

Acquisition, not assembly: the classical manufacture problem for a boost-holonomy time machine, and the three routes to obtaining one

Jonathan Schemoul

LibertAI, France

js@libertai.io

October 27, 2026

Abstract

The matter side of a composite time machine is now explicit. A vacuum pp-wave core with a compact boost-identified direction can be wrapped in a dust envelope, matched to an asymptotically flat exterior, and evolved from regular compact initial data; every layer is closed-form, numerically gate-verified on the gatefast interior band $r \geq w + 0.0904$ and, at $a = 0.02$, gate-positive over the full C4 envelope layer $r > w$ after the one-knob collar retune, or cited from Ori’s construction, and no exotic matter is required. What remains is not a material but a global structure: the identification $z \sim z + L$, whose holonomy is a Lorentz boost. We argue that the known classical routes to manufacturing this organ are blocked or constrained. The sharpest classical statement is the Carroll–Farhi–Guth–Olum energy obstruction: in an open universe, initial data on an edgeless noncompact slice with timelike total momentum cannot evolve into the boost-holonomy (Gott) configuration. The theorem is rigorous in $2 + 1$ dimensions; the $3 + 1$ statement for our data class is not a published theorem, and we state that gap explicitly rather than closing it by assertion. Three independent results reinforce the conclusion, with different hypotheses: Geroch’s topology-change dichotomy, the Deser–Jackiw–’t Hooft/Cutler singular-constraint analysis of the two-defect route, and topological censorship, which makes the identification causally cloistered before the chronology horizon. The machine is legal general relativity; on the known classical routes, however, the identification is not installed by a causal process. The thesis is an argument with a $3 + 1$ proof gap, not a theorem. The identification must be *acquired*. We map three routes—primordial inheritance through foam, inflation, survival and survey; quantum nucleation; and reproduction inside an existing chronology-violating region, where Geroch’s price is already paid—and we separate what is established from what is scenario. The live empirical task is the survey: a bare relic is a massless, frequency-multiplicative lens, producing co-spatial spectral double images with a fixed ratio and no lens mass. We summarize the two executed channels, their nulls and their weak limits, and state what would change the verdict. No impossibility theorem and no detection are claimed.

1 Introduction and thesis

The program represented by this paper and its companions has reached an unusual point. On the side of matter and local geometry, the time machine is no longer a slogan: the composite initial data exist. A vacuum pp-wave core carrying a compact boost-identified direction is truncated on a

torus, matched through an explicit umbilical slice to a dust envelope, and handed off to Ori’s outer layers, which complete the data to an asymptotically flat exterior [1, 2, 37]. The construction is regular, compactly supported on the slice, and made of vacuum and dust; the weak, dominant and strong energy conditions hold on the gatefast interior band $r \geq w + 0.0904$ and, at $a = 0.02$, over the full C4 envelope layer $r > w$ in the retuned construction, including the innermost collar. The quantum stress on the chronology horizon of the 2005 core has been computed in the image series and is negligible at the design points examined [36]. The engineering ledger is astronomical but finite, and its scaling laws are known [37].

On the side of global structure, the situation is different. The working organ of the machine is not its dust, nor even its strong field. It is the identification

$$z \sim z + L \tag{1}$$

of the core, whose holonomy is a Lorentz boost by rapidity $\Lambda = L/2$ in the Ori convention [1, 38]. Before the chronology horizon the z -circles are spacelike; on it they are null; after it they are timelike and closed. The dust envelope exists to embed, match and complete this identified vacuum region. Remove the identification and the dust is an inert shell; keep the identification and the dust is a machine. The identification is the organ.

This paper asks the assembly question in its sharpest form: can the identification be *made*, from separated, independently prepared parts, by a classical process? Our answer is that the known classical routes cannot do it, and the argument is not a single theorem but a convergence of four classical results, each stated below with its hypotheses. The matter and the local geometry can be assembled; the global identification cannot be installed on those routes. It must be *acquired*: inherited from the quantum-gravity era, nucleated quantum-mechanically, or reproduced inside an existing chronology-violating region. The three routes are the subject of Sec. 4; the survey that tests the first and most conservative of them is the subject of Sec. 6.

The thesis is deliberately narrower than “time travel is impossible” and deliberately stronger than “we do not know how.” It is a statement about the boundary of what *building* can mean in classical general relativity. General relativity is a local theory: its field equations constrain the metric point by point and say nothing about global causal structure. That is why the theory is infested with exact solutions containing closed timelike curves, and why no local experiment can forbid them. But the same locality is why a global identification cannot be a local ingredient. One can specify it, embed it, stabilize it and compute its consequences; one cannot arrive at it by arranging ordinary matter on a simply connected slice, because the arrangement never contains the information that the identification requires. The no-go results make this intuition precise in the cases where they apply.

What is and is not claimed. We do not claim a theorem that no boost-holonomy spacetime can be created in $3+1$ dimensions. We do not claim that the Ori core is forbidden; on the contrary, it is a counterexample to any too-strong reading, which is exactly why the theorems must be quoted with their hypotheses. We do not claim that the survey has detected a relic; the executed channels return nulls. We do not claim a relic density bound; the current limits are conditional and weak. What we do claim is structural: the identification is a global quotient, not a local field; the classical routes considered here are blocked or constrained by known results; the remaining routes are acquisition routes; and the only one that is empirically accessible now is the survey.

Companion dependencies. Several results used here are established in companion papers that are in preparation and not yet public: the image-series stress on the 2005 core and the state-choice

band (P1) [36]; the explicit composite initial data and the mass law (P2) [37]; the quotient legality, monodromy and on-generator diagnostics of the traveling-wave core (P3) [38]; the gravitational-wave relic search (P4) [39]; and the Gaia radial-velocity census (P5) [40]. Where this paper relies on a companion result we say so; where it relies on a published theorem we give the hypotheses; where it relies on an estimate we label it an estimate. The program’s working notes are not cited as literature.

2 The machine and its organ

2.1 The composite

The core is the Ori (2005) vacuum plane-wave geometry [1]

$$ds^2 = dx^2 + dy^2 - 2 dz dT + [f(x, y) - T] dz^2, \quad f(x, y) = \frac{a}{2}(x^2 - y^2), \quad (2)$$

with z cyclic, $z \sim z + L$. The metric is Ricci-flat and $\det g = -1$ exactly. The surface $T = f$ is where $g_{zz} = f - T$ vanishes and the identification vector ∂_z becomes null; for $a \neq 0$ this surface is timelike, and it is the boundary ∂V of the chronology-violating region. The chronology horizon—more precisely, the local eikonal branch used in the companion analysis—is the null surface $T = F$ with $|\nabla F|^2 = f - F$, displaced below $T = f$ by $\delta_{\text{CH}} = a^2 \rho^2 + O(a^3)$ [36]. The distinction matters: the CTC boundary is not the horizon, and the stress is cut off at the horizon while still on its rising asymptote. The semiclassical status of such horizons is the subject of Hawking’s conjecture and the Kay–Radzikowski–Wald theorem [10, 11]; the companion computes the image-series stress in the eikonal/test-field approximation and finds it finite rather than divergent at the design points, with a true singularity remaining only on the single closed null geodesic. That is a local eikonal branch, not a proof that it is the global Cauchy horizon: the full tensor’s null energy is negative, opposite in sign to the positive energy-density proxy, and the value on ∂V is a formal past-horizon extrapolation rather than a boundary value. Olum’s analysis of the Ori–Soen spacetime is the relevant caution: in simplified models with the same causal structure he finds the Cauchy horizon compactly generated, and argues that building such a spacetime with normal matter may produce singular behaviour where the causality violation would otherwise occur [34]. The companion’s object is a local eikonal branch; neither compact generation nor a global Cauchy-horizon theorem is established for the composite.

The composite adds a dust envelope. On a z -invariant chart the data are matched across a torus to an envelope with an explicit profile family, then flattened onto Ori’s Layer-2 form and handed to his outer layers, which are core-independent and apply verbatim [2, 37]. The construction is initial data: a specified configuration of geometry and dust on a compact region of a spacelike hypersurface, which then evolves by itself. It is not a manufactured machine. The evolution into closed timelike curves is conditional on Ori’s existence theorem for the envelope, which the companion cites rather than re-proves.

The design has five dials, and they are decoupled [37]. The transverse radius ρ_0 sets the size and, through it, every mass and density. The parameter e sets when and where the CTC region forms, $t > e\rho^2$. The parameter a sets the wave strength, the ejection rate, and adversely the horizon stress. The slice depth $\tau_0 < 0$ sets the matching precision, the lead time and the mass budget. The period L sets the CTC loop scale and the wrap dynamics. The construction is scale-free: every length scales with ρ_0 , every time with ρ_0/c , every density with $c^4/G\rho_0^2$, every mass with $c^2\rho_0/G$.

2.2 The identification

The identification is not a field on the slice. It is a quotient of the core’s covering space: a deck transformation that acts on the (T, z) block as a Lorentz boost. In the $f = 0$ limit the (T, z) block has Riemann tensor identically zero and the core is Misner space [12]; the covering coordinates $\tau = 2\sqrt{-T}$, $\psi = z/2 + \frac{1}{2}\ln(-2T)$ map it to a flat two-dimensional block, and the deck map is the boost of rapidity $\Lambda = L/2$ [38]. The photon-transport holonomy measures $L/2$ to 10^{-14} , and the exact Misner gate $\sigma_n = 4|T_0|(\cosh(nL/2) - 1)$ confirms the dictionary. An earlier reading in which z itself was the rapidity ($\Lambda = L$) is wrong: z is twice the covering boost angle up to a logarithm. We use $\Lambda = L/2$ throughout. Quotient legality is not automatic: Rieger’s deck-equivariance classification of Misner-type spacetimes is the closest external treatment of when such a quotient is causally acceptable [22], and the state-choice problem on a time-periodic quotient has a precedent in the invariant vacuum of Melton et al. [23]. Neither replaces the construction; both bound what a quotient can consistently mean.

Three consequences follow, and they organize the rest of the paper.

First, the holonomy is discrete and global. In the $f = 0$ limit a local observer inside the core sees flat space; for $a \neq 0$ the local curvature is weak and does not by itself reveal the quotient. The identification is visible only to a curve that winds around the z -circle. There is no local density, current or field strength that can be dialed to “turn on” the holonomy. The parameters (a, e, L) are born with the quotient; they are not assembled from local data.

Second, the identification is causally hidden before the chronology horizon. The z -circle has proper length $L\sqrt{|e\rho^2 - T|}$; it is spacelike for $T < e\rho^2$, null on the CTC boundary, and timelike only after it. A causal curve cannot wind around a spacelike circle and return to its starting point. Before the horizon, therefore, no causal probe can complete a closed loop around the identification and return to the same event; the holonomy is hidden from generic probes, and only a chord that enters the thin entry shell can cross the seam. This is the covering-space form of the topological-censorship statement discussed in Sec. 3.4, and it is why the machine is legal general relativity: it hides its holonomy from generic probes until the chronology horizon opens it.

Third, the identification is what the matter does *not* supply. The dust envelope supplies mass, matching, asymptotic flatness and a well-posed initial-value problem. It does not supply the quotient. A dust shell around a non-identified vacuum region is a perfectly good spacetime; it simply never produces closed timelike curves. The organ is the quotient, and the quotient is the thing that has to be acquired.

2.3 What the matter side achieves

It is worth stating the positive result plainly, because the rest of the paper is about what it does not buy. The composite initial data are explicit and gate-verified. The canonical design is $(\rho_0, e, a, \tau_0, L) = (0.5 \text{ m}, 0.1, 10^{-4}, -1 \text{ m}^2, 1)$. The assembly-to-CTC lead time is $\sqrt{|\tau_0|}/c \simeq 3.3 \text{ ns}$; at $t = 2e\rho_0^2$ the loops are about 16 cm; the horizon stress at $a = 10^{-4}$ is $\sim 8 \times 10^{-9} \text{ J m}^{-3}$, more than fifty orders of magnitude below back-reaction relevance; the a -floor at metre scale is $\sim 2 \times 10^{-17}$, so the canonical value sits thirteen orders above it. The exact vacuum-tube (RW) integral to $R_1 = 4 \text{ m}$ gives $6.45 \times 10^{26} \text{ kg}$, about $0.34 M_J$, comparable to the compact-object wall estimate $c^2\rho_0/2G \simeq 3.4 \times 10^{26} \text{ kg}$. That is the tube alone, not the machine mass. The per-region ledger gives a C4 open layer of $1.33 \times 10^{27} \text{ kg}$ and a flattening band of $8.9 \times 10^{26} \text{ kg}$, so the full canonical envelope is $2.9\text{--}5.4 \times 10^{27} \text{ kg}$, or $1.5\text{--}2.8 M_J$ ($480\text{--}900 M_\oplus$), the upper bound applying Ori’s $R_2 > 2m$ condition at $R_2 = 8 \text{ m}$. The full envelope is therefore ~ 8 times the compact-object wall, not at it. The dust is ordinary matter in composition, but the canonical corner is not ordinary in density:

$\varepsilon_0 = 4.0 \times 10^{25} \text{ kg m}^{-3}$ at the hand-off, with a peak of $1.17 \times 10^{26} \text{ kg m}^{-3}$ —about 1.7×10^8 times nuclear saturation at the hand-off and 5×10^8 at the peak. The density scales as k_0^2 and becomes 778 kg m^{-3} at the deep rung $k_0 \simeq 2.2 \times 10^{-12}$; at fixed ρ_0 and L_z the total dust mass scales as k_0 , not k_0^2 , giving $\sim 1.7 \times 10^{16} \text{ kg } L_z$ at that rung. The earlier “tens of tons” figure applied the k_0^2 law to a volume that grows as $1/k_0$; it is retracted. The correct statement is that the machine is specified, not that it is cheap: the matter budget is Jupiter-scale at the canonical corner ($1.5\text{--}2.8 M_J$), and the deep rung is ordinary-density but asteroid-scale.

The Xavier (2026) traveling-wave core improves the acquisition picture without changing its logic [21, 38]. Its identification is still the global organ that has to be acquired, but the core can be found *off*—no CTCs, no horizon, locally flat in the $f \approx 0$ limit—and armed on demand by sending the ignition wave around the loop. The companion reports that the winding compression of its horizon generators may be one, which would make the Hawking divergence mechanism inapplicable and the core effectively silent even when armed; that report is unpublished and contested, and its state-choice trichotomy is not exhaustive, so we cite it qualitatively rather than as an established property. If it holds, a natural machine can wait, quiet and invisible, until it is switched on. This does not make it easier to manufacture. It makes it easier to *inherit*: the relic need not be an armed time machine, only a quotient pocket waiting for a trigger. Scalar quasisresonance is the current instability-side comparator [28]; the vacuum core’s stability is the companion’s calculation, not an assumption.

2.4 What “manufacture” would mean

The assembly question is not “can the final data exist?” They can; they are written in Eq. (2) and its envelope. The question is “can a classical process take data that do not carry the identification and produce data that do?” Three readings of that question are possible, and it is important not to conflate them.

(i) *Local assembly.* Take ordinary matter and geometry on a simply connected slice, arrange it, and let it evolve. This is the reading relevant to a laboratory or a construction site. It is the reading of the Carroll–Farhi–Guth–Olum theorem and of Geroch’s dichotomy.

(ii) *Global preparation.* Specify the identification as a boundary condition or a quotient from the outset, then prepare the matter that completes it. This is what Ori’s construction does, and what the companion’s initial data do. It is not manufacture in reading (i); it is specification. The identification is an input, not an output.

(iii) *Acquisition.* Obtain a region that already carries the quotient—from the early universe, from a quantum nucleation event, or from an existing machine—and then wrap it. This is the only reading that requires no new identification-creating mechanics; Routes A and B are non-circular as first-machine routes, while Route C is a bootstrap that presupposes a machine.

The no-go results below are statements about reading (i). They do not forbid reading (ii), which is why the machine exists as a solution; and they motivate reading (iii), which is the program’s actual position. The distinction is the paper’s thesis in one paragraph.

3 Why classical manufacture fails

3.1 The energy obstruction: Carroll–Farhi–Guth–Olum

The sharpest classical statement is the energy-momentum obstruction of Carroll, Farhi, Guth and Olum [3, 4]. In $2 + 1$ dimensions they prove that Gott time machines cannot exist in an open universe with timelike total momentum; equivalently, initial data on an edgeless, noncompact,

spacelike hypersurface whose total momentum is timelike do not have enough energy to evolve into the boost-holonomy (Gott) configuration. The theorem is rigorous in $2 + 1$ dimensions, and it is the closest thing in the literature to a general creation no-go.

The hypotheses matter. The slice must be edgeless and noncompact; the total momentum must be timelike; the conclusion concerns the Gott configuration, a pair of moving defects whose combined holonomy is a boost. In a closed universe the theorem’s own exception appears: Gott pairs, although not CTCs, can be produced from the decay of stationary particles [3, 4]. That exception is not a door for construction. It requires the entire connected spacetime to be spatially compact, a fact about the universe rather than a laboratory operation; the Planck constraint on spatial curvature, $\Omega_k = 0.001 \pm 0.002$, does not make our universe closed at the relevant scale; and manufacturing compactness around the laboratory is topology change on the largest possible scale, which returns us to Geroch. A compact slice around the apparatus does not help either: a compact spatial slice cannot carry a well-defined nonzero conserved timelike total momentum, so the data leave the theorem’s class only by not existing as specified.

Our composite’s matter and local geometry belong to the theorem’s data class—regular, asymptotically flat, dust at rest, timelike ADM momentum—once the identification is treated as the global datum it is. The identification itself is not a local datum on that slice; it is a quotient of the core’s development. The theorem therefore does not forbid the Ori data, which already carry the quotient. It forbids *reaching* those data from data that do not. That is the assembly question, and the theorem answers it in $2 + 1$ dimensions.

The $3 + 1$ extension is the subject of Sec. 5. We flag here only the logical status: the $2 + 1$ theorem is a theorem; the $3 + 1$ statement for our data class is a gap, not a result. The thesis is an argument with a $3 + 1$ proof gap, not a theorem.

3.2 Topology change: Geroch

Geroch’s topology theorem [5] is the second classical constraint. In a globally hyperbolic spacetime with compact Cauchy surfaces, all Cauchy surfaces are homeomorphic: the spatial topology cannot change. Equivalently, a topology-changing evolution that remains globally hyperbolic, singularity-free and energy-condition-respecting does not exist; topology change must be paid for by closed timelike curves, a singularity, or a violation of the energy conditions.

Strictly, Geroch’s theorem is stated for compact Cauchy surfaces. The composite’s asymptotic slice is noncompact, so the compact theorem does not apply verbatim, and topology change at infinity is not at issue here. The topology change relevant to the identification is localized in the core’s development. We therefore use the dichotomy as a structural constraint on the transition, not as a direct theorem about the composite. On the reading in which the identification is a compact direction of the spatial sections, creating it from data that do not carry it is a topology-changing operation in the sense of the theorem. The three prices are each unattractive. Closed timelike curves are the product one is trying to make, so paying with them is circular. A singularity destroys the machine. An energy-condition violation requires matter that no known laboratory source supplies; the quantum-inequality literature bounds the available negative energy, and the bounds are far below the required macroscopic configuration (see, e.g., Ref. [19] for the averaged inequalities and Ref. [20] for the current state-independent null-energy bounds). One can also try to pay by hiding the change behind an event horizon—a black-hole interior already contains a singularity, so the price is formally pre-paid—but the product is then interior to the horizon forever (a one-way membrane), can never be exported, surveyed or operated, and the only singularities humans can access are inside astrophysical black holes with no controllable interior. The loophole is formally real and causally worthless; Minguzzi’s area theorem, together with the stable dominant

energy condition in non-empty spacetimes with source content, closes the horizon-first ordering trick more generally [33]. Andréka et al. sharpen what a causal extension can guarantee: they construct an extendible spacetime without CTCs whose every extension contains them [24], which is the precise sense in which a local patch can be causally benign while its global completion is not.

The theorem does not forbid the machine’s existence. The machine’s spacetime is not globally hyperbolic; it has a chronology horizon, and the horizon is exactly the price. What the theorem forbids is a topology-changing *transition* from a globally hyperbolic pre-machine phase to the machine phase without one of the three prices. That is the assembly transition. The distinction between forbidding the final state and forbidding the transition is the same distinction as in Sec. 2.4, and it is the reason the machine can exist while remaining unbuildable.

3.3 The two-defect route: Deser–Jackiw–’t Hooft and Cutler

The simplest classical route to a boost holonomy is a pair of moving cosmic strings, Gott’s construction [9]. Two parallel strings moving past each other in opposite directions produce a conical identification plus a boost; a path circling both can close as a closed timelike curve when

$$\gamma v \alpha \gtrsim 1, \quad \alpha = \frac{8\pi G\mu}{c^2}, \quad (3)$$

where μ is the string tension. The route is attractive because it needs no exotic matter: motion and topology suffice.

It is killed twice. Deser, Jackiw and ’t Hooft [6] showed that physical cosmic strings do not generate closed timelike curves: the two-string spacetime with CTCs cannot arise from regular initial data, and geodesically complete Gott pairs exist only in spacetimes with closed spatial topology. Cutler [7] analyzed the global structure and found the initial-value constraints singular: no acceptable Cauchy hypersurface generates the Gott spacetime. The two-defect assembly is therefore not merely hard; it is constraint-singular.

The resource gap makes the point independently. The tension of a field-theory defect is set by its symmetry-breaking scale, $G\mu/c^2 \simeq (\eta/M_{\text{Pl}})^2$. A GUT-scale string has $G\mu/c^2 \sim 10^{-6}$, requiring $\gamma v \gtrsim 4 \times 10^4$; an LHC-scale defect has $G\mu/c^2 \sim 7 \times 10^{-31}$, short by about 29 orders at $\gamma v \sim 1$; an electroweak string is shorter still; a laboratory Abrikosov flux tube is shorter by about 49 orders. The Gott condition is scale-free, so a weaker string does not give a metre-scale machine; it pushes the CTC region out to astronomical distances, the opposite of a compact core. Velocity cannot rescue the tension: the required Lorentz factor for a sub-GUT string is astronomical. Observationally, the required $G\mu/c^2 \gtrsim 10^{-2}$ is excluded by 8–12 orders by CMB, lensing and pulsar-timing data [29, 30]; pulsar timing has revived strings as a possible gravitational-wave source at $G\mu/c^2 \sim 10^{-10}$, which is still eight orders too weak for Gott.

One string alone gives a rotation-type spatial holonomy, not the boost holonomy that the core embeds. The boost needs the pair, and the pair’s CTC region is an extended exterior region, not a compact identified core with a chronology horizon. The two-defect route is therefore both resource-excluded and structurally mismatched. BEC vortices and superconductor flux tubes carry real topological defect structure in the field and real analogue-gravity kinematics, but no spacetime identification; they are analogues, not the organ.

3.4 Topological censorship and causal cloistering

Topological censorship [8] is usually stated as a statement about wormholes: under the null energy condition, in a globally hyperbolic spacetime with a Cauchy surface, asymptotic flatness and a

simply connected asymptotic region, every causal curve from past to future null infinity is homotopic to a trivial curve, so nontrivial topology is hidden behind horizons and cannot be probed from infinity. The hypotheses hold in the pre-machine phase, not for the full machine, which is not globally hyperbolic. In the present context the theorem plays a different and more precise role. Before the chronology horizon the z -circle is spacelike; no closed causal curve can wind around it and return to the same event. The identification is therefore causally cloistered from the envelope: a generic causal probe does not enter the seam. The machine hides its holonomy from generic probes until the chronology horizon opens it.

This is not an independent no-go theorem, and we do not present it as one. It is a compatibility statement that sharpens the others. It explains why the machine is legal: topological censorship is not violated, because the topology is hidden from generic probes. It also explains why the organ cannot be installed by a causal process: a process that never traverses the loop cannot use the loop to bootstrap its own creation. The same fact is the survey’s central obstacle. A bare relic is not a lens in the ordinary sense; it is a region that light can enter only through a thin entry shell, and the shell is star-poor for small pockets. We return to this in Sec. 6.2.

3.5 What the four results do and do not establish

The four results are independent in their hypotheses, and none of them is a $3 + 1$ no-go theorem for every boost holonomy. It is worth saying this plainly, because overstating the case would be a disservice to the thesis.

- CFG/CFGO is a rigorous $2 + 1$ theorem about Gott configurations from open-universe data with timelike total momentum. It does not by itself cover the Ori core, which is Misner-type and already carries the identification, and its $3 + 1$ extension to our data class is not published.
- Geroch’s dichotomy is a classical statement about topology change in globally hyperbolic spacetimes. It does not forbid a spacetime with a chronology horizon; it forbids a topology-changing transition without a price. The machine pays the price by having a horizon.
- DJHT/Cutler is specific to the two-defect route. It is decisive for that route and silent about others.
- Topological censorship is a statement about causal probing. It does not forbid the identification; it hides it from generic probes until the horizon opens.

What the four results do establish, collectively, is a structural conclusion: the classical routes considered here for *creating* the identification either violate a constraint theorem, require a resource that is observationally excluded, or produce the identification only by already assuming it. The conclusion is not a theorem; it is a statement about the boundary of the classical construction space considered here, and it is strong enough to change the question from “how do we build it?” to “how do we get one?” The rest of the paper follows that change.

4 The three acquisition routes

If the identification cannot be manufactured, it must be acquired. There are exactly three logical possibilities: it was there from the beginning; it was created by a non-classical process; or it was created by a process that operates where the classical constraints do not apply. We call them Route A (primordial inheritance), Route B (quantum nucleation) and Route C (reproduction).

They are not equally plausible. Only Route A is classical-compatible and non-circular: its first step begins in the quantum-gravity era, so it is not classical physics, but it requires no new identification-creating mechanics.

4.1 Route A: primordial inheritance

The pipeline is

$$\text{foam} \longrightarrow \text{inflation} \longrightarrow \text{survival} \longrightarrow \text{survey} \longrightarrow \text{capture} \longrightarrow \text{wrap} \longrightarrow \text{armed machine.} \quad (4)$$

Every stage after “survey” is engineering that the companion papers have already specified; the first four are physics that must be argued.

Foam. In the quantum-gravity era the metric fluctuations are $O(1)$ and the topology fluctuates. Wheeler’s spacetime foam is the standard picture, and the modern review by Carlip collects the evidence that geometry and topology fluctuate at the Planck scale [13]. Locally quotient pockets with all holonomies—including boost-like ones—are natural foam products unless the full quantum theory forbids them. No known theory does. This is the weakest link in the chain and we label it as such: it is a plausibility argument, not a theorem. The foam may be dominated by configurations with no stable quotient pockets; the measure is unknown.

Inflation. Inflation stretches a Planck-scale pocket to any size. The standard statement that expansion preserves spatial topology applies to the spatial sections of a homogeneous cosmology [14]; it is not by itself a statement about a boost-holonomy quotient. The pocket’s identification mixes time and the compact z direction, the pp-wave metric is not a comoving perfect fluid, and inflation is not a symmetry of the quotient. The dimensionless machine parameters (a, e, L) are born with the pocket, but whether the boost-identified vacuum region survives inflation with its holonomy intact is an argument that has not been supplied: one must show that the quotient is not destabilized, diluted, or converted to ordinary spatial topology by the inflationary expansion and the reheating that follows. We treat this as an open conceptual gap in Route A, and we do not infer it from spatial-topology preservation. The foam samples the machine window $a \in (2 \times 10^{-17}, \sim 0.1)$ (the a -floor of the engineering ledger), $0 < e \lesssim 1/4$, $L = O(1)$. A pocket that is Planck-scale before inflation can be metre-scale or astronomical today, depending on its size at the end of inflation and on the subsequent expansion. The distribution of relic sizes is not predicted by the standard inflationary model; it depends on when the pocket exits the horizon and on the topology-changing dynamics, which is exactly the unknown physics.

Survival. A bare relic is vacuum, nearly Minkowski locally, and gravitationally negligible. The companion’s stability analysis finds that transverse perturbations are ejected, the closed null geodesic is attracting, and the core self-cleans [36, 38]. In the traveling-wave core the companion reports a winding compression of one, which would make the Hawking divergence mechanism inapplicable and the core effectively silent even when armed; that report is unpublished and contested, and we cite it only qualitatively [38]. The survival argument is therefore not “nothing happens for 13.8 Gyr”; it is “the known instabilities do not operate on this object.” That is a stronger and more honest statement. It is still conditional: it assumes no destabilizing coupling to ambient matter or radiation, and it assumes that the companion’s linear and leading nonlinear analysis extends to cosmological timescales.

Survey. A quotient pocket is a lensing object, but not a mass lens. It produces co-spatial double images with a fixed frequency ratio and no deflection. This is a search problem, and it is the live empirical task of Sec. 6. The relic may be off (no CTCs, no horizon) and still produce the seam signature for light that enters the entry shell; the trigger epoch opens the image cascade. The survey is the only stage of Route A that can be executed now.

Capture and wrap. Once a relic is found, the program takes over. The C4-envelope construction, the flattening hand-off and Ori’s outer layers convert a relic core into an armed machine with a pressureless dust envelope. The mass budget is the full canonical envelope, $1.5\text{--}2.8 M_J$, whose gravitational radius is comparable to its outer radius: it is a compact object, not a dust halo. The arming time is nanoseconds at the canonical corner; the loop size is set by L . The capture step itself—approaching a relic, determining its parameters, and building the envelope around it—is the least developed part of the engineering, and we do not claim otherwise. The point is that the *physics* after acquisition is specified.

4.2 Route B: quantum nucleation

The classical kills do not bind quantum creation. A Euclidean tunneling process can nucleate an identified pocket, in the same sense that Farhi–Guth–Guvén tunneling can nucleate a baby universe [15, 16]. The amplitude is expected to be exponentially suppressed at anything above the Planck scale, but at the Planck scale Route B *is* Route A’s first step. No known theorem forbids it. The relevant question is not “is it forbidden?” but “is the rate nonzero in the full theory?” and that is a question for quantum gravity.

The present-day ledger is not close. Creating a Planck-scale packet requires the Planck energy per particle, $\sim 1.95 \times 10^9$ J; the LHC delivers $\sim 2.2 \times 10^{-6}$ J per proton, fifteen orders short. The energy density at the event would need to be $\sim M_{\text{Pl}}^4 \simeq 5 \times 10^{113} \text{ J m}^{-3}$; heavy-ion collisions reach $\sim 8 \times 10^{35} \text{ J m}^{-3}$, seventy-eight orders short. A false-vacuum instanton at the Planck scale needs $\rho \sim M_{\text{Pl}}^4$; the electroweak false vacuum is $\sim 10^{44} \text{ J m}^{-3}$, about sixty-eight to sixty-nine orders short. These numbers are illustrative, not a theorem: the shortfall is not a small rate but a different regime, and every order is exponential in the wrong direction. The honest statement is that Route B is a route for nature, not for any present or near-future apparatus.

There is also a structural lesson from Farhi and Guth [15]: creating an inflationary pocket from regular laboratory initial data fails on energy grounds before the quantum amplitude is even discussed. Even a successful vacuum-decay catalysis would produce a true-vacuum bubble of ordinary topology, not an identified quotient pocket with boost holonomy. The wrong product is a more fundamental problem than the suppression.

4.3 Route C: reproduction

Geroch’s price for topology change is CTCs, a singularity, or an energy-condition violation. Inside a working machine’s CTC region the price is already paid: the spacetime is already non-globally hyperbolic, so the topology-change theorem does not apply. A civilization with one machine can, in principle, perform identification-creating operations inside its CTC region and carry new identified pockets out; transporting a pocket out is not itself a topology change. Machines reproduce. Gott and Li’s self-creation scenario is the cosmological version of this logic [35]: a region of CTCs bounded by a Cauchy horizon from which the universe inflates. This is the bootstrap route, and it is logically the most economical: it builds every machine after the first, and it requires no new physics beyond the existence of one machine.

Route	Mechanism	Status	Weak link
A	Foam $\rightarrow$ inflation $\rightarrow$ survival $\rightarrow$ survey $\rightarrow$ capture	Classical-compatible (begins in the quantum-gravity era), non-circular; survey executable now	Foam abundance, inflation/quotient survival and survival are not theorems
B	Euclidean nucleation of an identified pocket	No known prohibition; exponential suppression	15–78 orders beyond present apparatus; may produce ordinary topology
C	Reproduction inside an existing CTC region	Geroch’s price pre-paid; consistency models exist	Requires a first machine; operation inside the horizon is not engineered

Table 1: The three acquisition routes. Only Route A is classical-compatible and non-circular. The table classifies mechanisms, not probabilities.

The consistency layer for what can be done through CTCs is already studied. Deutsch’s fixed-point prescription [17] defines evolution near closed timelike lines; postselected CTCs and self-consistency conditions give alternative prescriptions; the program’s simulations of these models find no logical obstruction. The recent multiple-histories framework of Shoshany and Wogan [18] drops the self-consistency constraint in favor of a probabilistic ensemble and resolves the usual paradoxes without it. None of these is a construction of the identification-creating operation. They are consistency checks on the logic of operating inside a CTC region.

Route C has three honest limitations. First, it is circular as a first-machine route: it presupposes a machine. Second, the CTC region is causally hidden until the chronology horizon opens, so the operation must be performed inside the horizon, where the semiclassical stress and the stability of the core are the companion’s open problems, not solved engineering. Third, the transported pocket must be stable and must not be destroyed by the same backreaction that threatens the host. These are not reasons to discard the route; they are reasons to classify it as a bootstrap, not a first cause.

4.4 Comparison and the bootstrap

Table 1 summarizes the three routes. Route A is the only classical-compatible, non-circular route; its weak links are the foam premise, the survival of the boost quotient through inflation, and the survival argument. Route B is allowed by known physics but is exponentially out of reach and produces the wrong product in the simplest catalysis scenarios. Route C is the most economical in the long run but requires a first machine. The logical structure is a bootstrap: the first machine must come from A or B; every later one can come from C. The survey tests A, because A is the only route whose first step is accessible to present-day instruments.

5 The 3+1 CFG gap and what would reopen it

5.1 The status of the theorem

The Carroll–Farhi–Guth–Olum result is often quoted as if it were a $3 + 1$ no-go theorem. It is not. The rigorous statement is $2 + 1$ -dimensional: Gott time machines cannot exist in an open universe with timelike total momentum [3, 4]. The proof uses the $2 + 1$ -dimensional momentum and the conical geometry of point masses; it does not directly apply to a $3 + 1$ pp-wave core with a dust

envelope. The $3 + 1$ statement for our data class is expected to hold by structural analogy; no counterexample is known, but it is not a published theorem. We treat it as a gap.

This is not a technicality to be waved away. A referee is entitled to ask: “Where is the $3 + 1$ theorem?” The honest answer is: it does not exist in the form we would need. The energy accounting that makes the $2 + 1$ proof work—the total momentum of an open universe is timelike, and the Gott configuration requires a momentum configuration that such data cannot supply—is expected to generalize, but expectation is not proof. The gap is the only energy no-go for the Ori data class; Geroch, DJHT/Cutler and topological censorship are different kinds of support and do not independently forbid that class. The thesis is therefore an argument with a $3 + 1$ proof gap, not a theorem.

5.2 What would reopen it

A rigorous $3 + 1$ counterexample to the CFG momentum inequality for a non-Gott data class would reopen a proof gap. It would not, by itself, give a build route: one would still have to exhibit a classical process that creates the identification without violating Geroch, DJHT/Cutler or topological censorship. But it would remove the strongest energy obstruction and would force the program to take the construction route seriously.

A consistent $3 + 1$ Gott-pair formation from open data with timelike total momentum would be stronger: it would be a counterexample to the structural analogy, not just to the $2 + 1$ theorem. It would also have to evade DJHT/Cutler’s singular-constraint analysis, which is independent of the dimension.

A physically realizable loophole in the hypotheses would reopen the theorem itself. The candidates are the closed-universe exception, a compact slice, and non-timelike total momentum. The first requires the whole universe to be closed; the Planck curvature constraint does not support it, and manufacturing compactness is Geroch’s problem on the largest scale. The second cannot carry a nonzero conserved timelike momentum. The third is not a property one can choose for a machine at rest in an asymptotically flat region.

A quantum-gravity mechanism would not reopen the classical gap; it would bypass it. That is Route B, and it is a different route, not a counterexample to CFG.

5.3 What would not reopen it

Several “loopholes” that circulate in informal discussions do not survive scrutiny. The closed-universe exception gives Gott pairs, not CTCs, and requires a cosmological fact we do not have. A compact slice around the apparatus cannot carry the momentum the theorem assumes. Hiding the topology change behind a horizon is formally allowed by Geroch but causally worthless: nothing crosses outward, so nothing can be exported or operated. A sub-GUT string pair does not evade CFG; it makes the resource gap worse and pushes the CTC region to astronomical distances. A modified-gravity embedding does not help unless it changes the energy accounting in a controlled way; the current modified-gravity CTC literature finds that chronology-violating regions tend to survive rather than disappear, which is the opposite of a construction route [26, 27].

5.4 Why the gap is a formality but a real one

We call the gap a formality because no known physics exploits it, and we call it real because a conceptual paper about classical manufacture must be explicit about which part of the claim is theorem and which is analogy. The paper’s thesis is not “CFG proves that our machine cannot be built.” The thesis is: the identification is a global quotient; the classical routes considered here

for creating it are blocked or constrained; the $2 + 1$ theorem is the sharpest energy statement; the $3 + 1$ version is open; and the acquisition routes are what remain. That thesis is an argument with a $3 + 1$ proof gap, not a theorem; a stronger claim would not survive the gap.

6 The survey and capture pipeline

6.1 The signature

A bare relic is, in the $f \approx 0$ limit, a locally flat region whose covering space is quotiented by a boost. It has no local curvature and no Newtonian mass, but it has a global holonomy: a path that crosses the seam once returns with a Lorentz-transformed momentum. For a photon crossing the seam once, the frequency ratio between the wound and direct images is

$$r \equiv \frac{\nu_{\text{wound}}}{\nu_{\text{direct}}} = \cosh \Lambda - n_z \sinh \Lambda, \quad (5)$$

where $n_z \in [-1, 1]$ is the direction cosine along the boost axis. The ratio lies between $e^{-\Lambda}$ and $e^{+\Lambda}$, with the endpoints attained on axis, and for axial photons it is exact:

$$r(n_z = +1) = e^{-\Lambda}, \quad r(n_z = -1) = e^{+\Lambda}. \quad (6)$$

A source therefore shows a *co-spatial spectral double image*: one position on the sky, two copies of one spectrum at a fixed ratio, with no lens mass and no source dependence. Off axis the wound image is additionally tilted by up to $\arctan(\sinh \Lambda)$, which is 27.5° for $\Lambda = 0.5$; the observable is chiefly spectroscopic, not astrometric. A multi-wound exit gives ratios $e^{\pm n\Lambda}$: integer powers of one fixed base, still a fixed-family signature. There is also a time echo of order $2\tau_{\text{wall}} \sinh(n\Lambda/2)$, about 3.4 ns per circuit for the canonical metre pocket.

For a gravitational wave the same boost multiplies the frequency of the whole waveform by $e^{\pm\Lambda}$. The detector-frame chirp mass is inferred from the frequency and its derivative, so a common scaling of all frequencies becomes a common scaling of the inferred source-frame masses: both masses are multiplied by the same factor, and the mass ratio and χ_{eff} are unchanged. The transverse-traceless amplitude h is invariant, so the image is not magnified. But $h \propto M_c^{5/3} f^{2/3}/D_L$: with h , the mass ratio and χ_{eff} fixed, the inferred luminosity distance scales with the inferred masses, $D_L \propto M_c$ (M_c/D_L invariant), rather than staying equal. A true holonomy pair therefore lies on a mass–distance scaling degeneracy, not at a common D_L . For the axial channel the two images are co-spatial; off axis the wound image is tilted, as for photons. The invariant discriminators that no ordinary lens can fake are the fixed ratio $e^{\pm\Lambda}$ and its integer-power ladder, the unchanged mass ratio and χ_{eff} , and the common scaling of mass and distance ($D_L \propto M_c$); an ordinary lens preserves the source frequencies and magnifies the flux instead. The earlier GW cut of Ref. [39] required equality of D_L ; it is therefore a null on that cut, not on the scaling signature derived here. The corrected M_c/D_L cut has since been run and is also null, as quantified below.

6.2 Causal cloistering and the two windows

The pre-horizon pocket is not a generic lens. Its world-tube is the Rindler band $0 < \tau < \tau_{\text{wall}}$. A null chord from a background star crosses the tube’s spatial location at covering radius $\tau \sim \tau_{\text{now}}$, far outside the band, and never enters. Entry requires the chord to graze the birth-cone null surfaces $W = \pm Z$ inside the tube: a shell of thickness $\sim r_{\text{wall}}$ around those surfaces. The star count in the shell, using a local stellar density of 0.004 stars/ly³ and a disk-shadow volume, is 3.8×10^{-51} for a

0.5 m pocket, 1.6×10^{-17} for a 1 AU pocket, and 1.6×10^{-8} for a 1000 AU pocket. The OFF bare relic is invisible to a generic stellar survey.

There are two windows where light does enter. The first is a large pocket: if τ_{wall} is astronomically large, the entry shell is star-populated and the pair channel opens. The second is the trigger: at $t > e\rho^2$ the z -circles turn timelike, winding becomes free, and the image cascade opens—the classical Misner pile-up, many images of one source at geometrically compressed frequencies crowding the horizon direction. The detectable epoch is the CTC epoch, or a pocket with $\tau_{\text{wall}} \gtrsim 0.1$ ly. This is the covering-space form of the topological-censorship statement, and it is the central obstacle to the survey.

The numbers above are order-of-magnitude estimates, not a theorem. They depend on the local stellar density, on the disk-shadow geometry, and on the assumption that the seam is sharp. They are enough to tell us that the EM and GW channels are large-pocket or trigger-epoch channels, not all-sky point-source searches.

6.3 The two channels and the executed cuts

The EM channel uses stellar radial velocities. For the small rapidities accessible to stars, $\Lambda \sim 10^{-4}$ – 10^{-3} , the wound image is a second spectrum at the same sky position with a radial-velocity offset $\Delta\text{RV} = c(e^{\pm\Lambda} - 1)$. The Gaia DR3 census has been executed on the 33,812,183 sources with published radial velocities, tiled by source identifier into 768 level-3 HEALPix tiles with level-4/5 splits, producing 1,032 pair CSVs and 16,358 unordered pairs at separations $\leq 1''$ [40]. An astrometry match (parallax and proper motion consistent at 3σ) labels 7,217 pairs (44.1%); 1,177 of these have $|\Delta\text{RV}| > 3\sigma$, and 939 satisfy the candidate-zone cut $|\Delta\text{RV}| > 30 \text{ km s}^{-1}$ and $> 3\sigma$. Requiring all six astrometric values leaves 10,318 pairs, of which 1,177 are matched, 171 above 3σ , and 135 in the candidate zone; the permissive and strict 1,177 are different samples, and their equality is a coincidence of the data. The matched pair ΔRV distribution has median 7.4, 90th percentile 41.6, 99th percentile 126.8, and maximum 489.5 km s^{-1} ; the strict distribution is similar. The axial- Λ window is 1.0×10^{-4} to 1.6×10^{-3} , complementary to the GW channel’s $0.1 < |\Lambda| < 3$.

The shared-base family test returns a largest value of 173 of 939 permissive candidates (24 of 135 strict) at the low- Λ threshold $\Lambda = 1.29 \times 10^{-4}$. This is not a cluster: the tolerance is a half-width $0.1b$, a full 7.74 km s^{-1} window at the maximum, the local density is 22.4 candidates per km s^{-1} , and the smooth expectation is 173.4 against 173 observed (strict: 24.0 against 24). The statistic is CDF-dominated, exactly as in the GW channel. A composition-matched bootstrap gives $p_{\text{excess}} = 0.737$ (permissive; null median 179, 90% 164–197) and 0.969 (strict; 28, 24–34), with smooth models and controls agreeing; no excess survives. We report no significant signal, no detection, and no exclusion, and we claim no density bound. The candidate set is consistent with an astrophysical background: co-moving binaries and multiples, line-of-sight superpositions, and composite spectra. Gaia DR3 exposes mean RVS spectra and no epoch RVS spectra: 9 of the 1,132 candidate sources in the EM-2 snapshot have a mean RVS, 0 have published epoch RVs, and 0 candidate pairs have both. The full epoch-series test is not executable on DR3; only a weak epoch-constancy proxy is available, validated at 1/300 trials with a wide confidence interval.

The GW channel uses the mass-shifted pair. A holonomy pair has $\Lambda_1 = \ln(m_1^j/m_1^i)$ and $\Lambda_2 = \ln(m_2^j/m_2^i)$ equal to the same Λ , because both masses are scaled by the same factor; the mass ratio and χ_{eff} are unchanged. The cut chain is pair scaling, χ_{eff} (where available), the invariant M_c/D_L (equivalently $\Lambda_D \equiv \ln(D_L^j/D_L^i) = \Lambda$, not zero), sky co-localization, posterior-level Λ , and the shared-base family test [39]. On the GWTC-5.0 public catalog (391 events, 282 with source-frame masses, 39,621 pairs), with the lensing-search baseline and the catalog methods of Refs. [31, 32], the Stage-1 count at the one-degree-of-freedom 3σ cut is 32,434 against a label-

permutation null median of 32,806 (two-sided $p \approx 0.10$, a mild deficit, not an excess). The corrected M_c/D_L scaling cut at 3σ gives 30,343 candidates against a matched independent-pair null median of 30,349 (two-sided $p \approx 0.96$), a null that recovers 5,688 pairs the archived equality- D_L cut discarded; the 3–6 σ corrected sweep is likewise null. The corrected sky co-localization pilot gives $p = 0.75$; the full volume gives $p = 0.95$ over 30,311 mapped candidates, with largest candidate co-localization $C = 0.9513$. On the corrected co-located set (3,669 candidates at $C > 0.1$), the family statistic is $S = 2,697$; the nominal sky-permutation value $p = 0.005$ falls to $p = 0.539$ under the composition-matched null, so the excess is a composition shift, not a family. At posterior precision the corrected M_c/D_L residual leaves 82 of 120 pairs (63 both, 19 corrected-only, 2 equality-only; equality 65, labelled only). The family test is degenerate at $\sigma_\Lambda = 0.08$ –0.5, so the shared-base cluster over the corrected 120 saturates (64 of 82 at $|\Lambda_0| = 0.467$, $p = 1.000$; equality 57 of 65). The winding-ladder diagnostic is exploratory and not part of the archived chain.

6.4 What the nulls establish

The GW null establishes that a relic population which would write a large number of detectable, posterior-confirmed, co-located, shared-base pairs would have produced an excess over the sky-permutation null. The archived cut required equality of D_L , so its null is a null on that cut, not on the scaling signature. The scaling-corrected cut has now been run: it is null at Stage 1.5 (30,343 against a matched null of 30,349, two-sided $p \approx 0.96$), the corrected sky stages are null ($p = 0.75$ pilot, $p = 0.95$ full volume), and the corrected family excess does not survive the composition-matched null ($p = 0.539$). Under the calibrated counting assumptions on that corrected co-located set, the 90% upper limit is $\mu < 274$ at the operational $C > 0.1$ threshold (205 under the Gamma–Poisson model, which is not adequate at $C > 0.1$), or $\mu < 4.50$ (4.95 conservatively) for the $C > 0.9$ strong-co-localization subset. The often-quoted $\mu < 2.3$ is only the ideal zero-observed case, not a data-driven bound. An injection/selection model gives $\varepsilon_{\text{family}} \simeq 1.8 \times 10^{-36}$, so the count limit implies $N_{\text{relic}} < 1.5 \times 10^{38}$ (1.1×10^{38} model-dependent) and no astrophysically meaningful density bound. The canonical metre-scale pocket is invisible to the GW pair channel because its image delay (~ 3 ns) is far below the event-separation threshold; the channel opens only for $\tau_{\text{wall}} \gtrsim 0.002$ –0.02 AU. The EM census likewise gives no detection and no exclusion. What the nulls establish is a methods result and a set of weak conditional limits, not the absence of relics.

This is the correct place to repeat the paper’s honesty constraint: a null in a survey with an unmeasured per-relic efficiency cannot exclude a population whose expected detectable yield lies below the bound. The current surveys are sensitive to large pockets and trigger-epoch cascades, not to the canonical metre-scale relic. The survey is a live empirical task precisely because the theory does not predict the relic abundance, and the instruments are only now reaching the relevant volumes and precisions.

6.5 Capture and wrap

The capture step converts a relic into an armed machine. The relic supplies the identification; the program supplies the envelope. The C4 construction matches the core’s umbilical boundary data to a dust envelope, flattens onto Ori’s Layer-2 form, and hands off to his outer layers [37, 2]. The dust is a pressureless fluid of ordinary composition, not exotic matter; the energy conditions hold on the gatefast interior band $r \geq w + 0.0904$ and, in the retuned construction, over the full C4 envelope layer $r > w$ including the innermost collar at $a = 0.02$; the initial-value problem is well posed; the evolution into CTCs is conditional on Ori’s existence theorem. The canonical arming time is nanoseconds; the full canonical envelope is 1.5–2.8 M_J (the 0.34 M_J vacuum tube is a component,

Observation or result	What it would change	Current status
3 + 1 CFG counterexample for a non-Gott data class	Reopens the energy proof gap; not a build route by itself	No such result; 2 + 1 theorem only
Topology change observed at accessible scales	Breaks Geroch’s dichotomy as a practical law	No such observation
Cosmic strings at $G\mu/c^2 \gtrsim 10^{-2}$	Reopens the Gott pair route	Excluded by 8–12 orders
Macroscopic NEC violation	Reopens exotic-matter routes	Bounded by quantum inequalities; no lab source
Confirmed boost-holonomy relic	Confirms Route A; does not test the no-go	No detection
Quantum-gravity no-go for foam pockets	Kills Route A’s first step	No such theorem
Failure of the boost quotient to survive inflation with its holonomy intact	Kills Route A’s second step	Open; spatial-topology preservation alone does not settle it

Table 2: Falsifiers and what they would change. The table separates proof-gap falsifiers from build-route falsifiers.

not the machine mass); the deep rung is ordinary-density but asteroid-scale. The capture mission itself—transport, placement, envelope construction around a natural object—is not designed here. We state it as the remaining engineering, not as a solved problem.

7 Falsifiers and observational strategy

7.1 Falsifiers of the thesis

The thesis is falsifiable in specific, named ways. Table 2 lists them with their current status. The most important distinction is between falsifiers that would reopen a *proof gap* and falsifiers that would open a *build route*. A 3 + 1 CFG counterexample is the former; a confirmed relic is neither (it confirms Route A, not the no-go); a topology-change event at accessible scales would be the latter, because it would break Geroch’s dichotomy as a practical law.

7.2 Falsifiers of the routes

Route A fails if the foam does not produce stable quotient pockets, if the boost-identified pocket does not survive inflation with its holonomy intact, or if survival fails on cosmological timescales. The first is a quantum-gravity question; the second is an open question about the quotient under inflation, not a consequence of standard spatial-topology preservation; the third is the companion’s stability program. A theorem that the full quantum theory forbids local quotient pockets would be decisive against Route A. A demonstration that the known instabilities grow on cosmological timescales would be decisive against survival. Neither exists.

Route B fails if the full theory forbids Euclidean nucleation of an identified pocket, or if the amplitude is exactly zero. No such result is known. The practical falsifier is not a no-go but a null: if no relic is found and the foam abundance is independently bounded, the route is constrained.

Route C fails if the CTC region cannot support the operation, or if the transported pocket is unstable. The companion’s backreaction and stability programs are the relevant calculations. As a first-machine route it is already excluded by circularity; as a bootstrap it is limited by the host.

7.3 Observational strategy

The strategy follows from the signature and the cloistering. There are four lines.

Multi-epoch spectroscopy of co-spatial pairs. The decisive EM test is not a single-epoch radial-velocity offset; it is a constant, exactly multiplicative shift across all spectral features and all epochs. A binary’s offset varies with orbital phase; a holonomy shift does not. Gaia DR3 lacks epoch RVS spectra, so the full test is data-limited; the next generation of multi-epoch surveys (Gaia DR4 and beyond, 4MOST, WEAVE, SDSS-V, ELT/HIRES) is the natural home. The current 939-candidate sample is a background sample and a training set for the selection function, not a detection sample.

Gravitational-wave pairs at higher precision. The GW channel needs better mass precision, because the family test is degenerate when σ_Λ is comparable to the Λ structure. Next-generation detectors (Einstein Telescope, Cosmic Explorer, LISA) will improve the precision and the volume. The co-localization cut is already the discriminator; the family test becomes decisive only when σ_Λ is small compared with the spacing of the Λ ladder.

Large pockets and the trigger epoch. The cloistering calculation says that small pockets are invisible to generic stellar surveys. The two windows are large pockets ($\tau_{\text{wall}} \gtrsim 0.1$ ly) and the trigger epoch (the image cascade near the chronology horizon). The first is a lensing-like search for massless double images with a fixed ratio; the second is a search for anomalous image multiplicities and frequency ladders in transient or variable sources. Both are worth doing, and neither is currently systematic.

The massless-lens discriminator. Any candidate must pass the test that no mass lens can fake: the fixed frequency ratio and its integer-power ladder, with no deflection and no lens convergence. A boost is co-spatial; its image differs in frequency and hence in flux by a spectral remapping, not by an achromatic magnification. A pair with a fitted lens mass is not a relic; a pair with a fixed ratio across independent sources is. The family test is the statistical version of this discriminator, and its degeneracy is the main methodological obstacle. Analogue experiments, such as the optical chronology-protection proposal of Bermudez and Leonhardt [25], test the divergence mechanism rather than the existence of a quotient; they are a complementary laboratory check, not a relic search.

7.4 A decision tree

The program’s logic can be summarized as a decision tree. If a relic is found, Route A is confirmed and the capture/wrap pipeline is the engineering path. If no relic is found and the foam abundance is bounded below the survey’s reach, Route A is constrained but not killed; the bound is only as good as the selection model. If a $3 + 1$ CFG counterexample is found, the construction route reopens as a proof problem, but Geroch and DJHT/Cutler still stand. If a topology-change event is observed, the classical no-go landscape changes qualitatively. If a macroscopic NEC violation is engineered, the exotic-matter routes reopen, but they are a different program. The survey is the only branch that is currently executable, and it is the branch the program is pursuing.

8 Conclusion

The time machine’s matter side is buildable in the only sense that matters for a conceptual program: the initial data exist, explicitly, with a pressureless dust of ordinary composition rather than exotic matter, energy conditions on the gatefast interior band $r \geq w + 0.0904$ and, in the retuned construction, over the full C4 envelope layer $r > w$ including the innermost collar at $a = 0.02$, a well-posed initial-value problem, and a verified envelope. The full canonical envelope is 1.5–2.8 M_J ; its density reaches about 5×10^8 times nuclear saturation, so “ordinary” describes the composition, not the density. Only the extrapolated deep rung reaches ordinary density, and there the mass is asteroid-scale. The machine’s organ is not matter. It is the boost-holonomy identification, a global quotient whose holonomy is a Lorentz boost. That organ is not classically manufacturable on the known classical routes. The sharpest classical statement is the Carroll–Farhi–Guth–Olum energy obstruction in $2 + 1$ dimensions; Geroch’s topology-change dichotomy, the Deser–Jackiw–’t Hooft/Cutler singular-constraint analysis of the two-defect route, and topological censorship reinforce the conclusion with different hypotheses. The $3 + 1$ CFG statement for our data class is a gap; the thesis is an argument with a $3 + 1$ proof gap, not a theorem.

The conclusion is not that time machines are impossible. It is that the question has changed. The known classical routes considered here are closed at the organ, not at the matter. What remains is acquisition: primordial inheritance through foam, inflation, survival and survey; quantum nucleation; or reproduction inside an existing chronology-violating region. Route A is the only classical-compatible, non-circular route, and its first step is the survey. The survey is the live empirical task: a bare relic is a massless, frequency-multiplicative lens, producing co-spatial spectral double images with a fixed ratio and no lens mass. The two executed channels return nulls and weak conditional limits—the GW equality- D_L and corrected mass–distance scaling cuts are both null, and the family excess does not survive a composition-matched null—and they do not exclude the canonical relic or establish a density bound.

The program’s remaining work is therefore divided cleanly. On the theory side, the $3 + 1$ CFG gap and the foam/survival premises are the honest open problems. On the observational side, the survey needs multi-epoch spectroscopy, better mass precision, and a systematic search for large pockets and trigger-epoch cascades. On the engineering side, the capture mission is the least developed part of the wrap. None of these is a reason to abandon the question. They are the question, stated in the form in which it can be answered.

A The boost quotient and the photon ratio

This appendix collects the algebra used in the text. The core metric is Eq. (2). In the $f = 0$ limit the (T, z) block is

$$ds_{(2)}^2 = -2 dz dT - T dz^2, \quad (7)$$

which is Misner space. The covering coordinates

$$\tau = 2\sqrt{-T}, \quad \psi = \frac{z}{2} + \frac{1}{2} \ln(-2T) \quad (8)$$

bring the block to

$$ds_{(2)}^2 = -d\tau^2 + \tau^2 d\psi^2, \quad (9)$$

the Milne wedge. The identification $z \sim z + L$ acts on ψ as $\psi \sim \psi + L/2$, which is the Lorentz boost of rapidity $\Lambda = L/2$ in the (τ, ψ) plane. The deck transformation is therefore a boost, and the quotient is a Misner-type spacetime.

A photon crossing the seam once has its four-momentum acted on by the deck boost. Write the photon momentum in the covering plane as $k^\mu = (\omega, \omega n_z)$ in the axial case, where $n_z = \pm 1$. A boost of rapidity Λ multiplies the frequency by $e^{\pm\Lambda}$, depending on the sign of n_z . For a general direction cosine, the frequency ratio measured by a comoving observer is Eq. (5). The axial cases are exact because the boost does not change the spatial direction of a photon moving along the boost axis.

For a gravitational wave the same argument applies to the frequency of the whole waveform. The detector-frame chirp mass is inferred from the frequency and its derivative, so a common multiplicative factor r on all frequencies ($f \rightarrow rf$, $\dot{f} \rightarrow r^2 \dot{f}$) becomes a factor r^{-1} on the inferred masses. The mass ratio and χ_{eff} are invariant, and the transverse-traceless amplitude h is invariant, so there is no magnification. The amplitude relation $h \propto M_c^{5/3} f^{2/3} / D_L$ then forces $D_L \rightarrow D_L / r$: the inferred distance scales with the inferred mass, and the invariant combination is M_c / D_L , not D_L by itself. This is the algebra of the corrected GW cut; the archived cut of Ref. [39] used equality of D_L and is a null on that cut.

The identification is not a field on the slice. It is a deck transformation of the covering space, and it is invisible to any local observer. This is the precise sense in which the organ is global. The metric Eq. (2) is a local solution of the vacuum equations; the quotient is an additional global datum. The matter envelope can be specified locally; the quotient cannot.

B Theorem hypotheses and scope

Table 3 lists the classical results used in the paper with their hypotheses and their scope. The table is deliberately conservative: it states what each result does not cover, because the paper's argument depends on the combination, not on any single result being stronger than it is.

Result	Hypotheses	What it does and does not cover
CFG/CFGO [3, 4]	2 + 1 dimensions; edgeless noncompact spacelike slice; timelike total momentum; open universe; Gott configuration	Proves the energy obstruction. Does not cover 3 + 1, the Ori Misner-type core, or non-Gott holonomies.
Geroch [5]	Classical, globally hyperbolic spacetime with compact Cauchy surfaces; topology change between Cauchy surfaces	Forbids topology change without CTCs, singularity or energy-condition violation. Does not forbid a spacetime with a chronology horizon; the machine pays by having one.
DJHT/Cutler [6, 7]	Two-defect Gott spacetime; regular initial data; acceptable Cauchy surface	Shows the two-string CTC spacetime is constraint-singular or requires closed spatial topology. Specific to the two-defect route.
Topological censorship [8]	Null energy condition; globally hyperbolic spacetime (a Cauchy surface); asymptotic flatness; simply connected asymptotic region	Applies to the pre-machine phase, not the machine, which is not globally hyperbolic. Hides nontrivial topology from causal probes. Does not forbid the identification; makes it causally cloistered until the horizon.

Table 3: The classical results and their scope. The paper's conclusion is the combination, not any single theorem.

C Claims ledger

Table 4 classifies the paper’s claims by status. The classification follows the program’s practice: a claim is a theorem, a cited result, a computed result, an estimate, a scenario, or an open problem. The table is the paper’s answer to the question “which sentence is load-bearing, and on what?”

Claim	Status	Basis
Composite initial data exist, explicit, vacuum + dust (not exotic)	Constructed	Companion P2; Ori 2007
Canonical envelope $1.5\text{--}2.8 M_J$; density reaches $\sim 5 \times 10^8$ nuclear	Computed	Companion P2; Sec. 2.3
Energy conditions on the gatefast band $r \geq w + 0.0904$ and, retuned, the full C4 layer $r > w$	Computed	Companion P2; collar retuned at $a = 0.02$
Horizon stress is finite at design points, eikonal/test-field	Computed (local branch)	Companion P1; global CH not proved
Identification is a global boost quotient, $\Lambda = L/2$	Derived	Appendix A; companion P3
Identification is not a local field	Structural	Definition of quotient; App. A
CFG 2 + 1 energy obstruction	Theorem	Refs. [3, 4]
CFG 3 + 1 for our data class	Open (gap)	Sec. 5
Geroch topology-change dichotomy	Theorem	Ref. [5]
DJHT/Cutler two-defect obstruction	Theorem	Refs. [6, 7]
Pre-horizon causal cloistering	Derived	Sec. 3.4; Ref. [8]
Route A foam abundance	Scenario	Ref. [13]; not a theorem
Route A inflation/quotient survival	Open	Sec. 4.1; not implied by spatial-topology preservation
Route A survival	Conditional	Companion P1/P3
Route B nucleation allowed	Open	No known prohibition
Route C reproduction	Conditional	Requires a host machine
Photon ratio $r = \cosh \Lambda - n_z \sinh \Lambda$	Derived	Appendix A
GW masses and D_L scale together; M_c/D_L invariant	Derived	Appendix A; companion P4
Survey nulls	Computed	Companions P4/P5
Survey limits	Conditional	Counting + selection model
Relic density bound	Not claimed	Efficiency unmeasured

Table 4: Claims ledger. “Constructed”, “computed” and “derived” refer to results in this paper or its companions; “theorem” refers to a published result with the hypotheses of Appendix B; “scenario” and “conditional” are not theorems.

References

- [1] A. Ori, “A class of time-machine solutions with a compact vacuum core,” *Phys. Rev. Lett.* **95**, 021101 (2005), arXiv:gr-qc/0503077.
- [2] A. Ori, “Formation of closed timelike curves in a composite vacuum/dust asymptotically-flat spacetime,” *Phys. Rev. D* **76**, 044002 (2007), arXiv:gr-qc/0701024.
- [3] S. M. Carroll, E. Farhi, A. H. Guth, “An obstacle to building a time machine,” *Phys. Rev. Lett.* **68**, 263 (1992); erratum **68**, 3368 (1992), arXiv:hep-th/9207037.
- [4] S. M. Carroll, E. Farhi, A. H. Guth, K. D. Olum, “Energy-momentum restrictions on the creation of Gott time machines,” *Phys. Rev. D* **50**, 6190 (1994), arXiv:gr-qc/9404065.
- [5] R. P. Geroch, “Topology in general relativity,” *J. Math. Phys.* **8**, 782 (1967).
- [6] S. Deser, R. Jackiw, G. ’t Hooft, “Physical cosmic strings do not generate closed timelike curves,” *Phys. Rev. Lett.* **68**, 267 (1992).

- [7] C. Cutler, “Global structure of Gott’s two-string spacetime,” *Phys. Rev. D* **45**, 487 (1992).
- [8] J. L. Friedman, K. Schleich, D. M. Witt, “Topological censorship,” *Phys. Rev. Lett.* **71**, 1486 (1993); erratum **75**, 1872 (1995).
- [9] J. R. Gott, “Closed timelike curves produced by pairs of moving cosmic strings: exact solutions,” *Phys. Rev. Lett.* **66**, 1126 (1991).
- [10] S. W. Hawking, “Chronology protection conjecture,” *Phys. Rev. D* **46**, 603 (1992).
- [11] B. S. Kay, M. Radzikowski, R. M. Wald, “Quantum field theory on spacetimes with a compactly generated Cauchy horizon,” *Commun. Math. Phys.* **183**, 533 (1997), arXiv:gr-qc/9603012.
- [12] C. W. Misner, “Taub-NUT space as a counterexample to almost anything,” in *Relativity Theory and Astrophysics I: Relativity and Cosmology*, Lectures in Applied Mathematics Vol. 8, edited by J. Ehlers (American Mathematical Society, Providence, 1967), pp. 160–169.
- [13] S. Carlip, “Spacetime foam: a review,” *Rep. Prog. Phys.* **86**, 066001 (2023), arXiv:2209.14282.
- [14] M. Lachièze-Rey, J.-P. Luminet, “Cosmic topology,” *Phys. Rep.* **254**, 135 (1995).
- [15] E. Farhi, A. H. Guth, “An obstacle to creating a universe in the laboratory,” *Phys. Lett. B* **183**, 149 (1987).
- [16] E. Farhi, A. H. Guth, J. Guven, “Is it possible to create a universe in the laboratory by quantum tunneling?” *Nucl. Phys. B* **339**, 417 (1990).
- [17] D. Deutsch, “Quantum mechanics near closed timelike lines,” *Phys. Rev. D* **44**, 3197 (1991).
- [18] B. Shoshany, J. Wogan, “Wormhole time machines and multiple histories,” *Gen. Rel. Gravit.* **55**, 44 (2023), arXiv:2110.02448.
- [19] L. H. Ford, T. A. Roman, “Averaged energy conditions and quantum inequalities,” *Phys. Rev. D* **51**, 4277 (1995).
- [20] J. R. Fliss, A. Rolph, “Curious QNEIs from QNEC: new bounds on null energy in quantum field theory,” *JHEP* (2026), to appear, arXiv:2510.26247.
- [21] S. Xavier, “Closed timelike curves from a vacuum traveling wave,” arXiv:2607.00788 (2026).
- [22] N. E. Rieger, “Causal classification of pathological Misner-type spacetimes,” *Annals of Physics* **489**, 170455 (2026), arXiv:2501.00841, doi:10.1016/j.aop.2026.170455.
- [23] W. Melton, A. Strominger, T. Wang, “Quantum fields on a time-periodic $\text{AdS}_3/\mathbb{Z}$ background,” arXiv:2510.15036 (2025).
- [24] H. Andr  ka, J. Madar  sz, J. Manchak, I. N  meti, G. Sz  kely, “An extendible spacetime without closed timelike curves whose every extension contains closed timelike curves,” *Found. Phys.* (2026), to appear, arXiv:2510.26902.
- [25] D. Bermudez, U. Leonhardt, “Optical time travel: proposal for testing Hawking’s chronology protection conjecture in an optical analogue,” *Class. Quantum Grav.* (2026), to appear, arXiv:2607.22056.
- [26] B.-H. Lee, N. A. Nilsson, S. Thakur, “Global causality constraints in rotating scalar-tensor spacetimes,” *Phys. Rev. D* (2026), to appear, arXiv:2602.17724.
- [27] F. Ahmed, A. Al-Badawi,  . Sakalli, “Are Petrov type-N and D spacetimes admitting CTCs valid in $f(R, \mathcal{L}_m, \Phi, X)$ gravity?” *Eur. Phys. J. C* (2026), to appear, arXiv:2605.26696.

- [28] D. Senjaya, “Probing the chronology protection conjecture via scalar quaresonance,” *Eur. Phys. J. C* **86**, 34 (2026).
- [29] Planck Collaboration, “Planck 2013 results. XXV. Searches for cosmic strings and other topological defects,” *Astron. Astrophys.* **571**, A25 (2014), arXiv:1303.5085.
- [30] NANOGrav Collaboration, “The NANOGrav 15 yr data set: evidence for a gravitational-wave background,” *Astrophys. J. Lett.* **951**, L8 (2023).
- [31] LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration, “GWTC-4.0: searches for gravitational-wave lensing signatures,” arXiv:2512.16347 (2025; v3 2026).
- [32] LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration, “GWTC-5.0: methods for identifying and characterizing gravitational-wave transients,” arXiv:2605.27224 (2026).
- [33] E. Minguzzi, “Area theorem and smoothness of compact Cauchy horizons,” *Commun. Math. Phys.* **339**, 57 (2015), arXiv:1406.5919.
- [34] K. D. Olum, “The Ori–Soen time machine,” *Phys. Rev. D* **61**, 124022 (2000), arXiv:gr-qc/9907007.
- [35] J. R. Gott, L.-X. Li, “Can the universe create itself?” *Phys. Rev. D* **58**, 023501 (1998), arXiv:astro-ph/9712344.
- [36] J. Schemoul, “Protection without teeth: semiclassical chronology protection on the Ori time machine,” in preparation (2026).
- [37] J. Schemoul, “Initial data for a time machine from dust: an explicit composite construction with a certified ordinary-density scaling regime,” in preparation (2026).
- [38] J. Schemoul, “Quotient structure and on-generator diagnostics of vacuum traveling-wave time machines,” in preparation (2026).
- [39] J. Schemoul, “A search for boost-holonomy relics in the gravitational-wave catalog,” in preparation (2026).
- [40] J. Schemoul, “A Gaia search for boost-holonomy relics: the full-sky co-spatial radial-velocity pair census,” in preparation (2026).