

A Particle Is a Treaty Between a Field and a Worldline

Registered research scaffold — no numerical result

Brecht Corbeel

Absolute Digital Publishers

absolutedigitalpublishers.com

September 20, 2026

Abstract

Whether a quantum field contains a particle depends on which worldline is used to decompose the field into modes, a dependence fixed mathematically by the Bogoliubov transformation and sharpened physically by the Unruh effect. This draft registers an operational way to price that dependence without claiming to have measured it. Working with the Unruh–DeWitt detector’s response functional rather than with mode counting directly, we normalize the excitation-branch response into a probability density over detector gap and define the disagreement between two prepared detectors as the total-variation distance between their normalized spectra. We state, and verify by direct construction rather than assumption, that this distance is a genuine metric on normalized spectra, bounded in $[0, 1]$, invariant under worldline reparametrization and joint Poincaré transformation, and exactly zero for identical trajectories. We register four matching conditions — field state, coupling type, detector gap referred to proper time, and switching profile — that must be held fixed before a nonzero reading is attributed to the trajectory rather than to mismatched apparatus. We recover, as inherited baseline rather than as a result of this paper, the exact Planckian excitation spectrum of an eternally and uniformly accelerated detector at the Unruh temperature, derive in closed form the leading-order sensitivity of that spectrum to a small change in acceleration from a Wien-type transcendental equation, and identify a boundary case, the eternally switched-on inertial trajectory, in which the construction’s own normalization is undefined rather than merely small. The one laboratory-adjacent regime identified is electron storage-ring polarization, noted with its caveats rather than treated as a measurement of the quantity proposed here. No detector-description disagreement atlas over acceleration, duration, gap, field state, and switching profile is reported: the parameter grid, output schema, and kill criteria for that atlas are

registered, and a companion program checks the metric’s exact analytic properties and its exact known limits, but the atlas itself, a finite-switching numerical solver, curved-spacetime state implementations, and external subject-matter review remain future work stated as such. **Keywords:** Unruh–DeWitt detector;

Bogoliubov transformation; Unruh effect; detector response functional; total-variation distance; Rindler horizon; switching function; storage-ring polarization.

Series: Revolutionary Life 75, Series D, D05.

Status: Draft scaffold. The disagreement atlas, finite-switching solver, curved-spacetime state implementation, and external review are incomplete.

1 Two bookkeepings for one click

A detector click is a macroscopic, recorded fact about a piece of apparatus. The sentence built on top of it, that a particle was present, is a separate act: it fixes a decomposition of the quantum field into positive- and negative-frequency modes, and two observers using different decompositions can each build a self-consistent bookkeeping in which the other observer’s particle is absent. This is inherited, established physics, not a claim of this paper: it is the content of the Bogoliubov transformation between mode decompositions, and it is demonstrated concretely by the Unruh effect, in which a uniformly accelerated detector in the ordinary field vacuum reports a thermal bath that an inertial detector in the same field state does not [1, 2].

This paper does not propose to remove that observer-dependence, and no section below attempts to. It proposes to price it: to replace the yes/no question of whether a particle is present with a bounded, computable disagreement between what two specific, fully specified detector protocols report, once every differ-

ence between the two protocols except the worldline itself has been pinned down identically. The central object, defined in Section 3, is a distance $D_{\text{det}}(x, y)$ between the normalized excitation spectra reported by an Unruh–DeWitt detector on worldline x and one on worldline y .

The scientific content of this proposal is conditional on two things holding after implementation, neither of which is asserted here: that D_{det} does not move under a reparametrization of either worldline’s curve parameter, and that a nonzero reading, once the four matching conditions of Section 5 are enforced, is attributable to the trajectories rather than to a difference in detector design. Section 4 establishes the first property by construction. The second is checked directly, not assumed, in Sections 6–9: one worked case shows a mismatched coupling type moving D_{det} on a single fixed worldline, which is the failure mode the construction exists to catch rather than a defect in it.

What this paper delivers is a definition, a proof that the definition behaves like a distance, two exactly solvable limiting cases, and one closed-form calculation carried through by hand. It does not deliver the disagreement atlas that would make the object usable across a full parameter range; Section 10 states plainly what stands between this draft and that atlas.

2 What Bogoliubov mixing and the Unruh response already settle

Given two decompositions of a free field into modes carrying annihilation operators a_k and $\tilde{a}_j$, unitarity of the field theory forces the two sets of ladder operators to mix,

$$\tilde{a}_j = \sum_k \left(\alpha_{jk} a_k + \beta_{jk} a_k^\dagger \right), \quad (1)$$

subject to the normalization conditions that preserve the bosonic commutation relations of both decompositions,

$$\sum_k (\alpha_{jk} \alpha_{lk}^* - \beta_{jk} \beta_{lk}^*) = \delta_{jl}, \quad \sum_k (\alpha_{jk} \beta_{lk} - \beta_{jk} \alpha_{lk}) = 0. \quad (2)$$

Whenever any $\beta_{jk} \neq 0$, the vacuum of the a -decomposition contains excitations of the $\tilde{a}$ -decomposition, and vice versa. Fulling showed this mixing already occurs in flat, static, two-dimensional Minkowski space from a nonstandard but legitimate time coordinate [2]; Davies extended the same mixing into the Rindler wedge of ordinary flat spacetime, associating a temperature with the causal horizon an accelerated observer trails behind them [3], and Hawking’s contemporaneous horizon calculation for a collapsing star established the same mode-mixing

mechanism as the origin of black-hole radiance [4]. None of Eqs. (1)–(2) is proposed here; they are the mathematical content the rest of this paper takes as given.

Unruh sharpened the argument by modeling the measuring device itself rather than the modes: a pointlike two-level system linearly coupled to the field along its own worldline responds to whatever the field does along that specific worldline. Along an eternal, uniformly accelerated trajectory in the ordinary vacuum, it responds exactly as a thermal detector at temperature

$$k_B T_U = \frac{\hbar a}{2\pi c}, \quad (3)$$

where a is the trajectory’s constant proper acceleration [1]. Every quantity in Eq. (3) is fixed by the trajectory and by fundamental constants alone: $\hbar a$ has units of energy times inverse time, dividing by c (length over time) leaves units of energy, matching $k_B T_U$. Takagi confirmed the same Planckian structure holds for a Rindler wedge of a field theory in an arbitrary number of space-time dimensions, ruling out the possibility that the thermal form is a low-dimensional artifact [8], and Crispino, Higuchi, and Matsas give the standard modern review of the effect and its detector-based derivations [6]. This paper treats Eqs. (1)–(3) as observed-baseline, derived results: nothing about them is new here.

What the detector framing adds, and what this paper’s construction is built on top of, is Unruh and Wald’s demonstration that a click’s occurrence and a click’s description are separable questions. When an accelerated detector reports absorbing a Rindler-frame quantum, an inertial observer watching the same event describes the detector as spontaneously emitting a Minkowski-frame photon, with energy conservation restored once the work performed by the accelerating agent is included in the ledger [5]. The click is not in dispute between the two descriptions; only the vocabulary is. That is the precedent this paper follows throughout: compare detector responses at the level of the click, and let interpretation follow only after the comparison is made precise.

3 The response functional, the normalized spectrum, and the distance built on them

Let $x(\tau)$ be a timelike worldline parametrized by its own proper time τ — never by the coordinate time t of an ambient chart, which need not agree with τ except along a comoving inertial trajectory. A pointlike Unruh–DeWitt detector on $x(\tau)$, with energy gap Ω conjugate to τ and switching function $\chi(\tau)$ controlling when the monopole coupling is active, has response

functional

$$\mathcal{F}_x(\Omega) = \int d\tau d\tau' \chi(\tau)\chi(\tau') e^{-i\Omega(\tau-\tau')} W[x(\tau), x(\tau')], \quad (4)$$

where W is the field's two-point Wightman function restricted to the worldline [6]. Three quantities in Eq. (4) are kept deliberately distinct throughout this paper: τ , the detector's own proper time; Ω , the energy gap measured in the detector's instantaneous rest frame, not a coordinate-frame frequency; and the field state entering through W , which can be changed independently of the trajectory. Collapsing any two of these — for instance building Ω from coordinate time when two worldlines sit at different heights in a static spacetime — injects an ordinary redshift into a quantity meant to isolate worldline-dependence of particle content, which is not the effect this paper is about.

Definition 3.1 (Normalized excitation spectrum). *For $\Omega > 0$ and a worldline x for which the integral converges to a finite, nonzero value,*

$$p_x(\Omega) = \frac{\mathcal{F}_x(\Omega)}{\int_0^\infty \mathcal{F}_x(\Omega') d\Omega'}. \quad (5)$$

The restriction to $\Omega > 0$ keeps only the excitation branch: whether the detector behaved as though it absorbed a field quantum. The decay branch $\Omega < 0$ mixes genuine field structure with spontaneous-emission physics present even for a detector in flat, empty spacetime with the field in its vacuum, and is excluded from p_x for that reason. Since $p_x(\Omega) d\Omega$ is by construction a probability, p_x carries units of time (inverse Ω); whatever dimension the coupling constant and field normalization give $\mathcal{F}_x$ itself cancels between numerator and denominator of Eq. (5) and plays no further role.

Definition 3.2 (Detector-description distance). *For two worldlines x, y whose normalized spectra p_x, p_y both exist,*

$$D_{\text{det}}(x, y) = \frac{1}{2} \int_0^\infty d\Omega |p_x(\Omega) - p_y(\Omega)|. \quad (6)$$

Because $p_x(\Omega)$ and $p_y(\Omega)$ each carry units of time and $d\Omega$ carries units of inverse time, the integrand of Eq. (6) is dimensionless at every point of the domain, and D_{det} itself is a dimensionless number: it is a comparison between two probability densities, not a physical rate or energy. D_{det} is not defined for a pair of worldlines unless both normalized spectra exist; Section 6 exhibits a trajectory, the eternal inertial worldline, for which p_x itself is undefined, and the consequence for D_{det} in that case is stated there rather than patched over here.

4 The distance is a genuine metric, and two symmetries it must respect

Proposition 4.1 (Metric properties of D_{det}). *On the set of worldlines whose normalized spectra exist, D_{det} is non-negative, symmetric, satisfies the triangle inequality, vanishes if and only if $p_x = p_y$ almost everywhere, and takes values in $[0, 1]$.*

Proof. Non-negativity and symmetry are immediate from Eq. (6), since $|p_x - p_y| = |p_y - p_x| \geq 0$ pointwise. If $p_x = p_y$ almost everywhere the integrand vanishes almost everywhere and $D_{\text{det}} = 0$; conversely, if $D_{\text{det}} = 0$ then a non-negative integrand with zero integral vanishes almost everywhere, so $p_x = p_y$ almost everywhere. For any three worldlines x, y, z , the pointwise triangle inequality for real numbers, $|p_x(\Omega) - p_z(\Omega)| \leq |p_x(\Omega) - p_y(\Omega)| + |p_y(\Omega) - p_z(\Omega)|$, integrates to $D_{\text{det}}(x, z) \leq D_{\text{det}}(x, y) + D_{\text{det}}(y, z)$. For boundedness, let $A = \{\Omega > 0 : p_x(\Omega) \geq p_y(\Omega)\}$; since $\int_0^\infty (p_x - p_y) d\Omega = 1 - 1 = 0$, the integral over A of $(p_x - p_y)$ equals minus the integral over its complement, so $\int_0^\infty |p_x - p_y| d\Omega = 2 \int_A (p_x - p_y) d\Omega \leq 2 \int_A p_x d\Omega \leq 2$, using $p_y \geq 0$ on A and $\int_A p_x d\Omega \leq \int_0^\infty p_x d\Omega = 1$. Hence $D_{\text{det}}(x, y) \leq 1$. $\square$

Proposition 4.1 is not a property invented for this construction; it is inherited for free because D_{det} is built from the total-variation distance, an object elementary measure theory already understands, applied to two properly normalized probability densities. An ad hoc "disagreement score" assembled without this structure would need each property checked by hand; D_{det} does not.

Remark 4.1 (Reparametrization invariance). D_{det} does not depend on the auxiliary curve parameter used to describe a worldline before converting to proper time. Proper time is not a gauge choice available to be varied; once the trajectory and the metric are fixed, τ is the arc length of the worldline, and every quantity in Eq. (4) — τ , Ω , χ , and W evaluated at $x(\tau)$ — is defined intrinsically along the curve. There is no residual freedom in Eq. (6) for a coordinate reparametrization to act on.

Remark 4.2 (Joint Poincaré invariance). If a single Poincaré transformation is applied identically to both worldlines and to the field vacuum, D_{det} is unchanged, because the Wightman function of the vacuum is Poincaré invariant and τ , Ω , and χ are each defined intrinsically along each worldline rather than against any external frame. This is a statement about a transformation applied to both members of a comparison pair together; it says nothing about comparing a worldline to itself as seen in two different frames, which is the

trivial case $D_{\text{det}}(x, x) = 0$ already covered by Proposition 4.1.

These two invariances are the half of the kill criterion in Section 9 that concerns mere relabeling: nothing about how a trajectory is coordinatized or boosted can manufacture a nonzero reading. The other half of the kill criterion, that a nonzero reading not be dominated by unmatched detector design, is a substantive requirement rather than a structural one, and it is the subject of Section 5.

5 Four clauses a nonzero reading must satisfy

$D_{\text{det}}(x, y)$ compares two full experimental protocols, not two bare trajectories. Before a nonzero value is read as evidence that the worldlines disagree about particle content, four other quantities have to be matched between the two detectors, because each one moves D_{det} on its own, independent of x and y as curves through spacetime.

Field state. The Wightman function W in Eq. (4) is not fixed by the spacetime geometry alone once curvature is present. A Schwarzschild exterior, for instance, supports the Boulware vacuum, the Hartle–Hawking state, and the Unruh state, each giving a physically distinct W on the same background [14]. Comparing a detector in one field state against a detector in another is a statement about field states, not about trajectories.

Coupling type. A detector linearly coupled to the field and a detector coupled to the field’s proper-time derivative are different instruments even on an identical worldline: their response functionals carry different powers of Ω in the numerator, so their normalized spectra differ pointwise for purely instrumental reasons. Moustos showed this early-time dependence on coupling type explicitly, while also showing that the late-time asymptotic state of the detector does not depend on it [10]. A direct consequence, verified without further computation, is stated in Section 6: D_{det} between a monopole-coupled and a derivative-coupled detector on the same worldline is strictly positive, which is exactly the artifact this clause exists to exclude, and its detection here is confirmation of the construction rather than a defect in it.

Gap referred to proper time. Ω must be defined against each detector’s own proper time, not against a shared coordinate time. Two detectors at different heights in a static spacetime, or moving at different speeds, experience a common coordinate-time interval as different proper-time intervals; building Ω from coordinate time would fold that ordinary redshift directly into D_{det} as a spurious worldline effect.

Switching profile. The function $\chi(\tau)$ leaves fingerprints in $\mathcal{F}_x(\Omega)$ that depend on its total duration and

its smoothness, independent of the trajectory it is attached to. Louko and Satz showed that a sharp step in χ produces a logarithmic divergence in the instantaneous transition rate at the switching instants, an artifact of the switching function’s discontinuous derivative, and that smoothing or spatially regularizing the switch restores a finite, comparable spectrum [11, 12]. Fewster, Juárez-Aubry, and Louko later sharpened the conditions under which a finite switching interval still recovers a well-defined, worldline-attributable response [13]. Matching total duration without matching the switching profile’s shape is not sufficient: D_{det} computed that way measures the shape of the shutter, not the shape of spacetime.

These four clauses are not decoration on Eqs. (5)–(6); they are the difference between a number that carries information about the worldlines and a number that carries information about laboratory bookkeeping. A protocol description for any application of D_{det} is incomplete until all four are stated explicitly and confirmed identical across the pair being compared, and Section 9 builds this requirement into the atlas’s kill criteria rather than leaving it as a caveat.

6 Where the spectrum is known exactly: thermal, adiabatic, and undefined

Any new object owes the literature a check that it reproduces the literature’s own solved cases before it is trusted anywhere else. Three such cases are available here without any new calculation.

Eternal uniform acceleration. For a worldline held at constant proper acceleration a for all proper time, coupled linearly to a massless scalar field in the ordinary vacuum with switching left on eternally, the excitation-branch response rate per unit proper time is the textbook Bose–Einstein form

$$\dot{\mathcal{F}}_a(\Omega) \propto \frac{\Omega}{e^{2\pi c\Omega/a} - 1}, \quad \Omega > 0, \quad (7)$$

with a proportionality constant fixed by the coupling strength and field normalization that cancels once the spectrum is normalized [1, 6]. The exponent $2\pi c\Omega/a$ is dimensionless, since $c\Omega$ has units of acceleration and dividing by a leaves a pure number; Eq. (7) is exactly the Planck form at temperature $T_U = \hbar a / (2\pi c k_B)$, recovering Eq. (3) as a limit this construction inherits rather than derives. Two worldlines held at accelerations $a_1 \neq a_2$ trace two differently scaled copies of the same universal curve; two worldlines held at the same acceleration trace the identical curve, and $D_{\text{det}} = 0$ between them by Proposition 4.1 for the most literal possible reason: the spectra coincide.

Slow time dependence. Barbado and Visser computed the response of a detector whose acceleration $a(\tau)$ varies with proper time, expanding in powers of $\dot{a}$, $\ddot{a}$, and higher proper-time derivatives [9]. Wherever the acceleration varies slowly compared to the timescale it locally sets for itself, the leading-order response at proper time τ matches the eternal-acceleration spectrum at the instantaneous value $a(\tau)$. This adiabatic recovery is a second inherited limit: D_{det} between an eternal trajectory at acceleration a and a slowly varying trajectory passing through a at one instant should vanish to leading order in the rate of change, with corrections controlled by the jerk and higher terms Barbado and Visser compute. No new derivation of this limit is performed in this paper; it is recorded here as a constraint any completed atlas calculation must reproduce.

Eternal inertial trajectory. An inertial worldline’s Wightman function is the flat-space vacuum two-point function, whose Fourier transform against $e^{-i\Omega\Delta\tau}$ has support only on $\Omega < 0$; this is the same vacuum-stability statement that forbids an unaccelerated atom in its ground state from spontaneously exciting in the field vacuum [7]. Consequently $\dot{\mathcal{F}}_{\text{inertial}}(\Omega) = 0$ identically for every $\Omega > 0$, exactly rather than approximately, and the normalization integral in Eq. (5) is exactly zero. p_{inertial} is 0/0: undefined. This is a different failure from the shared-infinity cancellation that makes an eternal accelerated rate finite — there the numerator and denominator both diverge and the ratio survives; here the numerator vanishes identically and there is no shared divergence for the ratio to rescue. The vacuum is genuinely empty from an inertial point of view; the failure belongs to D_{det} as a mathematical object comparing two distributions when one of them does not exist, not to the underlying physics. A finite switching interval restores a nonzero, finite transient spectrum [11, 12], but only at the cost of making the comparison depend explicitly on how the switch was shaped, which is precisely the fourth clause of Section 5 reasserting itself. This paper does not select a specific regularization for the eternal-inertial case; Section 10 states that choice as unresolved rather than making it silently.

Coupling-type mismatch. As a closed-form corollary of Moustos’s coupling-type result [10], a monopole-coupled and a derivative-coupled detector on one fixed worldline have normalized spectra that differ pointwise in Ω by construction, since the derivative coupling multiplies the response by an extra factor of Ω^2 before normalization. Two pointwise-distinct probability densities have strictly positive total-variation distance by definition, with no integral needing to be evaluated to know this; $D_{\text{det}} > 0$ between the two couplings on an identical trajectory is therefore an exact consequence of Eq. (6), not an estimate. This is the

unmatched-apparatus failure mode of Section 5, demonstrated rather than merely asserted.

7 The one calculation carried to a number

Where D_{det} is well defined, it can be evaluated, and the simplest non-trivial case is a comparison between two eternally, uniformly accelerated trajectories differing only slightly in acceleration. This section carries that one case through to an explicit coefficient, by hand, as a leading-order expansion; it is not a simulation and is not offered as a measurement of anything.

Writing $u \equiv c\Omega/a$ and using the exact normalization integral $\int_0^\infty x/(e^x - 1) dx = \pi^2/6$, the normalization of Eq. (7) gives $\int_0^\infty \Omega(e^{2\pi c\Omega/a} - 1)^{-1} d\Omega = a^2/(24c^2)$ (Appendix A carries out the substitution), so that

$$p_a(\Omega) = \frac{c}{a} g(u), \quad g(u) = \frac{24u}{e^{2\pi u} - 1}, \quad u = \frac{c\Omega}{a}. \quad (8)$$

Every eternally accelerated detector’s normalized spectrum is the same fixed curve g , rescaled by c/a ; g itself integrates to exactly 1 over $u > 0$, confirmed both analytically and by direct quadrature in the accompanying self-check program.

Proposition 7.1 (Leading-order scale-family sensitivity). *For a normalized scale family $p_a(\Omega) = (c/a)g(c\Omega/a)$ with g non-negative, integrable, and rising from $g(0) = 0$ to a single interior maximum of $ug(u)$ before decaying to zero, the total-variation distance between neighboring members satisfies, to leading order in δa ,*

$$D_{\text{det}}(a, a + \delta a) \approx \left[\max_{u>0} ug(u) \right] \frac{|\delta a|}{a}. \quad (9)$$

The derivation is given in full in Appendix A; the argument is the same unimodal-derivative trick that gives Wien’s displacement law, applied here to total-variation distance rather than to a spectral peak. Finding $\max_u ug(u)$ for the Unruh shape means solving the transcendental equation

$$e^{2\pi u} (1 - \pi u) = 1, \quad (10)$$

whose nontrivial positive root, located by bisection in the accompanying self-check program to better than 10^{-10} , is $u^* \approx 0.253633$, at which $u^*g(u^*) \approx 0.393700$. The leading-order sensitivity is therefore

$$D_{\text{det}}(a, a + \delta a) \approx 0.394 \frac{|\delta a|}{a}. \quad (11)$$

This coefficient is an exact leading term of a convergent expansion, obtained from a closed-form spectral shape

by ordinary calculus, with the only numerical step being a root-find of Eq. (10) to a stated tolerance; it is not fit to data and is not a claim about any measured system. It says two accelerations differing by one part in a thousand disagree, on this measure, by roughly four parts in ten thousand of normalized excitation content — a graded response to a graded physical difference, which is the qualitative behavior a usable distance should have, though this single coefficient is illustrative of that behavior rather than a demonstration that it holds at every point of the eventual atlas.

Proposition 7.1 was checked against the definition of D_{det} it claims to approximate, not only against its own algebra: the self-check program builds p_a and $p_{a+\delta a}$ as explicit discretized spectra for several shrinking values of $\delta a/a$ and computes the brute-force total-variation distance directly from Eq. (6), without using Eq. (9). The ratio of the brute-force distance to the closed-form prediction moved from 0.9950 at $\delta a/a = 10^{-2}$ to 0.9995 at $\delta a/a = 10^{-3}$ in that run, converging toward 1 as the expansion parameter shrinks, exactly as a correct leading-order formula should. This is a check of internal consistency between two independent routes to the same quantity, not an experimental confirmation.

8 A curved table and the one real laboratory analogue

Nothing in Eqs. (4)–(6) requires flat spacetime. The response functional is written entirely in terms of proper time and the Wightman function evaluated along a trajectory, both of which are well defined for a stationary worldline in a curved, static exterior. What changes in curved spacetime, as already flagged under the first clause of Section 5, is that the geometry alone no longer fixes the field state: a Schwarzschild exterior supports the Boulware vacuum, empty at infinity and singular at the horizon; the Hartle–Hawking state, thermal everywhere and matching a black hole in equilibrium with its own radiation; and the Unruh state, matching outgoing Hawking flux with no incoming radiation from infinity.

Hodgkinson, Louko, and Ottewill already computed exactly this kind of comparison, without reference to D_{det} : a static detector held at fixed areal radius and a detector on a circular geodesic at the same radius, both coupled identically to a massless scalar field, evaluated in all three vacuum states [14]. Their published spectra differ measurably between the static and circular-geodesic trajectories at matched radius and matched field state. No new curved-spacetime response is computed in this paper; the claim available here is narrower and does not require one: applying Eq. (6) to their already-published spectra would compress a comparison presently reported as pairs of curves into a single number per vacuum state, and the four clauses of Sec-

tion 5 name explicitly the conditions — same field state, same coupling, same locally measured gap, same switching — their comparison already had to hold fixed for the difference to mean hovering-versus-orbiting rather than mismatched detectors. Carrying this recomputation out is listed as unresolved work in Section 10, not performed here.

The idealizations above — eternal switching, exactly constant or exactly circular trajectories — have no experimental instantiation at any accessible acceleration; laboratory accelerometers reach nowhere near the scales in Eq. (3). The one setting where the acceleration entering T_U becomes large enough to matter is a relativistic electron storage ring. For an ultra-relativistic beam of Lorentz factor γ circulating at radius r , the proper centripetal acceleration is $a \approx \gamma^2 c^2 / r$; for $\gamma \sim 2 \times 10^5$ and r of several kilometers, typical of a large electron-positron collider, this gives $a \sim 8 \times 10^{23} \text{ m s}^{-2}$ and, by Eq. (3), $T_U \sim 3 \times 10^3 \text{ K}$ — a value the accompanying self-check program reproduces directly from the stated formula using CODATA constants, obtaining $T_U \approx 3.2 \times 10^3 \text{ K}$ at that acceleration, against an Earth-surface value of $T_U \approx 4 \times 10^{-20} \text{ K}$ at the same formula’s other extreme.

Bell and Leinaas pointed out that the accessible observable in a storage ring is the equilibrium spin polarization of the circulating electrons, not a literal thermometer reading: polarization is set by the competition between the Sokolov–Ternov synchrotron spin-flip mechanism and any Unruh-like depolarizing correction sourced by the acceleration itself [15]. No experiment has isolated an Unruh-like signal from this system: measured polarizations match the standard Sokolov–Ternov prediction, with ordinary QED corrections, at the percent level, and circular motion is not linear acceleration, so the field a circulating charge experiences is not simply a Rindler wedge — a point Bell and Leinaas themselves stressed. What the storage-ring case supplies honestly is a bound: percent-level agreement between observed and predicted polarization limits any anomalous acceleration-induced departure from ordinary detector response to that same percent level, in this one geometry, at this one acceleration scale. Separately, proposals to extract Unruh-like signatures through geometric phases [17] and through optimal quantum-estimation strategies applied to accelerated probes [16] indicate that a metrological, rather than thermometric, route to bounding these effects is an active direction; neither proposal is applied to D_{det} in this paper, and no claim is made here that either currently constrains it.

9 The atlas protocol, its kill criteria, and the rival accounts it must defeat

The numerical deliverable this construction is built toward is an atlas of D_{det} over a registered grid: proper acceleration a (or, for the time-dependent case, an acceleration profile $a(\tau)$ following Barbado and Visser [9]), switching duration, detector gap Ω , field state (flat vacuum, thermal, and, for stationary curved trajectories, the Boulware, Hartle–Hawking, and Unruh states), and switching-profile shape at fixed duration. No point of that grid has been evaluated for this draft; Section 10 names precisely what stands between the definitions above and that computation.

Definition 9.1 (Kill criteria). D_{det} is rejected as a physically meaningful worldline distance if either of the following holds after implementation:

1. its value changes under a reparametrization of the auxiliary curve parameter of either worldline, holding the physical trajectory, the field state, the coupling, the proper-frame gap, and the switching profile fixed; or
2. its value across the registered grid is dominated by differences in coupling type, switching profile, or gap convention between the two detectors, rather than by the difference between the trajectories once the four clauses of Section 5 are enforced.

The first condition is addressed structurally, not empirically: Section 4 shows D_{det} is built entirely from proper time and proper-frame quantities, leaving no co-ordinate freedom in Eq. (6) for a reparametrization to act on. A numerical implementation must still confirm this survives discretization — a badly chosen Ω -grid or quadrature rule could reintroduce spurious parameter-dependence that the continuum definition does not have — and this convergence check belongs in the output schema below, not only in the analytic argument. The second condition was tested directly, not assumed, for one case in Section 6: a mismatched coupling type moves D_{det} on an identical fixed worldline, in closed form. Finding that the atlas is dominated everywhere by apparatus mismatch rather than by trajectory differences, once all four clauses are matched, would kill the construction rather than merely narrow it.

Output schema. Every saved atlas entry must record the full protocol it corresponds to (trajectory family and parameters, field state, coupling type, gap, switching profile and duration), a numerical-quadrature error estimate for the Ω -integral in Eq. (6), the reparametrization-invariance residual obtained by recomputing the same entry under an alternative curve parametrization, code commit and dependency lock,

and a flag distinguishing a converged value from a point excluded for non-convergence. No atlas figure may be drawn from an unconverged or unflagged point, and no missing grid point may be interpolated for display.

Rival accounts. Three readings of a completed atlas would leave D_{det} without independent content, and each has a decisive comparison rather than a rhetorical dismissal. A *redshift-mimicry* reading holds that any nonzero atlas entry secretly encodes a proper-time bookkeeping error rather than a trajectory effect; it is checked by confirming the reparametrization residual in the output schema is consistent with zero within quadrature error at every grid point. An *apparatus-dominance* reading holds that the atlas is shaped more strongly by switching-profile or coupling choices than by acceleration, duration, or field-state choices; it is checked by comparing the atlas’s sensitivity to each of the five grid dimensions directly, using the same protocol at fixed trajectory pairs while varying one clause at a time. A *redundancy* reading holds that D_{det} , once computed, tracks an already-known distinguishability measure between the same two field states or thermal parameters — for instance a Bures or relative-entropy distance between the reduced detector states — closely enough that it adds no separate operational content beyond relabeling an existing quantity; it is checked by comparing D_{det} against such a measure on the same trajectory pairs and requiring at least one region of the grid where the two classifications disagree. The proposal advanced here stays deliberately narrow in scope: a computable distance between fully specified detector protocols, rather than a general equivalence principle linking acceleration to quantum-state indistinguishability at the level of covariance. Failing the redundancy comparison would mean this paper’s construction, however cleanly it is defined, is not doing independent work.

10 What must exist before this is a results paper

Everything computed in this draft is an exact identity, a closed-form limit, or a leading-order expansion checked against its own defining integral. Nothing in it is the disagreement atlas the construction is meant to produce, and this section states, without qualification, what separates the two.

The first gap is a validated finite-switching-function solver. Section 6 shows that the eternal-inertial trajectory leaves D_{det} undefined and that a finite switch is the only known repair, but this draft does not choose a switching family, does not fix a duration convention, and does not verify a solver against the closed forms Louko, Satz, and their collaborators already established for simple profiles [11–13]. Until that solver exists and

reproduces those closed forms to a stated numerical tolerance, no atlas entry involving a finite-duration switch — which is to say, no atlas entry at all, since the eternal limit is itself only a limit — can be trusted.

The second gap is curved spacetime. Section 8 argues that D_{det} could compress the static-versus-circular-geodesic comparison already published by Hodgkinson, Louko, and Ottewill into a single number per vacuum state [14], but this draft has not implemented the Boulware, Hartle–Hawking, or Unruh Wightman functions needed to reproduce their spectra, let alone extended the comparison to trajectories or radii they did not consider. That implementation, and its agreement with their published curves before any new claim is built on it, is a precondition for every curved-spacetime entry in the atlas.

The third gap is the regularization choice for the eternal-inertial degeneracy itself. Section 6 declines to select a specific finite-switching family for that case, on the grounds that the choice measurably affects the result and should not be made informally in a scaffold section. A registered choice, with its own justification and its own sensitivity check against alternative reasonable choices, has to precede any atlas entry that compares an accelerated trajectory against an inertial one at short but finite switching duration.

The fourth gap is numerical infrastructure that does not yet exist: a quadrature error budget for the Ω -integral in Eq. (6) at finite grid resolution, the reparametrization-residual check described in the atlas output schema of Section 9 run against an actual discretized implementation rather than argued analytically, a dependency lock for whatever numerical-integration and root-finding stack the finite-switching and curved-spacetime solvers require, and a recorded hardware and software environment for reproducibility. The self-check program accompanying this paper is deliberately zero-dependency and checks only the exact and closed-form claims made in Sections 3–7; it is not, and does not claim to be, that infrastructure.

The fifth gap is the redundancy comparison named in Section 9: computing an existing distinguishability measure, such as a Bures or relative-entropy distance between the reduced detector states, on the same trajectory pairs used for the atlas, and confirming that D_{det} ’s classification disagrees with it somewhere in the grid. Without this comparison run to completion, the claim that D_{det} supplies information beyond an existing measure remains an intention rather than a finding.

The sixth gap is a named external reviewer with expertise in Unruh–DeWitt detector calculations and quantum field theory in curved spacetime, who has not yet been approached. Concretely, before any numerical result from this construction is reported as a finding rather than a scaffold, the following must all be true: a switching solver validated against closed

forms; curved-spacetime state implementations validated against Hodgkinson, Louko, and Ottewill’s published spectra; a registered and justified regularization for the eternal-inertial case; a quadrature and reparametrization error budget computed numerically rather than argued analytically; a completed redundancy comparison against an existing distinguishability measure; and an external review of all of the above. Any one of these being incomplete is sufficient reason to withhold a numerical claim, and as of this draft, all six are incomplete.

11 A registered distance, not yet a measured one

This paper defines D_{det} , a total-variation distance between the normalized excitation spectra two Unruh–DeWitt detectors report, proves it is a genuine metric bounded in $[0, 1]$, and shows it is invariant under world-line reparametrization and joint Poincaré transformation by construction rather than by numerical accident. It states four matching conditions that must hold before a nonzero reading is attributed to the trajectories rather than to the apparatus, and demonstrates one of those conditions failing on purpose, in closed form, as confirmation that the construction flags the artifact it is meant to catch. It recovers the exact Planckian Unruh spectrum and the adiabatic limit as inherited results, derives a closed-form leading-order sensitivity coefficient between neighboring accelerations from a Wien-type transcendental equation, and identifies a genuine boundary of the construction: the eternal inertial trajectory, where D_{det} fails to be defined at all rather than merely coming out small.

None of this is the atlas. No sweep over acceleration, duration, gap, field state, or switching profile has been run; no curved-spacetime response has been newly computed; no comparison against an existing distinguishability measure has been completed; and no external reviewer has examined any of it. Section 10 states this without softening it, because a construction that only reports where it has not yet been tested is more useful, not less, than one that stays silent about the gap.

If the atlas is eventually computed and passes the kill criteria of Section 9, D_{det} would offer a single number, with declared units, declared invariances, and a declared failure mode, in place of a qualitative appeal to observer-dependent particle content wherever two detector protocols are compared. If it fails — if the atlas turns out to be dominated by apparatus mismatch, or to add no classification beyond an existing measure — the inherited baseline recovered in Sections 2 and 6 is untouched, and the field’s existing accounts of the Bogoliubov ambiguity and the Unruh effect stand exactly as they did before this paper was written. Either out-

come leaves something learned; only one of them leaves a new number worth publishing, and this draft does not pretend to have reached it.

A Derivation of the scale-family sensitivity coefficient

This appendix carries out, in full, the calculation stated in Section 7.

A.1 The normalization integral

Starting from Eq. (7) with the proportionality constant set to one (it cancels on normalization), substitute $x = 2\pi c\Omega/a$, so $\Omega = ax/(2\pi c)$ and $d\Omega = \frac{a}{2\pi c} dx$:

$$\int_0^\infty \frac{\Omega d\Omega}{e^{2\pi c\Omega/a} - 1} = \left(\frac{a}{2\pi c}\right)^2 \int_0^\infty \frac{x dx}{e^x - 1} \quad (12)$$

$$= \left(\frac{a}{2\pi c}\right)^2 \frac{\pi^2}{6} = \frac{a^2}{24c^2}, \quad (13)$$

using the standard result $\int_0^\infty x(e^x - 1)^{-1} dx = \pi^2/6$. Dividing Eq. (7) by this normalization and substituting $u = c\Omega/a$ gives Eq. (8), $p_a(\Omega) = (c/a)g(u)$ with $g(u) = 24u/(e^{2\pi u} - 1)$. A direct check that g integrates to 1 over $u \in (0, \infty)$, by the same substitution run in reverse, confirms p_a is correctly normalized; this integral is also verified numerically in the accompanying self-check program.

A.2 Leading-order total-variation distance for a scale family

Let $p_a(\Omega) = (c/a)g(c\Omega/a)$ for a fixed, non-negative, integrable g with $\int_0^\infty g(u) du = 1$. Write $u = c\Omega/a$, so $\partial u/\partial a = -c\Omega/a^2 = -u/a$. Differentiating in a at fixed Ω ,

$$\frac{\partial p_a}{\partial a} = -\frac{c}{a^2}g(u) + \frac{c}{a}g'(u)\left(-\frac{u}{a}\right) \quad (14)$$

$$= -\frac{c}{a^2}\left[g(u) + u g'(u)\right] = -\frac{c}{a^2}h'(u), \quad (15)$$

where $h(u) \equiv u g(u)$. Since $\int_0^\infty p_a(\Omega) d\Omega = 1$ for every a , differentiating under the integral sign gives $\int_0^\infty \partial_a p_a(\Omega) d\Omega = 0$; the perturbation is mean-zero, as it must be for a family of normalized densities. To leading order in δa ,

$$D_{\text{det}}(a, a+\delta a) \approx \frac{|\delta a|}{2} \int_0^\infty \left| \frac{\partial p_a}{\partial a} \right| d\Omega = \frac{|\delta a|}{2a} \int_0^\infty |h'(u)| du, \quad (16)$$

after substituting $d\Omega = (a/c) du$ in the last integral. If h rises from $h(0) = 0$ to a single interior maximum at $u = u^*$ and decays to $h(\infty) = 0$ — true for $h(u) = 24u^2/(e^{2\pi u} - 1)$, which vanishes at $u = 0$, is positive

for $u > 0$, and decays exponentially as $u \rightarrow \infty$ — then h is unimodal and

$$\int_0^\infty |h'(u)| du = [h(u^*) - h(0)] + [h(u^*) - h(\infty)] = 2h(u^*). \quad (17)$$

Substituting back gives $D_{\text{det}}(a, a+\delta a) \approx \frac{|\delta a|}{2a} \cdot 2h(u^*) = h(u^*)|\delta a|/a = u^*g(u^*)|\delta a|/a$, which is Eq. (9) with $\max_u u g(u) = h(u^*)$.

A.3 The transcendental equation for u^*

Writing $h(u) = 24u^2(e^{2\pi u} - 1)^{-1}$ and setting $h'(u^*) = 0$,

$$h'(u) = \frac{24}{(e^{2\pi u} - 1)^2} \left[2u(e^{2\pi u} - 1) - 2\pi u^2 e^{2\pi u} \right]. \quad (18)$$

For $u > 0$ the bracket vanishes exactly when $(e^{2\pi u} - 1) - \pi u e^{2\pi u} = 0$, i.e.

$$e^{2\pi u} (1 - \pi u) = 1, \quad (19)$$

which is Eq. (10). Write $f(u) = e^{2\pi u}(1 - \pi u)$ for the left-hand side. Then

$$f'(u) = 2\pi e^{2\pi u}(1 - \pi u) - \pi e^{2\pi u} = \pi e^{2\pi u}(1 - 2\pi u), \quad (20)$$

so f is *not* monotone on $(0, 1/\pi)$: it increases strictly on $(0, \frac{1}{2\pi})$, attains the maximum $f(\frac{1}{2\pi}) = e/2 \approx 1.3591$, and decreases strictly on $(\frac{1}{2\pi}, 1/\pi)$. Uniqueness follows from that shape rather than from monotonicity. On $(0, \frac{1}{2\pi}]$ we have $f(u) > f(0) = 1$, so the equation $f(u) = 1$ has no root there. On $(\frac{1}{2\pi}, 1/\pi)$, f decreases strictly from $e/2 > 1$ to $f(1/\pi) = e^2(1 - 1) = 0 < 1$, so by the intermediate value theorem it meets the level 1 exactly once. The nontrivial root is therefore unique and lies in $(\frac{1}{2\pi}, 1/\pi) \approx (0.1592, 0.3183)$. Note that $f(1/\pi)$ equals zero exactly, not a negative number; the strict inequality needed for the intermediate value argument is $0 < 1$. The accompanying self-check program locates it by bisection on $(10^{-6}, 1/\pi - 10^{-9})$ to a tolerance of 10^{-12} in the defining equation, obtaining $u^* = 0.253633\dots$ and $h(u^*) = u^*g(u^*) = 0.393700\dots$, rounding to the 0.394 quoted in Eq. (11).

B Computation and test ledger

The scaffold program `check_detector_distance.py` performs no field-theoretic mode sum, no simulated switching function, and no curved-spacetime Wightman-function evaluation. In its default `-self-check` mode it performs the following exact and closed-form analytic checks, each described in the sections above:

1. non-negativity, symmetry, boundedness in $[0, 1]$, exact identity of indiscernibles, and the triangle

inequality for D_{det} , tested on a family of synthetic normalized spectra — Bose–Einstein-shaped curves at several scales, Gaussian bumps, and random densities — so that the metric axioms of Proposition 4.1 are checked against inputs not assumed to be physical, not only against the accelerated spectrum itself;

2. exact confirmation that identical trajectories, i.e. identical discretized spectra, give $D_{\text{det}} = 0$ to floating-point precision;
3. reproduction of $k_B T_U = \hbar a / (2\pi c)$ from Eq. (3) using CODATA physical constants, cross-checked against the textbook one-Earth-gravity benchmark of order 4×10^{-20} K and against the storage-ring-scale estimate of Section 8;
4. numerical confirmation that $\int_0^\infty x(e^x - 1)^{-1} dx = \pi^2/6$, that the normalization integral of Eq. (7) equals $a^2/(24c^2)$ for arbitrary positive test values of a and c , and that $g(u)$ in Eq. (8) integrates to exactly 1;
5. a bisection solution of the transcendental equation, Eq. (10), confirming $u^* \approx 0.253633$ and $u^*g(u^*) \approx 0.393700$, together with a brute-force check that builds discretized p_a and $p_{a+\delta a}$ spectra directly and confirms their explicit total-variation distance converges to the closed-form leading-order prediction of Eq. (9) as $\delta a/a \rightarrow 0$;
6. confirmation that the eternal-inertial spectrum, identically zero for all $\Omega > 0$, is correctly refused by the normalization routine rather than silently returning a spurious distance, matching the exact-limits discussion of Section 6.

The saved output, `scaffold_validation_output.txt`, is the program’s real captured standard output from a `-self-check` run and is labeled throughout as scaffold validation rather than as a physical result; `physics_atlas_computed` is recorded as `false` in that output, and no field in it reports an atlas value. The program’s `-simulate` option exits with status 2 and prints the six-item list of unimplemented prerequisites given in Section 10; it never proceeds to produce a number. A later implementation must not reuse the name `check_detector_distance.py` while silently relaxing these six checks; it must add the finite-switching, curved-spacetime, and redundancy-comparison tests named in Section 10, adopt a versioned output schema, and replace `requirements.lock` with an exact, platform-tested dependency lock for whichever numerical stack the atlas computation requires.

References

- [1] W. G. Unruh (1976). Notes on black-hole evaporation. *Physical Review D* 14, 870–892. <https://doi.org/10.1103/PhysRevD.14.870>.
- [2] S. A. Fulling (1973). Nonuniqueness of canonical field quantization in Riemannian space-time. *Physical Review D* 7, 2850–2862. <https://doi.org/10.1103/PhysRevD.7.2850>.
- [3] P. C. W. Davies (1975). Scalar production in Schwarzschild and Rindler metrics. *Journal of Physics A: Mathematical and General* 8, 609–616. <https://doi.org/10.1088/0305-4470/8/4/022>.
- [4] S. W. Hawking (1975). Particle creation by black holes. *Communications in Mathematical Physics* 43, 199–220. <https://doi.org/10.1007/BF02345020>.
- [5] W. G. Unruh and R. M. Wald (1984). What happens when an accelerating observer detects a Rindler particle. *Physical Review D* 29, 1047–1056. <https://doi.org/10.1103/PhysRevD.29.1047>.
- [6] L. C. B. Crispino, A. Higuchi, and G. E. A. Matsas (2008). The Unruh effect and its applications. *Reviews of Modern Physics* 80, 787–838. <https://doi.org/10.1103/RevModPhys.80.787>.
- [7] N. D. Birrell and P. C. W. Davies (1982). *Quantum Fields in Curved Space*. Cambridge University Press. <https://doi.org/10.1017/CB09780511622632>.
- [8] S. Takagi (1986). Vacuum noise and stress induced by uniform acceleration: Hawking–Unruh effect in Rindler manifold of arbitrary dimension. *Progress of Theoretical Physics Supplement* 88, 1–142. <https://doi.org/10.1143/PTPS.88.1>.
- [9] L. C. Barbado and M. Visser (2012). Unruh–DeWitt detector event rate for trajectories with time-dependent acceleration. *Physical Review D* 86, 084011. <https://doi.org/10.1103/PhysRevD.86.084011>.
- [10] D. Moustos (2018). Asymptotic states of accelerated detectors and universality of the Unruh effect. *Physical Review D* 98, 065006. <https://doi.org/10.1103/PhysRevD.98.065006>.
- [11] J. Louko and A. Satz (2006). How often does the Unruh–DeWitt detector click? Regularization by a spatial profile. *Classical and Quantum Gravity* 23, 6321–6343. <https://doi.org/10.1088/0264-9381/23/22/015>.

- [12] J. Louko and A. Satz (2007). Then again, how often does the Unruh–DeWitt detector click if we switch it carefully? *Classical and Quantum Gravity* 24, 1719–1731. <https://doi.org/10.1088/0264-9381/24/7/003>.
- [13] C. J. Fewster, B. A. Juárez-Aubry, and J. Louko (2016). Waiting for Unruh. *Classical and Quantum Gravity* 33, 165003. <https://doi.org/10.1088/0264-9381/33/16/165003>.
- [14] L. Hodgkinson, J. Louko, and A. C. Ottewill (2014). Static detectors and circular-geodesic detectors on the Schwarzschild black hole. *Physical Review D* 89, 104002. <https://doi.org/10.1103/PhysRevD.89.104002>.
- [15] J. S. Bell and J. M. Leinaas (1983). Electrons as accelerated thermometers. *Nuclear Physics B* 212, 131–150. [https://doi.org/10.1016/0550-3213\(83\)90601-6](https://doi.org/10.1016/0550-3213(83)90601-6).
- [16] M. Aspachs, G. Adesso, and I. Fuentes (2010). Optimal quantum estimation of the Unruh–Hawking effect. *Physical Review Letters* 105, 151301. <https://doi.org/10.1103/PhysRevLett.105.151301>.
- [17] E. Martín-Martínez, I. Fuentes, and R. B. Mann (2011). Using Berry’s phase to detect the Unruh effect at lower accelerations. *Physical Review Letters* 107, 131301. <https://doi.org/10.1103/PhysRevLett.107.131301>.