closed interval $[I^1] = [s_0, s_1]$ of the do-

⁵Here we use abstract indices, in order to align with Fletcher. In the following sections of this paper, when we perform physical calculations, we switch to coordinate indices. In addition, unless otherwise stated, in this paper we work in natural units in which $c = G = 1$.

⁶Both Maudlin’s and Fletcher’s discussions are situated wholly within the context of classical general relativity. We therefore use the word ‘massless particle’ not to represent a quantum of some quantized field, but rather to represent the geometrical-optical limit of a classical field defined on a spacetime manifold. For more on the geometrical-optical limit, see (Goldstein [2011], Chapter 10; Misner *et al.* [1973], pp. 570–83; Perlick [2000], Chapters 5–6), and extensive discussion below.

⁷That is, an inertial frame of special relativity, in which the laws of physics are understood to take their simplest form.

main of γ ; given its reparametrization-invariance, we can assume that the intervals of parameter values are intervals of proper time along a worldline. The ‘halfway point’ on the image of γ (assessed with respect to η_{ab}) corresponds to $t_1 \in [I^1]$. Let h be the spatial distance, in a particular frame, between a point on $\gamma[I]$ and a point on ${}^\alpha\gamma[J]$, and let t be the proper time between s_0 and t_1 . By construction, this makes t the proper time between t_1 and s_1 as well. Let S be the spatiotemporal distance between $\gamma(s_0)$ and ${}^\alpha\gamma(j_0)$.

Minkowski geometry tells us that

$$S^2 = -t^2 + h^2. \quad (1)$$

By construction, the spatiotemporal distance between any two points on a null ray is always zero. This means that

$$t^2 = h^2. \quad (2)$$

So, if we shoot a massless particle from $\gamma(s_0)$, the proper time elapsed on $\gamma[I^1 \subsetneq I]$ between $\gamma(s_0)$ and $\gamma(s_1)$ is just $2h$. Clock ‘ticks’ correspond to points on $\gamma[I]$ where the oscillating particle ‘bounces off’ $\gamma[I]$ and travels towards ${}^\alpha\gamma[J]$. We label these as $\gamma(s_2), \gamma(s_3) \dots \gamma(s_n)$, where $n = {}^\alpha n \in \mathbb{N}$. The number of such points is referred to as the ‘bounce number’ for a segment of the trajectory, and denoted by ${}^\alpha n$. If we extend our interest to a larger segment $\gamma[I' \subset I]$ (where $[I' \supset I^1]$) of the image of $\gamma[I]$, then we will discover more points that correspond to clock ticks. The proper time elapsed on this larger segment of γ will then be

$$|I'| = |s_0 - s_n| = 2{}^\alpha n h. \quad (3)$$

What if we physically push our null clock, so that it changes its trajectory to one with a non-zero constant velocity, with respect to our original Lorentz coordinates? Call such an action a ‘Lorentz push’—a Lorentz push thus implements an active Lorentz transformation. On intuitions imported from classical spacetimes like Galilean spacetime, it might seem reasonable to assume that the spatial distance between the mirrors be preserved after the push, as measured in the original coordinate system. This hypothetical situation is represented on the left of Figure




Figure 2: Two configurations for possible null clocks, before and after a Lorentz push. (The region over which the Lorentz push is implemented is shaded in.) The clock on the left violates the relativity principle, for it ticks at different rates before and after the Lorentz push; the clock on the right satisfies the relativity principle, for it ticks at the same rate before and after the Lorentz push.

2. If such were the case, since our oscillating material traverses null rays, the post-Lorentz push clock would tick at a slower rate. This particular clock, then, would violate the relativity principle, for it would tick at different rates in different inertial frames.

The way out of this conundrum is well known: as Maudlin ([2012], pp. 112–3) states (essentially following the moral of Bell’s ([2001]) rocket experiment, a Lorentz transformation will bring the two mirrors together. (How exactly Maudlin argues to this conclusion is elaborated below.) Thus, the post-Lorentz push state of the clock under consideration will appear as on the right hand side of Figure 2. As a result, the null clock will continue to tick at the same rate, and so does satisfy the relativity principle. (In all of our considerations up to this point, we are ignoring the behaviour of the clock during the period of its acceleration over the course of the Lorentz push; we return to these matters shortly.)

In order to demonstrate the ‘approximate validity of the clock hypothesis’, Maudlin’s argument proceeds in two phases: (A) demonstrate that inertially moving clocks measure intervals of proper time along their worldlines as given by the metric field, even after a Lorentz push; (B) demonstrate that clocks in arbitrary states of motion approximately measure intervals of proper time on their worldlines as given by the metric field.

To achieve (A), Maudlin constructs a null clock by attaching the two mirrors to opposite ends of a rigid rod. A rod is said to be ‘rigid’ just in case, after a Lorentz push, it assumes

a configuration that, when expressed in the coordinates of the boosted frame (that is, its rest frame after the Lorentz push), is identical to its configuration in the coordinates of the original frame, before the Lorentz push.⁸ We say that the state of a rigid rod in its rest frame is its ‘equilibrium state’. The result of this construction is that the Lorentz push will cause the null clock to contract. The consequence, of course, is that such clocks satisfy (A), as given above.⁹

Now to (B). If we restrict our attention to segments of the null clock’s worldline in which the rigid rod is in an equilibrium state, then Lorentz pushes do not affect the clock—in the sense that our story regarding the behaviour of the clock is restricted to a story regarding its behaviour in its equilibrium states, in which (as we have already seen) the clock does read off intervals of proper time along its worldline. In the limit that we can approximate arbitrary motion as being composed of segments of inertial motion separated by arbitrarily short Lorentz pushes, Maudlin claims to achieve (B): a demonstration of the satisfaction of the clock hypothesis by null clocks in arbitrary states of motion.¹⁰

Fletcher ([2013], pp. 1370–1) correctly and reasonably points to the restricted scope of Maudlin’s argument: it relies on the notion of an equilibrium state of a rigid rod, and, moreover, is only approximate, in its application to null clocks in arbitrary states of motion. Motivated by these concerns, Fletcher generalizes Maudlin’s argument to account (he claims) for the satisfaction of the clock hypothesis by null clocks in arbitrary states of motion in both special and general relativity.

2.2 Fletcher’s result in special relativity

Let us tell the first part of Maudlin’s story from Fletcher’s more general perspective, using his notation. Consider the tangent space $T_p M$ at the point $p = \gamma(t_i) \in M$. Since Minkowski spacetime is a globally flat metric manifold, with a global smooth atlas, for an element $\rho^a \in T_p M$, the ‘exponential map’, $\exp_p(d\rho^a)$, where $d \in \mathbb{R}$, defines a smooth curve.¹¹ The metric

⁸Thus, a rigid rod would not be distorted by a Lorentz push in the manner of the null clock on the left of Figure 2, discussed above.

⁹One sense in which Fletcher generalizes Maudlin’s work is that he is not committed to the existence of such a rigid rod—see below.

¹⁰Note that this now includes the shaded regions in Figure (2), in which the clock is accelerated in order to implement the Lorentz push.

¹¹The exponential map at $p \in M$, $\exp_p : U_p \rightarrow M$, is defined on a subset U_p of the tangent space $T_p M$ as follows. First, $0 \in U_p$ and $\exp_p 0 = p$. Then any nonzero $\alpha^a \in U_p$ if and only if there is a geodesic $\gamma : [0, 1] \rightarrow M$

also defines an inner product¹² on each tangent space, which induces a norm.¹³ This norm generates an open-ball topology on $T_p M$ with respect to which one can define a ‘normal open neighbourhood’.¹⁴ If a normal open neighbourhood U_0 of $T_p M$ can be exponentiated to recover an open neighbourhood U of $p \in M$, then U is known as a ‘simply convex neighbourhood’ (Fletcher [2013], p. 1372).

For a timelike curve such as $\gamma[I]$, one can, at each point, identify an orthogonal unit spatial vector ρ^a , thus defining a vector field on $\gamma[I]$. Each vector can be exponentiated to map a point on $\gamma[I]$ to a point on another timelike curve, ${}^\alpha\gamma[J]$. If $\gamma[I]$ is an inertial trajectory, then multiplying each unit spatial vector by the same constant, ${}^\alpha d$, one can map $\gamma[I]$ to another timelike inertial trajectory. Each choice of α therefore generates a different parallel curve.

Call the exponential factor, ${}^\alpha d$, the ‘scalar radius’ of the curve ${}^\alpha\gamma$. In Minkowski spacetime, this can be taken to be equal to the spatial distance between the points $\gamma(t_1)$ and $\exp_p({}^\alpha d\rho^a) = {}^\alpha\gamma(j_0)$.^{15,16} So, Equation (3) can be rewritten as

$$|I'| = |s_0 - s_n| = 2n{}^\alpha d. \quad (4)$$

2.3 Fletcher’s theorem in general relativity

Maudlin’s argument applies only to Minkowski spacetime. We now generalize to arbitrary Lorentzian manifolds. To do so, we drop the requirement that the metric field be flat (indeed,

with tangent vector α^a at p such that $\gamma(0) = p$. Finally, for such nonzero $\alpha^a \in U_p$, $\exp_p \alpha^a = \gamma(1)$, which is well defined since the geodesic γ corresponding to α^a is unique (Fletcher [2013], pp. 1371–2).

¹²A real vector space V is a real inner product space just in case there is a function $\langle \cdot, \cdot \rangle$ from $V \times V$ to \mathbb{R} such that $\forall x, y, z \in V$, (i) $\langle \cdot, \cdot \rangle$ is positive-definite, that is, $\langle x, x \rangle \geq 0$ and $\langle x, x \rangle = 0$ iff $x = 0$, (ii) $\langle x, y + z \rangle = \langle x, y \rangle + \langle x, z \rangle$, (iii) $\langle x, \alpha y \rangle = \alpha \langle x, y \rangle$, and (iv) $\langle x, y \rangle = \langle y, x \rangle$.

¹³A norm $|\cdot|$ on a real vector space V is a real-valued function that associates a real number (a ‘length’) to elements of a vector space and satisfies the following three conditions: (i) absolute homogeneity: $\forall a \in \mathbb{R}, \forall v \in V, |a \cdot v| = |a||v|$, (ii) Triangle inequality: $\forall x, y \in V, |x + y| \leq |x| + |y|$, (iii) Non-degeneracy: $\forall x \in V, |x| = 0 \Rightarrow x = 0$. An inner product on a vector space induces a norm as $\forall x \in V, |x| = \sqrt{\langle x, x \rangle}$.

¹⁴Given $p \in M$, an open neighborhood $U_0 \subset T_p M$ containing the zero vector is called normal if and only if (i) $\exp_p|_{U_0}$ is well defined and a diffeomorphism onto its image, and (ii) U_0 is closed under scalar multiplication by $t \in [0, 1]$ (Fletcher [2013], p. 1372).

¹⁵This is because the tangent space at a point in Minkowski spacetime is isomorphic (as both a vector space and a manifold) to Minkowski spacetime. So there is a canonical one-one correspondence between vectors in $T_p M$ and straight lines in M .

¹⁶We should stress, as Fletcher does, that although in this case one can think of the scalar radius as being equal to the distance between the mirrors, this is not a privileged measure. ‘One could very well pick some other spacelike vector field on γ and some other scalar parameter to trace out the same companion curve, and this new pair would bear a systematic functional relationship to ρ^a and ${}^\alpha d$. The constraints on ${}^\alpha d$, as determined by the theorem, would then fix constraints on this new scalar parameter’ (Fletcher [2013], p. 1381).

we also drop the requirement that it have constant curvature). As a result, we lose global simple convexity—that is, the property that M is itself a simply convex neighbourhood. However, since we are dealing with a Lorentzian manifold, we still have ‘local simple convexity’,¹⁷ that is, with respect to the tangent space topology induced by the dynamical Lorentzian metric field g_{ab} , $\exists U_0 \subset T_p M, d \in \mathbb{R} : \exp_p(dU_0) = U_n \subset M$, and $\bigcup_n U_n = M$. If M is not flat, then $U \subsetneq M$ is an open neighbourhood of $p \in M$. In general, we can now only recover local patches of the manifold by exponentiating element of the tangent space at a point.

For a particular family of so-called convergent companion curves, $\{\gamma^\alpha\}_{\alpha \in \mathbb{N}}$, Fletcher’s theorem makes two assertions, which we dub ‘accuracy’ and ‘regularity’

$$\text{Accuracy: } \lim_{\alpha \rightarrow \infty} 2^\alpha n d = |I'|.$$

$$\text{Regularity: } \limsup_{\alpha \rightarrow \infty} \left\{ |(s_i^\alpha - s_{i-1}^\alpha) - (s_j^\alpha - s_{j-1}^\alpha)| : 1 \leq i, j \leq n \right\} = 0.$$

Accuracy is the statement that the times elapsed between ‘ticks’ measured on $\gamma[I]$ are equal¹⁸ to proper time intervals on the worldline. Regularity asserts that the proper time interval measured between any two arbitrary pairs of successive ticks is the same. If we restrict our interest to physical clocks, we see that it is regularity, not accuracy, that is significant. Maudlin ([2012], p. 106) expands:

An ideal clock is some observable physical device by means of which numbers can be assigned to events on the device’s worldline, such that the ratios of differences in the numbers are proportional to the ratios of [i]interval lengths of segments of the world-line that have those events as endpoints.

Accordingly, note that the null character of light rays does no operationally significant work here—we would still be able to construct ideal clocks in which the oscillating material traverses timelike paths. What is important is that the timelike curve $\gamma[I]$ is defined with respect to the same metric as the one surveyed by the oscillating matter—this is what ensures regularity. This is guaranteed by Fletcher’s assumption that light travels on null geodesics of g_{ab} , but would just

¹⁷See, for example, (O’Neill [1983], p. 131) for a proof that all Lorentz manifolds are locally simply convex.

¹⁸Note the difference compared to Maudlin’s statement of the clock hypothesis above. All real numbers are proportional to each other, so it is the equality of the clock’s proper time ascription and the metric’s proper time ascription that constitutes accuracy.

as easily be achieved by any oscillating material that travels at a constant velocity between the mirrors.¹⁹

Looking at Fletcher's theorem more closely, we see that, in general, we can still interpret d^α as being equal to the distance between the mirrors— d^α is constrained to be non-zero ([2013], p. 1376)—but only for the specific configurations of the null clock discussed below. As we will see, this is an important constraint on Fletcher's theorem: the theorem is only valid for clock configurations where this identification can be made.

Given $d^\alpha > 0$, the curve $\gamma[I]$ is not an element of $\{\gamma^\alpha\}_{\alpha \in \mathbb{N}}$. This ensures that a physical 'bounce' is always possible. The limit in the family of convergent curves is required to ensure that the two curves are sufficiently 'nearby' that γ^α can be arrived at by exponentiating the unit spacelike vector field, ρ^a , on $\gamma[I]$ —in other words, it ensures that they are in the same non-disjoint union of simply convex neighbourhoods. In Minkowski spacetime, global simple convexity ensures that these two curves can be arbitrarily far apart, and the null clock (in principle) still functions as an ideal clock.

Since spacetimes in general relativity can be arbitrarily curved, the bounce number, n^α , for a given timelike curve $\gamma[I]$ will, in general, depend upon the path traversed by the massless particle between $\gamma[I]$ and (the image of) the selected companion curve, $\gamma^\alpha[J]$. In order to avoid this dependence, one needs to be careful about which companion curve one chooses. This is where local simple convexity comes in. In regions in which it is possible, for any timelike curve, $\gamma[I]$, to find a companion curve, $\gamma^\alpha[J]$, which can be reached by exponentiating the spacelike tangent field on $\gamma[I]$, Fletcher's theorem holds.²⁰ In such regions, the scalar radius d^α does approximate the spatial distance between exponential map-related points across the curves. The fact that every point on a Lorentzian metric has a local simply convex neighbourhood guarantees that Fletcher's theorem holds for all Lorentzian spacetimes.

The physical interpretation that Fletcher gives of his use of the local simple convexity assumption is the following: it allows a light clock to expand and contract arbitrarily as it moves through the manifold, thereby accounting for possible gravitational tidal forces in generic gen-

¹⁹Fletcher ([2013], p. 1382) makes this observation.

²⁰Technically, what is required is that for any segment, $\gamma_k^k[I]$, of the curve $\gamma[I]$, such that $\bigcup_k \gamma_k^k[I] = \gamma[I]$, there exists a $\gamma_m^m[J]$ such that $\exp_p(\rho^a) = \gamma_m^m[J]$, $\forall p \in \gamma_k^k[I]$ and $\bigcup_m \gamma_m^m[J] = \gamma[J]$.

eral relativistic spacetimes. As long as the clock is sufficiently small that both mirrors are always within the same simply convex neighbourhood, Fletcher argues, it will measure proper time along $\gamma[I]$, and thereby satisfy the clock hypothesis. One way of putting our disagreement with Fletcher is the following: we argue (see Section 3) that the light rays themselves in physical light clocks will manifest spacetime-dependence. This dependence is what spoils the ideality of a light clock—see below.^{21,22}

Let us return now to Maudlin's null clock. Recall that the rigid rod between the mirrors ensures that their spatial separation, as measured in the rest frame of the pre-Lorentz push clock, is less after the push than before. But, as mentioned earlier, Maudlin says nothing about what goes on during the Lorentz push and in the time that it takes for the rod to reach its equilibrium state post-push. Therefore, his clock can only be guaranteed to measure proper times on trajectories to within the level of accuracy that the rigidity of the rod allows. One might abstract away from the use of the rigid rod, and instead impose a 'clock constraint', which restricts us to clocks whose mirrors' spatial separation changes in accordance with what an idealized rigid rod would have imposed on their configuration. But then one is presented with the non-trivial problem of showing that such systems exist.

In effect, the simple convexity of the open neighbourhoods around points on timelike trajectories in Lorentzian manifolds is what allows Fletcher to impose the clock constraint, and with it to prove the clock hypothesis for null clocks. More precisely, since, *ex hypothesi*, light travels on null geodesics and massive particles on timelike geodesics of g_{ab} , Fletcher's theorem proves that, within a given simply convex neighbourhood of a point on the manifold, an abstract version of Maudlin's construction holds, and any timelike trajectory can be approximated as a series of inertial trajectories linked by Lorentz pushes.

3 Electromagnetism and the Geometrical-Optical Limit

As discussed, Fletcher's assumption that light propagates on null geodesics—call this the 'relativistic null hypothesis'—is central to his theorem, for it suffices to ensure regularity. It is often

²¹It is important to note that we are conceding to Fletcher many things—for example, that it is possible to find a physical medium that emulates the behaviour of the mirrors of his light clock. In our view, there is good reason to doubt whether any such medium can be found—though we set the matter aside in the remainder of this paper.

²²Note that Maudlin's sense of the ideality of a clock is different from that deployed in Footnote 32.

taken for granted that the relativistic null hypothesis is satisfied in the ‘geometrical-optical limit’ of general relativity, in which, roughly speaking, the length scale over which the wavelength of the wave under consideration changes is much shorter than the length scale over which curvature changes in the ambient spacetime (cf. Misner *et al.* [1973], Section 22.5). In this section, we reconsider the physicality of the geometrical-optical limit, and thereby whether it is indeed true that light rays, qua solutions to Maxwell’s equations in curved spacetime, invariably traverse null geodesics.

It turns out that the answer is ‘no’—though a full explication of this answer will require some setup. First, some remarks on the field of optics. Research in this area is concerned with aspects of the behaviour of electromagnetic waves that are determined by how those waves propagate in the domain in which they can be approximated as rays. In this approximation, information about interference and diffraction is discarded, but, in compensation, one is able to derive a plethora of theorems and results regarding the observed behaviour of light. The approximation here is that the wavelength of the light being studied is significantly smaller than (a) the distances over which its amplitude varies, and (b) the distances over which curvature effects are non-negligible. Given that the wavelength of light studied in the context of optics is of the order of 100–1000nm, and typical experiments are conducted over distances at least 10^4 times as large, this optical approximation is generally a reasonable one.

In the following, we focus upon the behaviour of light in the limit that (a) and (b) are always satisfied. Since, for light of any wavelength, one can always find a Lorentzian manifold whose curvature varies on a scale comparable to that wavelength, the only way of guaranteeing the generality of a statement (such as Fletcher’s) made on the basis of such optical approximations, is by considering light in the limit that its wavelength tends to zero (equivalently, its frequency tends to infinity), and its amplitude is constant—this is what we mean, precisely, by the ‘geometrical-optical limit’.

It is worth pausing here, briefly, to discuss the physical intuitions behind taking this limit. Recall from Section 2 that what the construction of a null clock is intended to capture is the periodic motion of a particle bouncing back and forth between two mirrors. In classical physics, although light is not described by a particle, it is common to refer to the ‘path traversed by light’.

The standard interpretation of this locution is in terms of some limit. To the extent that we can interpret a solution to Maxwell's equations as instantiating a wave packet that is sufficiently localized and dynamically robust (that is, continues to exist as a wave packet on the time scales of interest), we can talk about a path traversed by light. Moreover, the wave packet can be approximated as a plane-wave solution—the divergence in behaviour between the two types of solution only manifests itself on large length and time scales.

By choosing a suitably short wavelength, then, this wave packet can be localized to well below the length scale of interest. The trajectory of the wave packet can then be thought of as an integral curve to some vector field—these integral curves are the light rays. Call each element of this vector field a 'wave vector', k^μ . By construction, the wave vector always points in the direction of propagation of the wave in spacetime. In the plane-wave approximation, the wave vector is perpendicular to the 'wave front'—the line or surface connecting all points on a wave that are in phase. Therefore, if we know the wave vector field and the metric, we can calculate the wave fronts, which then constitute the one-form, k_μ , which is orthogonal to the vector field.

We can now reverse this construction. If we begin with the wave fronts, we can therefrom construct the wave vector field, and with it the integral curves that represent the wave packet's trajectory. On this picture, it is clear that a coherent picture of light as traversing any kind of path presupposes a robust notion of a wave front. Therefore, the solutions to Maxwell's equations that describe light as traversing paths have to be such that the wave fronts of the initial configuration are preserved under the dynamical evolution of the field. In Minkowski spacetime, it is easy to come by such solutions—all plane wave solutions, for example, have this feature.

A plane wave in Minkowski spacetime satisfies what we refer to in Section 3.1 as a 'standard wave equation'—these are equations whose solutions are (possibly infinite superpositions of) plane waves, whose wave fronts move at a constant velocity, thus defining a constant wave front one-form k_μ . As a result, the wave vector field is also constant, since the metric is constant. Therefore, the following equation, known as a 'dispersion relation', describes the propagation

of light in Minkowski spacetime:

$$\eta_{\mu\nu}k^\mu k^\nu = 0. \quad (5)$$

The wave vector is determined by the Maxwell wave equation, which itself depends on the background geometry, and from the dispersion relation it is clear that the wave vector k^μ is null—so light propagates on null geodesics.

When considering the propagation of light in curved spacetimes, we still need to hold onto the notion of wave front preservation. But, as will be shown explicitly in Section 3.2, the generalization of the notion of a plane wave in Minkowski spacetime to arbitrarily curved spacetimes (call them ‘generalized plane waves’) loses this important feature. It is therefore crucial to discover whether there exists some approximation in which this feature is still valid—this is what motivates the use of the geometrical-optical limit. Note that, in a generically curved spacetime, neither the metric field nor the one-form k_μ corresponding to a solution of the curved-spacetime Maxwell equations is, in general, constant across the manifold. More importantly, these k_μ need not even carry a straightforward interpretation as wave fronts of a propagating wave packet.

In generically curved spacetimes, what we call the ‘generalized dispersion relation’ is of the form

$$g_{\mu\nu}(x)k^\mu(x)k^\nu(x) = f(x), \quad (6)$$

where $f(x)$ is some function of spacetime.

In the limit that the wavelength of light tends to zero, it can be shown that the evolution of generalized plane waves does preserve (at least locally) wave fronts and wave vectors ((Misner *et al.* [1973], pp. 574–5). In this limit, we have recovered a notion of the trajectory of light. In order for it to be the case that these light rays are null, a further condition needs to be met,

$$g_{\mu\nu}(x)k^\mu(x)k^\nu(x) = f(x) = 0. \quad (7)$$

Call such a dispersion relation a ‘null dispersion relation’. In a certain class of rotating spacetimes, an example of which is discussed in Section 3.3, if we consider an arbitrary solution to the vacuum Maxwell equations, which does not describe the propagation of light rays, and then

take the geometrical optical limit, the result is not guaranteed to be a solution to those vacuum Maxwell equations. In other words, $f(x)$ is not guaranteed to vanish for the dispersion relation associated with exact generalized plane wave solutions that preserve wave fronts, even though it might do so in the dispersion relation arrived at from geometrical-optical limit of an approximate solution. Although Fletcher never explicitly mentions geometrical optics, his assumption of the validity of the relativistic null hypothesis, irrespective of the background spacetime in which the wave is embedded, is equivalent to this assumption.

The purpose of this section, therefore, is to argue against Fletcher's use of this assumption, by demonstrating that, in a certain class of spacetimes, there is no solution of the Maxwell equations that gives rise to a null dispersion relation. We will show that this leads to a violation of the condition of regularity, and this is ultimately what undermines the purported generality of Fletcher's theorem. In Section 3.1, we introduce the vacuum Maxwell equations in curved spacetime, which describe the propagation of light. We then discuss, in Section 3.2, how the geometrical-optical limit can—if physically appropriate—demonstrate that light rays traverse null geodesics. Finally, we turn, in Section 3.3, to the vacuum Maxwell equations in certain rotating spacetimes. We see that such solutions deliver a startling verdict on the behaviour of light rays—that they violate the relativistic null hypothesis. We should stress that, in this section, our concern is with certain spacetimes violating the relativistic null hypothesis due to the misalignment between approximate and exact solutions that arise from rotation, rather than curvature effects.²³

3.1 Maxwell's equations in curved spacetime

It is worth distinguishing between two types of 'wave equation'—a 'standard wave equation' and a 'Maxwell wave equation'. A standard wave equation for, say, a scalar field ψ takes the form

$$\frac{\partial^2 \psi}{\partial x^2} - \frac{1}{c^2} \frac{\partial^2 \psi}{\partial t^2} = 0, \quad (8)$$

²³We intend to explore possible violations of the relativistic null hypothesis due to curvature couplings in dynamical equations for matter fields, such as those discussed in (Read *et al.* [2018], Section 2), in a future article.

where c is the wave propagation speed. The most general solution to this kind of equation is a superposition of ‘plane wave solutions’, that is, solutions of the form

$$\psi(x, t) = \psi_0 \exp \left\{ i \left(\vec{k} \cdot \vec{x} \pm \omega \cdot t \right) \right\}, \quad (9)$$

where ψ_0 is a constant. The associated four-dimensional wave vector is $k^\mu = (\omega, \vec{k})$, and the four-dimensional wave front is $k_\mu = \eta_{\mu\nu} k^\nu$. From this, we see that a general formula for a four-dimensional wave vector is

$$k^\mu = \nabla^\mu \theta, \quad (10)$$

where θ is the phase of the wave and ∇^μ is the Minkowski metric-compatible derivative operator. This definition of the wave vector applies even to position-dependent phases. For any given plane wave, therefore, it makes sense to talk about a trajectory—it is just the integral curve associated with k^μ . Since k^μ is a constant, these integral curves are a family of straight lines in Minkowski spacetime corresponding to trajectories of rays travelling at velocity c .

The Maxwell equations in curved spacetime are given by²⁴

$$\nabla_\mu F^{\mu\nu} = 0, \quad (11)$$

$$\nabla_\mu F^{*\mu\nu} = 0, \quad (12)$$

where ∇_μ is the derivative operator compatible with a generic Lorentzian metric field $g_{\mu\nu}$ (with respect to which index contraction is performed),

$$F_{\mu\nu} := \nabla_\mu A_\nu - \nabla_\nu A_\mu = \partial_\mu A_\nu - \partial_\nu A_\mu \quad (13)$$

is the antisymmetric Faraday tensor defined in terms of the electromagnetic four-potential A_μ ,²⁵ and $F^{*\mu\nu} := \frac{1}{2} \epsilon^{\mu\nu\alpha\beta} F_{\alpha\beta}$ the dual of $F^{\mu\nu}$.²⁶

As we are dealing with the source-free Maxwell equations, Equation (12) amounts to nothing

²⁴Throughout this section, we switch to (Greek) coordinate indices, since we will find it convenient to perform calculations in a particular coordinate basis.

²⁵(13) holds in any coordinate basis, since connection components in this equation vanish by the symmetry of the connection. For details, see for example (Dirac [2012], p. 39).

²⁶Of course, in general relativity, the dynamical equations for the $g_{\mu\nu}$ field are the Einstein field equations,

but a Bianchi identity. So, the only equations that are not identically satisfied are Equation (11); using Equation (13), we have (Misner *et al.* [1973], p. 569):

$$\nabla_\mu \nabla^\mu A_\nu + R_\nu{}^\mu A_\mu = 0. \quad (14)$$

This is the second type of wave equation discussed in this paper—a Maxwell wave equation.

Consider a solution of the Maxwell wave equation in curved spacetime of the following form—call it a ‘generalized plane wave’:

$$A_z(t, x) = \xi(x) \exp \{iS(x) \pm i\omega \cdot t\}. \quad (15)$$

Associated with this solution are $k^\mu = (\omega, \partial_i S(x))$, and $k_\mu = g_{\mu\nu} k^\nu$, both of which are spacetime dependent. Moreover, given the spacetime dependence of k^i , in general this will not have an interpretation as a wave vector for a wave packet state, since such states do not, in general, exist in curved spacetimes. In order for wave packet states to be guaranteed to exist, it must be the case that the spacetime-dependence of the amplitude and phase are negligible. So we go to the geometrical-optical limit in which ξ , S and $g_{\mu\nu}$ are approximately constant. In this case light rays are solutions to a standard wave equation, so guaranteed to travel at constant velocity, although the claim that this velocity is c requires further argument.

3.2 The geometrical-optical limit

The statement that light, even in arbitrarily curved spacetimes, can be taken to traverse null geodesics in the geometrical-optical limit²⁷ relies on an insufficiently general assumption about which take the form

$$G_{\mu\nu} := R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi T_{\mu\nu},$$

where $G_{\mu\nu}$ is the Einstein tensor, $R_{\mu\nu}$ is the Ricci tensor, R is the Ricci scalar, and $T_{\mu\nu}$ is the stress-energy tensor associated with the matter fields that serve as sources for the gravitational field. The contribution of the electromagnetic field to the energy-momentum tensor is given by

$$T_{\text{Maxwell}}^{\mu\nu} := F^{\mu\lambda} F^\nu{}_\lambda - \frac{1}{4}g^{\mu\nu} F_{\lambda\rho} F^{\lambda\rho}.$$

These equations are called the ‘Einstein-Maxwell equations’. In this article, we largely drop consideration of the Einstein equations, for we are not in general concerned with back-reaction effects of the electromagnetic field on the metric field.

²⁷Such a statement is certainly pervasive in the literature. See, for example, Wald ([2010], p. 71): ‘in this approximation (known as the geometrical-optical approximation), light travels on null geodesics, a suggestion

the behaviour of solutions to Maxwell's equations in curved spacetime. On this limit, Misner *et al.* ([1973], p. 571) state:

[...] the fundamental laws of geometric optics are: (1) light rays are null geodesics; (2) the polarization vector is perpendicular to the rays and is parallel-propagated along the rays; and (3) the amplitude is governed by an adiabatic invariant which [...] states that the number of photons is conserved.

They are careful to stress, however, that these laws are derived from the basic assumption of geometrical-optics (namely, (a) and (b)) mentioned above. The relativistic null hypothesis (that is, Misner *et al.*'s (1)), follows from the basic assumptions of the geometrical-optical limit, only when one further important criterion is met—that the solutions arrived at in this limit do, in fact, approximate exact solutions to arbitrary accuracy. In this section, beginning with the pedagogical setup of Misner *et al.*, we discover the conditions under which the relativistic null hypothesis is valid.

Recall that a light ray is, by definition, a curve that is perpendicular to a surface of constant phase (that is, a wave front) (Misner *et al.* [1973], p. 573). In what follows, we look in detail at the behaviour of approximate solutions that allow us to recover the ray-like characterization of light familiar from optics—perhaps the easiest way to think about the approximation at play is to see it as being motivated by a desire to make the curved spacetime model resemble the flat spacetime model by considering progressively smaller regions of spacetime.

So we begin with the picture in flat spacetime—here the electromagnetic equations are just the standard Maxwell equations, Equation (11) and Equation (12), defined with respect to the fixed Minkowski metric field $\eta_{\mu\nu}$.

Consider a vector potential A^μ , which is a solution to Equation (14). As a solution to a wave equation, it can always be decomposed into an ‘amplitude’ piece α^μ , and a phase piece $\theta \propto \frac{l}{\lambda}$, where l is the distance propagated by the wave and λ is its wavelenth. Thus,

$$A^\mu = \text{Re}(\alpha^\mu e^{i\theta}). \quad (16)$$

that can be confirmed by studies of the Green's function'. Or Malament ([2012], p. 147): '[The behaviour of] light [is determined by] the behavior of solutions to Maxwell's equations in a limiting regime ('the optical limit') where wavelengths are small [...] when one passes to this limit, packets of electromagnetic waves are constrained to move along (images of) null geodesics'.

Since we are in Minkowski spacetime, the curvature of the manifold, by definition, does not vary with distance. We can consider a wave-like solution the amplitude of which is constant across the manifold—such solutions are guaranteed to exist, since the Maxwell wave equation in Minkowski spacetime is a standard wave equation.

In curved spacetimes, the geometrical-optical approximation attempts to preserve the wave packet behaviour associated with certain solutions of the Maxwell equations. If we start with a specific solution to the Maxwell equations the amplitude of which does not change appreciably over the length scale of interest, and for which the ambient spacetime is approximately flat over the same length scale then, as the wavelength is decreased, these approximations apply to a larger class of spacetimes (that is, solutions to the Einstein field equations). Eventually, in the limit that the wavelength tends to zero, and the amplitude variation becomes negligible, it seems reasonable to assume that such solutions will coincide with exact ray-like solutions of the wave equation in any spacetime.

Let us examine one of the above-discussed limits—the constant amplitude limit. A general solution to the Maxwell equations in curved spacetime will describe a wave whose phase is a function of spacetime. If we restrict attention to regions of spacetime over which the amplitude varies very slowly, then minor corrections can be made to the amplitude at every point:

$$\alpha^\mu = a^\mu + b^\mu \lambda + c^\mu \lambda^2 + d^\mu \lambda^3 + \dots \quad (17)$$

Since the amplitude really depends on the choice of length scale L , the expansion is in powers of $\epsilon := \frac{\lambda}{L}$, so the vector potential in Equation (16) can now be expanded as

$$A^\mu = \text{Re} \left\{ \left(a^\mu + \epsilon b^\mu + \epsilon^2 c^\mu + \dots \right) e^{i\theta/\epsilon} \right\}. \quad (18)$$

Substituting our expansion into the source-free version of the Maxwell wave equations, Equation (14), and once again gathering terms that are to the order $(1/\epsilon^2)$, we get our familiar dispersion relation, $k^\mu k_\mu = 0$:

$$\begin{aligned}
0 &= -\nabla_\mu \nabla^\mu A^\nu + R^\nu{}_\mu A^\mu \\
&= \text{Re} \left\{ \left[\frac{1}{\epsilon^2} k^\mu k_\mu (a^\nu + \epsilon b^\nu + \epsilon^2 c^\nu + \dots) - 2 \frac{i}{\epsilon} k^\mu \nabla_\mu (a^\nu + \epsilon b^\nu + \dots) \right. \right. \\
&\quad \left. \left. - \frac{i}{\epsilon} \nabla_\mu k^\mu (a^\nu + \epsilon b^\nu + \dots) - \nabla_\mu \nabla^\mu (a^\nu + \dots) + R^\nu{}_\mu (a^\mu + \dots) \right] e^{i\theta/\epsilon} \right\}. \quad (19)
\end{aligned}$$

The relativistic null hypothesis is important for Fletcher's result only to the extent that it guarantees that light travels on trajectories of constant velocity—the actual value of this velocity is largely irrelevant. It is possible, therefore, that light rays travel at constant velocities, as judged by the affine connection, even if those geodesics are not null. In the terminology of Section 2, this would lead to a violation of accuracy but not regularity. So Fletcher's theorem would remain intact. If the dispersion relation were non-trivially spacetime-dependent, on the other hand, then the integral curves corresponding to the vector field defined by the wave vectors would not be straight lines according to the connection. This is what spells trouble for regularity, and so also for Fletcher's theorem.

In the class of curved spacetimes considered in (Asenjo and Hojman [2017b], Section 2), for example, in the geometrical-optical limit, light does propagate on null geodesics of the metric, determined by the coupled Maxwell and Einstein field equations. Consequently, in such spacetimes, both accuracy and regularity are preserved. The important assumption that allowed us to derive this result was that the perturbative expansion presented was treated as approximating the exact solution with arbitrary accuracy. As we will see in the following, this is not guaranteed to be the case.

3.3 Rotating spacetimes

Consider two possible means of arriving at solutions of the Maxwell wave equation, Equation (14), in curved spacetime for an arbitrarily small wavelength. The first is to solve it exactly by some technique. The second is to solve the wave equation for a relatively large wavelength, and then take the geometrical-optical limit. These two techniques are not guaranteed to agree on the space of solutions. More precisely, it is not guaranteed that a particular solution, in

the geometrical-optical limit, remains a solution to the Maxwell wave equations in curved spacetime. An example from (Asenjo and Hojman [2017b], Section 3), pertaining to Gödel spacetimes, demonstrates clearly this fact.

The Gödel metric, in Cartesian coordinates, reads

$$\begin{aligned} g_{00} &= -1 = -g_{xx} = -g_{zz}, \\ g_{yy} &= -2 + 4\exp(\sqrt{2}x\Omega) - \exp(2\sqrt{2}x\Omega), \\ g_{0y} &= \sqrt{2}[1 - \exp(\sqrt{2}x\Omega)], \end{aligned} \quad (20)$$

where Ω is a constant related to the angular velocity of the rotating universe.²⁸

Consider now the z -component Maxwell wave equation, Equation (14), in curved spacetime,

$$\partial_0^2 A_z + \frac{1}{\sqrt{-g}g^{00}}\partial_x(\sqrt{-g}\partial_x A_z) = 0. \quad (21)$$

For the Gödel metric, there is no choice of coordinates such that Equation (21) takes the form of a standard wave equation. There is, however, a choice of coordinates such that this equation takes the form

$$\partial_0^2 A_z + \sigma(\zeta)\partial_\zeta^2 A_z = 0, \quad (22)$$

where $\zeta = -e^{-\sqrt{2}x\Omega}/\sqrt{2}\Omega$. This looks somewhat similar to a standard wave equation, but with a position-dependent function determining the ‘velocity’ of the wave.

Let us approach Equation (22) with a generalized plane wave ansatz of the form

$$A_z(x, t) = \xi(x) \exp[i\omega t \pm iS(x)]. \quad (23)$$

In the geometrical-optical limit, this solution does have the form of a standard plane wave solution, for in that limit, $S(x) \propto x$ and $\xi(x) \approx \text{const.}$ However, outside that limit, the wave vector takes the form $k_0 = \omega; k_i = \pm\partial_i S(x)$, where $S(x)$ need not be a linear function of x . Therefore, $k^\mu k_\mu$ is not guaranteed to vanish. The exact form, derived from Equation (23), using

²⁸For a discussion of the difficulties in classifying a spacetime as ‘rotating’, as well as further foundational details on Gödel spacetime, see (Malament [2012], Chapter 3).

the formula for a wave vector for a generalized plane wave Equation (10), is

$$k_\mu k^\mu = \frac{1}{\xi \sqrt{-g}} \partial_x \left(\sqrt{-g} \partial_x \xi \right). \quad (24)$$

Equation (21) also allows us to derive the condition

$$\partial_x \left(\sqrt{-g} k_x \xi^2 \right) = 0. \quad (25)$$

Recall that, for the notion of a trajectory to be meaningful, the geometrical-optical approximation is used, in order to approximate solutions as plane waves. So, what we are looking for is a solution to Equation (21) for which (i) the geometrical-optical approximation holds, and which (ii) satisfies Equation (24) and Equation (25). This is demonstrably impossible—there are no exact solutions to Equation (21) for which $k^\mu k_\mu = 0$, along with the geometrical-optical constraints. In other words, the geometrical optical limit of the approximate solution is not itself a solution.

For the exact solution to Equation (21), one first uses Equation (25) to obtain

$$\xi(x) = \frac{\xi_0}{(-g)^{\frac{1}{4}} k_x^{\frac{1}{2}}}, \quad (26)$$

which can be plugged back into the dispersion relation Equation (24) (cf. Asenjo and Hojman [2017b], pp. 3–4), yielding the dispersion relation

$$k^\mu k_\mu = \frac{\Omega^2}{2} - \frac{k'_x}{2k_x} + \frac{3k_x'^2}{4k_x^2}. \quad (27)$$

Of central importance to the argument of this article is the fact that the expression is spacetime-dependent. The presence of first and second spatial derivatives of the spatial wave vector k_x indicates a dependence on (up to) third spatial derivatives of the function $S(x)$. The phase velocity $v_p := \frac{\omega}{k_x}$, and the group velocity $v_g := \frac{\partial \omega}{\partial k_x}$, will thus both generally deviate from c .

Consider for instance (following Asenjo and Hojman [2017b], Section 3) the case of very

small spacetime length scales ($\Omega x \ll 1$), such that the expression for the wave vector becomes

$$\omega \approx \left(1 + \frac{1}{2}\Omega^2 x^2\right)k_x. \quad (28)$$

We therefore find a phase velocity

$$v_p = \frac{\omega}{k_x} = 1 + \frac{1}{2}\Omega^2 x^2, \quad (29)$$

and thus a group velocity

$$v_g = \frac{\partial \omega}{\partial k_x} \approx 1 + \frac{1}{2}\Omega^2 x^2, \quad (30)$$

for $0 \leq \Omega^2 x^2 \ll 1$.

Both the phase and the group velocity of waves thereby exceed the speed of light. However, this does not necessarily mean that electromagnetic waves in Gödel spacetime can be used for faster-than-light signalling—it is a common misconception that the group velocity can straightforwardly be associated with the speed of information propagation (for a pedagogical clarification, see Chiao [unpublished]). Rather, it is the ‘front velocity’—the velocity with which sharp pulses modulated onto the wave can propagate—which denotes signalling speed.²⁹

But propagation of a signal requires that there can exist a well-defined signal in the first place. Thus, the pertinent question that arises now is the following: Does it make sense to talk about the propagation of wave front in Gödel spacetime at all? And thus, does it make sense to speak of a front velocity in such a spacetime? This would make sense only if the right-hand side of the dispersion relation Equation (27) varied only moderately with spacetime position. Only then would we be licensed to assume that there exists a well-defined package of information to be sent from one side of a clock to another, allowing us to realize a Langevin light clock in the first place.

We now face a dichotomy. Either (a) no reliable disturbance (against noise) can be propagated via light; or (b) if this is possible, then this signal will be propagated at a speed varying with spacetime location (as the front velocity varies with spacetime location). Both cases are

²⁹Even the claim that the front velocity denotes the signal speed in an actual experimental setting has not gone unchallenged; see (Garrison *et al.* [1998]).

problematic for Fletcher, for in scenario (a), one cannot construct a Langevin light clock at all, whereas in case (b), such a clock will violate regularity.

3.4 Aren't Gödel spacetimes unphysical?

An immediate response to the above argument suggests itself: if this result has only been shown to be applicable to Gödel spacetimes, and we have good reason to believe that our universe is not described by such a solution (on various conceptual grounds related to causality (Earman [1995], Chapter 6) and the ability to define a consistent quantum theory (for example, Schuller and Witte [2014])), then why should this result bother us? In more 'physical' spacetimes (Schwarzschild, FLRW, and de Sitter, for example), light rays do travel on null geodesics according to Maxwell's equations.

The most straightforward response to this objection is to note that it is plausible that the results presented in Section 3 generalize to more physical spacetimes—and indeed, Asenjo and Hojman ([2017b], Section 4) discuss the case of Kerr spacetimes. Even this notwithstanding, however, the above response overlooks the epistemological crisis to which the results of Section 3 give rise. In showing that, in a consistent solution to the Einstein field equations, a model of what is generally thought to be an ideal clock does not survey the metric, we have undermined any straightforward reasons for believing that we live in the sort of universe in which we can trust such a measuring device.³⁰ All empirical claims about the sort of universe in which we live are made based upon measuring devices that respond to the behaviour of matter and the metric field. This result shows that we have no way of using light clocks to determine whether the metric field we claim to have measured is, in fact, the metric field of the particular solution to the Einstein equations under consideration. In other words, it might be the case that our solution to the Einstein equations does contain, for example, closed timelike curves in the g_{ab} field, but the metric surveyed by our measuring devices (whose dynamics depends on g_{ab} , but

³⁰At this point, we are only referring to light clocks when discussing 'such measuring devices'. We do not rule out the possibility that other dynamical systems might exist that do satisfy the clock hypothesis. However, we must countenance the existence of possible worlds that consist only of material fields that consistently violate the clock hypothesis as light clocks do in Gödel spacetimes. It is to these spacetimes that the discussion in this subsection applies. In addition, two further points are in order here: (a) we find it plausible that our arguments regarding light clocks generalize to a wide class of clocks typically considered to be 'good' clocks (cf. Section 4.2); (b) once the epistemological crisis delineated in this section arises for light clocks, it plausibly generalizes to other clocks. Both of these matters are discussed in more detail below, and in Section 4.

are such that they do not survey g_{ab}) does not. The putative unphysicality of g_{ab} does not imply that the geometry associated with measuring devices governed by laws expressed with respect to g_{ab} is unphysical. This is another way of demonstrating the two distinct roles that the metric of general relativity—in any case equipped with own physical degrees of freedom—can play: (i) that it is a component of the formalism that allows us to articulate the laws, and (ii) that, in addition, it compels matter fields to behave in such a way as to survey it.

We find ourselves in a situation in which the very reason that we have for believing that we live in a particular spacetime is that we assume that ideal clocks always survey the metric field of the Einstein field equations, that is, we accept (ii). We cannot, therefore, be guaranteed by our own theory that there exists a reliable method of inference from the behaviour of light rays to the geometry of the g_{ab} field. It is fitting that the solution that demonstrates that there might be truths about a spacetime that cannot be determined from observations confined to that spacetime bears Gödel's name!³¹

4 The Clock Hypothesis and Chronogeometry

Consider a clock—realized as a particular configuration of matter fields—in a particular frame of reference. Call this clock ‘ideal’ in this frame just in case it can be used to read off intervals of proper time along its worldline, as given by the metric field. Now ask: under what further conditions does this clock read off intervals of proper time along its worldline in all frames, that is, in a frame-independent manner? As already elaborated in Section 2, a clock that satisfies this condition is one that satisfies the clock hypothesis.³²

We have seen that a necessary condition for a clock to satisfy the clock hypothesis is that

³¹Whether or not one finds this troubling may boil down to one's stance on what constitutes an adequate justification of this method of inference. This is analogous to debates in epistemology over the status of rules of inference such as induction and abduction—inference rules that are themselves rule circular. Adopting an externalist position would entail that the observations in this paper are not necessarily problematic, since all the externalist requires to be the case for us to be justified in believing that light surveys the metric field is that light does, in fact, survey geodesics of the metric field, whether or not we can point to some calculation or model that justifies (internally) our belief in its doing so.

³²A clock that reads off intervals along its worldline in all inertial frames may be called ‘ideal’ tout court. Note that ideality is a much weaker condition than satisfaction of the clock hypothesis. While it would be reasonable to claim that we do not need a clock satisfying the exact clock hypothesis in order to obtain operational access to the metric field, but only a clock that satisfies the ‘clock condition’ (cf. footnote 4), or (weaker still) an ideal clock (in this sense), in our view none of these apparatus are immune from the epistemological concerns raised in Section 3.4, and below.

it satisfy regularity—namely, the condition that the proper time interval measured between any two arbitrary pairs of successive ticks is the same. However, the results of Asenjo and Hojman ([2017a], [2017b]) presented in Section 3 demonstrate that this principle is not satisfied for light clocks in Gödel spacetimes, for in such cases the velocity of signal propagation is a function of spacetime coordinates. Thus, light clocks in such spacetimes do not satisfy the clock hypothesis.

What is the philosophical upshot of this work? Such results demonstrate that what are often considered the simplest, most reliable conceivable clocks may, in certain spacetimes, not accurately measure intervals as given by the metric field—that is, they may fail to be good clocks in these spacetimes. Such an observation gives rise to broader operational concerns: if such clocks, built using light rays, need not accurately survey the metric field, should we expect that the situation be any different for other clocks, built from different matter fields? If not, there arise pressing concerns regarding how one is to gain operational access to the metric field tout court.

The purpose of this section is to explore some of these philosophical concerns in more detail. In Section 4.1, we begin by framing the results of this paper in terms of Synge’s distinction between ‘natural observations’ and ‘mathematical observations’. In Section 4.2, we present a heuristic argument to the effect that one should not expect generic (Langevin) clocks to accurately record intervals along their worldlines. In Section 4.3, we consider the operational ramifications of situations in which different clocks read off different intervals along the same worldline, and in which we have no epistemic access to which of these readings, if any, correspond to the interval along this worldline as given by the metric field.

4.1 Natural and mathematical observations

It is helpful to view the work of this paper through the lens of Synge’s distinction between mathematical observations (MOs) and natural observations (NOs) (cf. Synge [1960], pp. 103–7)).³³ The distinction is roughly the following: while NOs are empirical observations,³⁴ MOs

³³We are grateful to an anonymous referee for this suggestion.

³⁴Synge ([1960], p. 103) distinguishes between uncontrolled, controlled, and imagined NOs; this more fine-grained distinction will not be important for our purposes.

are mathematical facts, constructs, and laws. Now, as Synge ([1960], p. 103) writes:

The peculiar fascination of theoretical physics lies in the art of forcing meaningful truth out of the meaningless equation $NO = MO$, which is a symbolic form of the assertion that natural phenomena obey exact mathematical laws. The true inequality $NO \neq MO$ should not be spoken above a whisper, because it is extremely dangerous. If believed, it would sever mathematics from physics, and reduce both to sterility through lack of mutual fecundation. It is whispered here only as an apology to those readers who expect to see the mathematics of relativity tied to the physics of relativity by a strong chain of clear thought. It cannot be done. We have to muddle through. And if this book is dishonest in confusing MO with NO , it is no more dishonest than all similar books are, and necessarily must be. This sad state of affairs is not peculiar to relativity; every branch of mathematical physics has in its cupboard the skeleton $MO \neq NO$.

Our work in the current paper brings to the fore an unexpected case of $MO \neq NO$.³⁵ To see this, consider a particular interval along a timelike path in a solution of general relativity, and consider the proper time along this path as given by the metric field. Since this is a theoretical construct, it is an MO ; call it MO_g . Now consider the time along this path as given by a theoretical model of a Langevin clock in which the oscillating matter is described by Maxwell fields. Again, this is a theoretical construct, and therefore a MO ; call it MO_F . In rotating spacetimes, we have seen in this paper, in light of the work of Asenjo and Hojman, that $MO_g \neq MO_F$.³⁶

Now, to which of MO_F or MO_g are associated the NO s of physical light rays? Naïvely, one might think, to both: to MO_F as light is described by Maxwell fields, and to MO_g given the mainstream view in relativistic physics that light rays propagate on null geodesics. But clearly, if in doubt, MO_F —being directly concerned with the physical nature of light—should be given

³⁵Of course, not one about which we should merely whisper!

³⁶Interesting questions arise even at this juncture. For example, could one define from the paths traversed by propagating light rays constructed from Maxwell fields—just as it is usually done for light rays assumed to move on null geodesics—in MO_F the conformal structure associated with the metric field g_{ab} in MO_g ? In light of the position-dependence of the behaviour of the rays in MO_F , one is inclined to say ‘no’. We are grateful to an anonymous referee for valuable input on these matters.

priority, so let us say that $MO_F = NO$. And since (as it turns out) $MO_g \neq MO_F$, we have that $MO_g \neq NO$. This result is more radical than standard cases of the kind $MO \neq NO$, precisely in virtue of its being radically unexpected. And to put Fletcher in these terms: he has provided a mathematical model to read off MO_g , but since (as discussed) we expect that $MO_g \neq NO$, the physicality of his model is questionable.

4.2 Clock registry discord

In this subsection, we argue that a generic Langevin clock should not be expected to record its worldline interval as given by the metric field in generic spacetimes. This result is significant, for it calls into question a basic assumption in relativity theory, again brought out in an illuminating discussion by Synge ([1960], p. 106):

It is necessary to expose here a certain physical assumption inherent in the structure of relativity. Let C [figure suppressed] be the world-line of a material particle, and B, A two events on it, with B before A . The particle carries two standard clocks consisting of atoms of different types, or two atoms of the same type but with the use of different energy levels. Each clock registers a definite number of ticks between B and A ; let these number be denoted by n_1 and n_2 . The physical assumption just referred to is the following hypothesis of consistency: *For two standard clocks, the ratio $n_1 : n_2$ is a natural constant, independent of the world-line on which the observations are made and of the events on that world-line.*

To bring out the sense in which our present work casts doubt upon this hypothesis of consistency, we reconsider the circumstances under which a given clock does indeed read off its worldline interval as given by the metric field. The worldline length of a timelike path $\gamma[I]$ is given by

$$\Delta s = \int_{\gamma} g_{\mu\nu} dX^{\mu} dX^{\nu}. \quad (31)$$

Dividing up the curve into equidistant segments $\{\gamma_i\}$ with respect to an arbitrary curve parameter $\lambda \in I$ gives

$$\Delta s = \sum_i \int_{\gamma_i} g_{\mu\nu} dX^{\mu} dX^{\nu}. \quad (32)$$

Denote the beginning of each segment γ_i by the point p_i . Choosing small enough segments, and taking each point p_i as the origin for normal coordinates, the metric around the origin is given by (cf. Viaclovsky [2011], p. 22)

$$g_{\mu\nu} = \eta_{\mu\nu} - \frac{1}{3}R_{\mu\lambda\nu\rho}q^\lambda q^\rho + \dots, \quad (33)$$

where q^ρ denotes the components of the vector to the point q considered in normal coordinates (the metric is Minkowskian at the origin). The worldline interval of the path γ thus splits into a curvature-free and a curvature-dependent part, that is,

$$\Delta s = \int_{\gamma} \eta_{\mu\nu} dX^\mu dX^\nu - \sum_i \int_{\gamma_i} \frac{1}{3} R_{\mu\lambda\nu\rho} q_i^\lambda q_i^\rho dX^\mu dX^\nu + \dots \quad (34)$$

Remember now that a Langevin clock is realized through a back-and-forth signalling process in an oscillating medium. Such a physical clock can effectively read off the worldline interval just in case, for each segment γ_i , it can be seen to evolve as if it were situated in local Minkowski spacetime, while also being sensitive to curvature in just such a way as to register the higher-order terms on the right hand side of Equation (34). However, different matter fields are governed by different dynamical equations, which in turn may feature different curvature couplings.³⁷ Thus, even for Langevin clocks, it is (as we see it, on the above heuristic grounds) to be regarded as implausible that clocks built from two or more different matter fields should correctly record—or even agree upon—the full interval along a given worldline, as given by the metric field Equation (34). But in that case, we have no guarantee—or even good reason to think—that the hypothesis of consistency will hold.

4.3 Chronogeometry

Consider cases in which regularity is lost—such as those discussed in Section 3—and in which Langevin clocks built from Maxwell fields accordingly do not correctly read off intervals as given by the metric field. Since such is the case for Maxwell fields, it is *prima facie* plausible that Langevin clocks built from other matter fields also do not correctly read off intervals as

³⁷Not to mention different sensitivities to rotation—cf. Section 3.

given by the metric field in such cases; moreover, in light of the reflections presented in Section 4.2, there exists no a priori reason to expect that such clocks will agree on the intervals read off along a particular section of a given worldline.³⁸ In this subsection, we reflect upon some operational concerns that would arise in such a scenario. The central question to be discussed is the following: if different clocks all read off different intervals along the same worldline, then how does one get a bead on the ‘true’ geometry of the metric field—that is, how is the metric field afforded its operational meaning?

Focusing upon Langevin clocks for simplicity, there exist two scenarios worthy of consideration here: (A) the dispersion relation of the oscillating medium manifests constant space-time dependence, and (B) the dispersion relation of the oscillating medium manifests variable spacetime dependence. In scenario (A), the matter fields constituting the oscillating media in the Langevin clocks under consideration possess dispersion relations of the form $k^\mu k_\mu = \text{const.}$ In this case, since this dependence is constant across spacetime, all clocks satisfy regularity, as discussed in Sections 2-3. Therefore, clock ticks remain proportional, and hence a universal notion of the time along a given worldline may be recorded. As mentioned above, this scenario is compatible with Fletcher’s theorem, for all Fletcher requires is that the signal in one’s clock travel at a constant velocity—he explicitly acknowledges that this may differ from c . In scenario (B), the matter fields constituting the oscillating media in the Langevin clocks under consideration possess dispersion relations of the form $k^\mu k_\mu \neq \text{const.}$ In this latter scenario, genuine operational concerns do arise—for in this case, the variable spacetime dependence in the dispersion relations of the oscillating media means that the ticks of the clocks need no longer be proportional to one another; hence, regularity is lost. In this case, one cannot take the ratios of such intervals and infer universally the proper time along a given worldline; indeed, there appears to be no way, using these matter fields alone, to gain epistemic access to the intervals of proper time along this worldline as given by the metric field.

What are the consequences of the above results for our notion of spacetime more generally? From a purely terminological point of view, we call a structure ‘spatiotemporal’ just if, at least

³⁸By extending the results of Asenjo and Hojman ([2017a], [2017b]) to other matter fields, we hope to make precise in a future technical paper whether such is, in fact, the case in rotating spacetimes of the kinds considered by Asenjo and Hojman.

to a satisfactory degree, it relates to what we think is (or rather would be) measured by rods and clocks (if present).³⁹

Now, the metric field of general relativity is usually considered to be spatiotemporal in this sense as it is expected that rods and clocks—if present—would at least to a satisfactory degree measure distances and times as given by the metric field.

The above findings, however, make a compelling case that Langevin clocks built from different matter fields may not agree on the interval along a given worldline.

If, however, all clocks differ significantly in their operational temporal readings, and the term ‘spacetime’ is associated with the behaviour of rods and clocks, then it becomes questionable whether the metric field of general relativity deserves the title of ‘spacetime’ at all—for this field ceases to play the operational role of codifying the behaviour of rods and clocks.

So, whether the metric field of general relativity can be conceived of as spatiotemporal is contingent not only upon the types of matter fields at play (cf. Read *et al.* [2018]), but also upon the nature of the solution to Einstein’s field equations under consideration: with respect to a flat metric, for example, matter fields may be used to realize a clock that measures the worldline interval linked to the metric—for in this case, scenario (A) obtains. With respect to other solutions, however, scenario (B) may obtain, and there may exist no straightforward means of procuring epistemic access to intervals of proper time as given by the metric field.

The current line of thought serves as a further argument (developing upon Brown [2005]; Read [unpublished]; Read *et al.* 2018) against the so-called geometrical approach to spacetime theories, according to which rods and clocks in general relativity invariably survey the metric field $g_{\mu\nu}$ (cf. Read [unpublished], Section 5).

Whereas, Brown ([2005], Section 9.5.2) presents various spacetime theories (such as the Jacobson-Mattingly theory (Jacobson and Mattingly [2001]), and Bekenstein’s ([2004]) tensor-vector-scalar theory (TeVeS)) in order to argue that, locally, metric and dynamical symmetries need not coincide, and so the metric field need not have chronogeometric significance, the

³⁹We take this to be the essence of Knox’s ([2013], [forthcoming]) ‘spacetime functionalism’. Of course, one may deny this view, and maintain (for example) that the metric field is inherently spatiotemporal. Though we do not find this view plausible, those readers who do embrace it (or another view diverging from that articulated above) are asked to consider the reflections in the body of this section in conditional form. Note also that our modal qualifications in the above mean that we are not committed in this paper to an extreme form of material operationalism.

results presented in this paper go further, for they demonstrate that, even in the presence of such symmetry coincidence, the metric field need not necessarily have chronogeometric significance. This constitutes further grist to the mill of the argument that the metric field does not possess its chronogeometric significance necessarily (as on the geometric approach), but ‘earns its spurs’ (to use Brown’s ([2005], p. 151) phrase) via considerations of the dynamical behaviour of matter fields.

5 Conclusion

While both Maudlin and Fletcher argue for the satisfaction of the clock hypothesis in the purely kinematical setting in which light rays necessarily propagate along null geodesics, in this paper we have called into question the extent to which such arguments remain sound, once it is recognized that light is a dynamical entity, which in generic spacetimes need not propagate at c . We have found particular trouble in this regard in the classes of rotating spacetimes considered by Asenjo and Hojman, in which regularity is lost. Consequently, light clocks cannot be regarded as ideal clocks in generic spacetime models.

These results lead to broad operational concerns: in certain spacetime models, it is not necessarily the case that we may have any operational access to the metric field. Such results also raise difficulties for the ‘geometrical approach’, for they provide further evidence that the metric field need not be surveyed by matter fields. In scenarios in which the metric field is not surveyed by rods and clocks built from matter fields, it is questionable whether this entity deserves the appellation ‘spacetime’ at all—thus, even in classical general relativity, spacetime may be significantly harder to come by than has hitherto been appreciated.

Acknowledgements

Very many thanks to Augustin Baas and Christian Wüthrich for constructive feedback on an earlier draft of this paper. We are also indebted to Wasilij Barsukow, Harvey Brown, Jeremy Butterfield, Juliusz Doboszewski, Sam Fletcher, Henrique Gomes, Nick Huggett, Vincent Lam, Dennis Lehmkuhl, Brian Pitts, Erhard Scholz and Nic Teh for valuable discussions. Niels Linemann is supported by a research grant provided of the Swiss National Science Foundation

(105212_165702); James Read by an Arts and Humanities Research Council doctoral scholarship at the University of Oxford, and a Senior Scholarship at Hertford College, Oxford. Tushar Menon and James Read are grateful to the University of Illinois at Chicago for visiting fellowships in the academic year 2017/18. These fellowships were funded by the John Templeton Foundation; the opinions expressed in this publication are those of the authors, and do not necessarily reflect the views of the John Templeton Foundation.

Tushar Menon

Faculty of Philosophy

University of Oxford

Oxford, UK

and

Department of Philosophy

University of Illinois at Chicago

Chicago, IL, USA

tushar.menon@balliol.ox.ac.uk

Niels Linnemann

Philosophy Department

University of Geneva

Geneva, Switzerland

niels.linnemann@unige.ch

James Read

Faculty of Philosophy

University of Oxford

Oxford, UK

and

Department of Philosophy

University of Illinois at Chicago

Chicago, IL, USA

james.read@philosophy.ox.ac.uk

Bibliography

- Asenjo, F. A. and Hojman, S. A. [2017a]: ‘Birefringent light propagation on anisotropic cosmological backgrounds’, *Physical Review D*, **96**, p. 044021.
- Asenjo, F. A. and Hojman, S. A. [2017b]: ‘Do Electromagnetic Waves Always Propagate along Null Geodesics?’, *Classical and Quantum Gravity*, **34**, p. 205011.
- Bekenstein, J. [2004]: ‘Relativistic Gravitation Theory for the MOND Paradigm’, arXiv:astro-ph/0403694.
- Bell, J. S. [2001]: ‘How to Teach Special Relativity’, in his *Speakable and Unspeakable in Quantum Mechanics*, Cambridge: Cambridge University Press.
- Brown, H. R. [2005]: *Physical Relativity: Space-time Structure from a Dynamical Perspective*, Oxford: Oxford University Press.
- Brown, H. R. and Read, J. [2016]: ‘Clarifying Possible Misconceptions in the Foundations of General Relativity’, *American Journal of Physics*, **84**, pp. 327–34.
- Brown, H. R. and Read, J. [2018]: ‘The Dynamical Approach to Spacetime Theories’, in E. Knox and A. Wilson (eds), *The Routledge Companion to Philosophy of Physics*, London: Routledge, pp. .
- Chiao, R. [unpublished]: ‘Superluminal phase and group velocities: A tutorial on Sommerfeld’s phase, group, and front velocities for wave motion in a medium, with applications to the “instantaneous superluminality” of electrons’, available at <arxiv.org/abs/1111.2402>.
- Dirac, P. A. M. [2012]: *Lectures on Quantum Mechanics*, Snowball Publishing.
- Earman, J. [1995]: *Bangs, Crunches, Whimpers, and Shrieks: Singularities and Acausalities in Relativistic Spacetimes*, Oxford: Oxford University Press.

- Fletcher, S. C. [2013]: ‘Light Clocks and the Clock Hypothesis’, *Foundations of Physics*, **43**, pp. 1369–83.
- Friedman, M. [1983]: *Foundations of Spacetime Theories*, Princeton, NJ: Princeton University Press.
- Garrison, J. C., Mitchell, M. W., Chiao, R. Y. and Bolda, E. L. [1998]: ‘Superluminal signals: causal loop paradoxes revisited’, *Physics Letters A*, **245**, pp. 19–25.
- Goldstein, H. [2011]: *Classical Mechanics*, Pearson.
- Jacobson, T. and Mattingly, D. [2001]: ‘Gravity with a Dynamical Preferred Frame’, *Physical Review D*, **64**.
- Knox, E. [2013]: ‘Effective Spacetime Geometry’, *Studies in History and Philosophy of Modern Physics*, **4**, pp. 346–56.
- Knox, E. [forthcoming]: ‘Physical Relativity from a Functionalist Perspective’, *Studies in History and Philosophy of Modern Physics*, available at <doi.org/10.1016/j.shpsb.2017.09.008>.
- Malament, D. B. [2012]: *Topics in the Foundations of General Relativity and Newtonian Gravitation Theory*, Chicago, IL: University of Chicago Press.
- Maudlin, T. [2012]: *Philosophy of Physics: Space and Time*, Princeton, NJ: Princeton University Press.
- Misner, C. W., Thorne, K. S. and Wheeler, J. A. [1973]: *Gravitation*, San Francisco, CA: Macmillan.
- O’Neill, B. [1983]: *Semi-Riemannian Geometry with Applications to Relativity*, Academic Press.
- Perlick, V. [2000]: *Ray Optics, Fermat’s Principle, and Applications to General Relativity*, Springer.
- Read, J. [unpublished]: ‘Explanation, Geometry and Conspiracy in Relativity Theory’.

Read, J., Brown, H. R. and Lehmkuhl, D. [2018]: ‘Two Miracles of General Relativity’, *Studies in History and Philosophy of Modern Physics*.

Schuller, F. P. and Witte, C. [2014]: ‘How quantizable matter gravitates: A practitioner’s guide’, *Physical Review D*, **89**, p. 104061.

Synge, J. L. [1960]: *Relativity: The General Theory*, Amsterdam: North-Holland.

Viaclovsky, J. [2011]: ‘Topics in Riemannian Geometry, University of Wisconsin lecture notes’, <www.math.uci.edu/~jviaclov/courses/865/_Fall_2011.pdf>.

Wald, R. M. [2010]: *General Relativity*, Chicago, IL: University of Chicago Press.