

There Is No Relativistic Position, Only an Atlas of Position Experiments

Registered research scaffold — no numerical result

Brecht Corbeel

Absolute Digital Publishers

absolutedigitalpublishers.com

September 20, 2026

Abstract

A boosted, spinning wavepacket does not have one position; it has several, each fixed by a named observer, slicing convention, or measurement protocol, and each correct within its own scope. This draft registers an operational way to compare those charts rather than adjudicate among them. Working from the hypersurface-dependent energy centroid $X^\mu(u, \Sigma)$ built on a body's stress-energy tensor, we hold the spin supplementary condition fixed and vary the observer, recovering the Møller world-tube shift as an exact identity of special relativity and deriving, by a direct Cauchy–Schwarz argument on the underlying cross product, that the shift saturates at the spin-to-mass length $\rho_M = |\mathbf{S}|/(Mc)$ precisely when spin and boost are transverse. Holding the observer fixed and varying the spin supplementary condition instead, we state the Pryce trade-off between the covariant centre of energy and the Newton–Wigner operator, and derive, from the exact commutativity of mixed partial derivatives rather than by assumption, that the Newton–Wigner momentum-space representation has commuting position components; the same rotation subalgebra is checked directly by matrix multiplication on a spin- $\frac{1}{2}$ representation. We tie the same length scale to the zitterbewegung frequency of the naive Dirac coordinate, recovering the nonrelativistic limit exactly rather than asymptotically in both cases. We state precisely, and do not overstate, what the Hegerfeldt and Reeh–Schlieder theorems and Malament's axiomatic argument forbid: a sharp, causally propagating, exactly localized position operator, for any of the charts above, not a prohibition on POVM-based operational localization or on superluminal signalling in the full field theory. We exhibit an explicit Gaussian-smeared localization POVM and verify its positivity and normalization by spectral construction and by quadrature, and we derive a closed-form family of postselected weak values for a spin observable, showing in closed form how the estimator diverges as the postselection

approaches orthogonality and departs the bound that constrains the sharp charts, without contradiction. We define a normalized disagreement functional between charts, prove it is non-negative, symmetric, and bounded, and register — without computing — the kill criteria, output schema, and parameter grid for the atlas the construction is built toward: a boosted Gaussian-wavepacket comparison of these charts together with a spin-Hall-of-light beam-centroid worked example. No such atlas is computed here; the spin-Hall effect and the trapped-ion Dirac simulation cited below are existing measurements, not results of this paper. A companion self-check program verifies the metric-relevant properties, the two commutators, the POVM, and the exact limits described above, and records `"scientific_result": false`; the atlas itself, the beam-centroid computation, the curved-spacetime extension, and external subject-matter review remain future work stated as such. **Key-**

words: Pryce classification; Newton–Wigner position operator; Møller world tube; Hegerfeldt theorem; Reeh–Schlieder theorem; localization POVM; weak value; spin Hall effect of light.

Series: Revolutionary Life 75, Series D, D07.

Status: Draft scaffold. The boosted-wavepacket atlas, the spin-Hall beam-centroid numerical example, the curved-spacetime extension, and external review are incomplete.

1 The object six instruments are asked to locate

Send a boosted, spin-carrying wavepacket past a bank of instruments and ask each one, separately, where the wavepacket is. An accelerometer riding with the packet's own rest frame reports a point fixed by the packet's mass distribution in that frame. A stress-

energy centroid computed on a tilted simultaneity surface reports a different point, shifted by an amount that depends on the surface's tilt as well as on the spin. A commuting-component position operator built for exactly this purpose reports a third point, frame-dependent in a way the first two are not. A postselected weak-value estimator can report a value that is not a spacetime point at all. A detector array reports neither an operator eigenvalue nor a classical centroid but a probability distribution over which cell fired. None of these five answers is a mistake. Each is the correct output of a fully specified protocol, and the protocols are not interchangeable.

This is not a new observation. Pryce classified the reasonable relativistic generalizations of the Newtonian centroid in 1948 [1]; Newton and Wigner built a position operator with commuting components the following year, at a stated cost [2]; Fleming made the cost explicit a generation later [4]. What this paper registers is not a sixth construction to add to that list, and not a verdict on which of the five is "really" the position. It registers a single operational object, a normalized disagreement between two named charts, together with the length scale that disagreement is measured against, the properties that object can be proved to have before any number is computed with it, and the specific numerical program — not run here — that would turn the object into an atlas.

Begin with the least controversial candidate. For a body made of mass elements m_a at positions $\mathbf{r}_a$, the Newtonian centroid is

$$\mathbf{R}_N = \frac{\sum_a m_a \mathbf{r}_a}{\sum_a m_a}. \quad (1)$$

Every term on the right is evaluated at one shared value of a time coordinate t , and which set of events counts as sharing a t -value is exactly the freedom special relativity revokes: two observers in relative motion slice the same worldlines into different families of simultaneous events, and Eq. (1) evaluated on one family need not agree with Eq. (1) evaluated on another. $\mathbf{R}_N$ is not wrong; it is silent about which slice it was taken on.

Naming the slice explicitly gives this paper's first genuine chart. Let Σ be a spacelike hypersurface and u^μ a unit timelike vector, $u^\mu u_\mu = -1$, normal to Σ , representing the simultaneity convention of one observer. For a body with stress-energy tensor $T^{\alpha\beta}$, define the hypersurface energy centroid

$$X^\mu(u, \Sigma) = \frac{\int_\Sigma x^\mu T^{\alpha\beta} u_\alpha d\Sigma_\beta}{\int_\Sigma T^{\alpha\beta} u_\alpha d\Sigma_\beta}. \quad (2)$$

Every symbol in Eq. (2) carries a declared type: $T^{\alpha\beta}$ is an energy-momentum density, whatever combination of

energy, momentum, and volume units a given convention for x^0 assigns it; $d\Sigma_\beta$ is a directed three-volume element, units of length cubed; u_α is dimensionless once normalized to $u^\mu u_\mu = -1$; x^μ is a spacetime coordinate, units of length. Whatever units the weighting factor $T^{\alpha\beta} u_\alpha d\Sigma_\beta$ carries under a given convention, it appears identically in the numerator and denominator of Eq. (2) and cancels exactly, leaving $X^\mu(u, \Sigma)$ with units of length regardless of that convention, a genuine spacetime point rather than an artifact of the weighting. That cancellation is a check passed by the structure of Eq. (2) itself, not by a coincidence of the numbers fed into it, and it is the first of several dimensional audits this paper performs explicitly rather than leaving implicit. Section 2 fixes the remaining convention this paper uses, that $P^\mu = \int_\Sigma T^{\mu\nu} d\Sigma_\nu$ carries the units of ordinary momentum in every component, the standard choice that makes $M = \sqrt{-P^\mu P_\mu}/c$ come out in units of mass.

$X^\mu(u, \Sigma)$ takes two pieces of auxiliary data, an observer and a slice, and returns a point. Different choices of (u, Σ) applied to the identical body — identical $T^{\alpha\beta}$ field, identical total four-momentum $P^\mu = \int_\Sigma T^{\mu\nu} d\Sigma_\nu$, identical total angular momentum $J^{\mu\nu}$ — can and generically do return different points. Section 2 states precisely what stays fixed while (u, Σ) varies, and Section 3 derives exactly how far the returned point can move.

This paper's scope is fixed here and does not widen later. It stays in flat spacetime throughout; it treats every chart's defining protocol as already established physics, none of it proposed here; and it does not compute the atlas that would compare these charts numerically across a range of boosts, spins, and hypersurface tilts. What it delivers is the length scale the atlas would be measured against, two commutators recomputed rather than copied, one closed-form saturation bound, one Gaussian localization POVM checked for positivity and normalization, one closed-form weak-value family exhibiting the estimator's own failure mode, a normalized disagreement functional with proved properties, and the registered — not run — numerical protocol and kill criteria that would turn all of this into a result. Section 10 states plainly what stands between this draft and that result.

2 Charges no chart is permitted to move

Before comparing what $X^\mu(u, \Sigma)$ returns for two different choices of (u, Σ) , it is worth being precise about what does not change between them. The total four-momentum $P^\mu = \int_\Sigma T^{\mu\nu} d\Sigma_\nu$ and the total angular momentum $J^{\mu\nu}$ are the Poincaré charges of one physical state: they are fixed by the field configuration itself, not

by which hypersurface is used to evaluate an integral over it, and Gauss's theorem applied to the conserved current $\partial_\alpha T^{\alpha\beta} = 0$ guarantees P^μ is the same number on every Cauchy surface of the same spacetime region. No relabeling of "the centroid" performed in this paper touches P^μ or $J^{\mu\nu}$. What moves, when (u, Σ) moves, is only the split of $J^{\mu\nu}$ into an orbital piece $L^{\mu\nu} = X^{[\mu} P^{\nu]}$ built on whichever centroid X^μ is current, and a spin piece $S^{\mu\nu} = J^{\mu\nu} - L^{\mu\nu}$ that absorbs the difference. This is the single invariant every chart in this paper sits on top of: physics agrees completely on what the body has, and the charts disagree only about which point should be credited with carrying the orbital share.

That fixed spin content supplies the length scale every disagreement below is measured against. Let $M = \sqrt{-P^\mu P_\mu}/c$ be the rest mass, with P^μ normalized so that every component carries units of momentum, $P^\mu P_\mu = -M^2 c^2$, and $|\mathbf{S}|$ the magnitude of the spin angular momentum evaluated in the body's own rest frame, the part of $J^{\mu\nu}$ left over once $L^{\mu\nu}$ is subtracted at the rest-frame centroid. Define

$$\rho_M = \frac{|\mathbf{S}|}{Mc}. \quad (3)$$

By inspection, angular momentum divided by momentum, ρ_M has units of length, and it is built from nothing but the two Poincaré-invariant labels every observer agrees on. It does not depend on which chart is used to locate the body, because it is built from M and $|\mathbf{S}|$, not from any X^μ . What ρ_M measures is established in Section 3: it is the radius of the region within which one particular pair of charts is forced to agree, to the extent it is forced to agree at all.

The remainder of this paper is organized around three questions about disagreement, in order of how tightly the disagreement can be controlled. Section 3 asks how $X^\mu(u, \Sigma)$ moves when the spin supplementary condition is held fixed and the observer varies, and derives an exact bound saturated at ρ_M . Section 4 asks how the centroid moves when the observer is held fixed and the supplementary condition itself is changed, comparing the covariant centre of energy against the Newton–Wigner operator, and Section 5 extends that comparison to the naive Dirac coordinate through zitterbewegung. Section 6 identifies where every chart built this way stops being exactly available at all. Sections 7 and 8 then introduce two charts of a structurally different kind, a detector-tied probability measure and a postselected estimator, neither of which is bound by the same disk. Section 9 collects the well-behaved charts into a single normalized disagreement functional and registers, without computing, the atlas that would test it against real boosts and real spins.

3 A disk with a radius derived, not assumed

Fix the spin supplementary condition that anchors the centroid to the body's own rest frame, $S^{\alpha\beta} P_\beta = 0$ in the standard notation of the spinning-body literature, and ask how $X^\mu(u, \Sigma)$ in Eq. (2) moves as the observer u is boosted away from the body's own rest frame at velocity $\mathbf{v}$, the supplementary condition held fixed throughout. Costa and Natário's synthesis of the classical spinning-body literature gives the displacement in closed form,

$$\Delta \mathbf{x} = \frac{\mathbf{S} \times \mathbf{v}}{Mc^2}, \quad (4)$$

transverse to both the spin and the boost [18]. This identity is not derived here; it is an exact result of special relativity applied to the classical spinning body, quoted as inherited baseline exactly as Section 2 inherits P^μ conservation. What this paper adds is the bound on Eq. (4) and the check that the bound is exact rather than approximate.

Proposition 3.1 (Saturation of the boost-transition shift). *For any spin vector $\mathbf{S}$ and any velocity $\mathbf{v}$ with $|\mathbf{v}| < c$, $|\Delta \mathbf{x}| \leq \rho_M (|\mathbf{v}|/c)$, with equality if and only if $\mathbf{S} \perp \mathbf{v}$.*

Proof. Lagrange's identity gives $|\mathbf{S} \times \mathbf{v}|^2 = |\mathbf{S}|^2 |\mathbf{v}|^2 - (\mathbf{S} \cdot \mathbf{v})^2$. Since $(\mathbf{S} \cdot \mathbf{v})^2 \geq 0$, $|\mathbf{S} \times \mathbf{v}| \leq |\mathbf{S}| |\mathbf{v}|$, with equality exactly when $\mathbf{S} \cdot \mathbf{v} = 0$. Dividing Eq. (4) by Mc^2 and using this bound,

$$|\Delta \mathbf{x}| = \frac{|\mathbf{S} \times \mathbf{v}|}{Mc^2} \leq \frac{|\mathbf{S}| |\mathbf{v}|}{Mc^2} = \frac{|\mathbf{S}|}{Mc} \cdot \frac{|\mathbf{v}|}{c} = \rho_M \frac{|\mathbf{v}|}{c},$$

with equality under the same condition, $\mathbf{S} \perp \mathbf{v}$. $\square$

As $|\mathbf{v}|/c \rightarrow 1$, the bound of Proposition 3.1 approaches ρ_M exactly, and no boost, however extreme, can push the hypersurface centroid of a fixed physical body further from its rest-frame value. As u ranges over every inertial observer with the supplementary condition held fixed, $X^\mu(u, \Sigma)$ therefore stays confined to a disk of radius ρ_M orthogonal to $\mathbf{S}$ — the world tube identified by Møller and reproduced in Costa and Natário's modern treatment of spin supplementary conditions [18].

Two checks come with this bound for free, and both are run rather than assumed. The first is the conservation check of Section 2: P^μ and $J^{\mu\nu}$ are identical for every observer compared here, since Eq. (4) only ever redistributes $J^{\mu\nu}$ between its orbital and spin pieces at a fixed total. The second is the nonrelativistic limit: as $|\mathbf{v}|/c \rightarrow 0$ at any fixed, finite $\mathbf{S}$, $\Delta \mathbf{x} \rightarrow 0$ exactly, not merely to leading order, because Eq. (4) is linear and homogeneous in $\mathbf{v}$ with no additive term surviving at $\mathbf{v} = 0$. Every observer's centroid collapses onto $\mathbf{R}_N$ in that limit, recovering Eq. (1). The self-check program

accompanying this paper confirms both the saturation bound and this exact linearity numerically, on a large sample of random spin and velocity vectors together with the exact perpendicular case.

A single evaluated number fixes the scale without pretending to describe an experiment. For a spin- $\frac{1}{2}$ wavepacket of electron mass, $|\mathbf{S}| = \hbar/2$, so $\rho_M = \hbar/(2m_e c) \approx 1.931 \times 10^{-13} \text{ m}$ by direct substitution into Eq. (3) — half the electron’s reduced Compton wavelength, a coincidence Section 5 makes exact rather than numerical. Boosted transverse to its own spin at $\beta = v/c = 0.6$, Proposition 3.1 is saturated and

$$|\Delta \mathbf{x}| = \rho_M \beta \approx 1.158 \times 10^{-13} \text{ m}, \quad (5)$$

about a hundred and sixteen femtometres. This is an exact closed-form evaluation of Eqs. (4)–(3), checked numerically in the accompanying program, not a simulated or measured quantity; it is reported only as the characteristic scale the rest of this paper’s charts are measured against.

4 Two ways to close the same algebra

Hold the observer fixed now and ask what changes when the spin supplementary condition itself is replaced. Pryce’s 1948 classification answered this before quantum field theory had settled on the machinery used elsewhere in this series: the reasonable relativistic generalizations of $\mathbf{R}_N$ collapse, once spin is included, into two structurally distinct families [1]. The covariant centre of energy transforms as the spatial part of a four-vector under a change of Lorentz frame, exactly the property Eq. (2) was built with, but its three spatial components fail to commute with one another once spin is present, by a term Fleming’s manifestly covariant reformulation shows is generated entirely by spin-momentum coupling [4]. The Newton–Wigner operator, built by Newton and Wigner for exactly this purpose, has strictly commuting components,

$$[\mathbf{X}_{\text{NW}}^i, \mathbf{X}_{\text{NW}}^j] = 0, \quad (6)$$

the property an operator supporting ordinary probability densities and wavefunction localization needs, at the price of no longer transforming as a four-vector: a boost does not carry the Newton–Wigner operator of one frame into the Newton–Wigner operator of another by simple Lorentz transformation, because the construction is tied to a specific rest-frame decomposition of the Poincaré algebra’s boost generator [2]. Fleming’s contribution was to show this trade-off is forced on any spin-carrying operator built this way, a general obstruction rather than a defect of one particular construction:

commuting components and simple four-vector transformation cannot both be had from the same object [4].

Eq. (6) is not asserted here without recomputation. In its momentum-space representation, the Newton–Wigner position operator acts as

$$\mathbf{X}_{\text{NW}}^i \psi(\mathbf{P}) = i\hbar \frac{\partial \psi(\mathbf{P})}{\partial P_i}, \quad (7)$$

on the appropriately normalized positive-energy wavefunction, spin carried as a passive label [2]. Acting twice,

$$[\mathbf{X}_{\text{NW}}^i, \mathbf{X}_{\text{NW}}^j] \psi = -\hbar^2 \left(\frac{\partial^2 \psi}{\partial P_i \partial P_j} - \frac{\partial^2 \psi}{\partial P_j \partial P_i} \right),$$

and the bracket vanishes for any twice-continuously-differentiable ψ by Clairaut’s theorem on the equality of mixed partial derivatives. This paper does not take that theorem’s applicability here on faith: Appendix A works out the exact mixed partial of a representative smooth test function by hand in both orders of differentiation, confirms they agree analytically, and the accompanying self-check program then confirms a central finite-difference estimate of the same mixed partial converges to that exact value at the rate, close to a factor of four for each halving of the step size, that a correct second-order approximation to an exactly commuting pair of derivatives must have. That convergence, not an appeal to the theorem’s name alone, is what licenses Eq. (6) in this paper.

Spin itself supplies the other algebra this paper checks directly rather than quotes. For a spin- $\frac{1}{2}$ representation, $\mathbf{S} = (\hbar/2)\boldsymbol{\sigma}$ in terms of the Pauli matrices, and the rotation subalgebra any realization of angular momentum must satisfy is

$$[S^i, S^j] = i\hbar \epsilon^{ijk} S^k. \quad (8)$$

The self-check program verifies Eq. (8) exactly, to floating-point precision, by direct 2×2 complex matrix multiplication on all six ordered pairs (i, j) with $i \neq j$, rather than assuming the Pauli matrices satisfy the relation they are known to satisfy. This is the algebra responsible for the noncommutativity Fleming attributes to the covariant centroid: whatever spin-momentum term generates $[X_{\text{cov}}^i, X_{\text{cov}}^j] \neq 0$ is built from $\mathbf{S}$ itself, so at $S = 0$ every term carrying that structure vanishes identically, and the covariant centroid’s noncommutativity and the Newton–Wigner operator’s frame-dependence collapse together onto the same $\mathbf{R}_N$ this paper started from — the second exact zero-spin recovery this construction is required to pass, alongside the zero-boost recovery of Section 3.

It is worth being precise about what kind of object each chart in this pairing is. $X^\mu(u, \Sigma)$, built from a classical or semiclassical stress-energy distribution, is

a number: one spacetime point per physical configuration. $\mathbf{X}_{\text{NW}}$ is an operator on a Hilbert space; what an experiment actually reports is expectation values and, eventually, click statistics built from it, never the operator itself. The atlas this paper registers in Section 9 has to keep that distinction visible rather than silently comparing a number to an operator as though they answered the same question.

5 A trembling coordinate and the length it points to

The Newton–Wigner operator was built to repair a specific pathology in the coordinate $\hat{\mathbf{x}}$ that appears literally in the free Dirac equation. A free Dirac wavepacket generically mixes positive- and negative-energy plane-wave components, and $\hat{\mathbf{x}}$ ’s Heisenberg equation of motion consequently carries, on top of the ordinary uniform drift, an oscillating term with no classical counterpart, at angular frequency

$$\omega_Z = \frac{2Mc^2}{\hbar}, \quad (9)$$

set purely by the rest mass, the effect usually called *zitterbewegung*. For an electron, direct substitution of CODATA constants into Eq. (9) gives $\omega_Z \approx 1.553 \times 10^{21} \text{ rads}^{-1}$, an oscillation period of about $4.05 \times 10^{-21} \text{ s}$ in the particle’s own rest frame, values the accompanying self-check program reproduces directly from Eq. (9). Foldy and Wouthuysen’s canonical transformation is what removes the oscillating term, by transforming to the representation in which the position operator is exactly $\mathbf{X}_{\text{NW}}$ and the term is absorbed into the difference between $\hat{\mathbf{x}}$ and $\mathbf{X}_{\text{NW}}$ [3].

The amplitude of that oscillation sets a length, and the length is not independent of the one this paper has already fixed. The reduced Compton wavelength is $\bar{\lambda}_C = \hbar/(Mc)$, and for a spin- $\frac{1}{2}$ particle,

$$\rho_M = \frac{\hbar/2}{Mc} = \frac{\bar{\lambda}_C}{2}, \quad (10)$$

an exact algebraic identity, not a numerical coincidence: both sides of Eq. (3) and the definition of $\bar{\lambda}_C$ share the same $\hbar/(Mc)$ factor once $|\mathbf{S}| = \hbar/2$ is substituted, and the self-check program confirms the ratio equals $\frac{1}{2}$ to machine precision for the electron’s CODATA mass. The Møller disk of Section 3 and the *zitterbewegung* amplitude of this section are, up to the factor of two fixed by Eq. (10), the same length scale viewed from a classical spin-transport identity and from a single-particle relativistic wave equation respectively.

Zitterbewegung is not directly observable in a free electron — the relevant scales sit far outside what a free-particle experiment resolves — but its structure has been observed exactly where it can be engineered:

in a single trapped ion whose internal and motional degrees of freedom are driven to reproduce the one-dimensional Dirac equation term for term, with the simulated position executing the predicted oscillation and the predicted crossover as the simulated mass parameter is tuned [20]. That is an observed fact about an engineered quantum simulation of the Dirac equation, not a measurement of a free electron’s own coordinate, and this paper keeps that distinction explicit rather than letting it blur; nothing in this section is computed or simulated here, only cited.

What *zitterbewegung* adds to this paper’s ledger of exact limits is a third, independent confirmation of non-relativistic recovery. $\hat{\mathbf{x}}$ and $\mathbf{X}_{\text{NW}}$ disagree by a time-dependent term whose amplitude is set by Eq. (10) and whose weight in a wavepacket’s dynamics is controlled by the fraction of negative-energy content the packet carries. As $v \ll c$ and that fraction is driven to zero, the oscillating term’s contribution falls relative to the packet’s own spatial width, and $\hat{\mathbf{x}}$, $\mathbf{X}_{\text{NW}}$, and $\mathbf{R}_N$ converge onto a single description — the same collapse already exhibited for the observer-transition shift of Section 3 and the supplementary-condition trade-off of Section 4, reached here from an entirely different starting operator and confirming, rather than merely repeating, the same nonrelativistic limit.

6 Where an analytic continuation refuses to sit still

Every chart built so far shares an assumption none of them states out loud: that "the particle is exactly here, right now" is at least momentarily a sharp, well-posed statement, even though which "here" depends on which chart answers. Hegerfeldt’s theorem removes that assumption for every chart in Sections 3–5 at once, and it is worth stating with some care exactly what the theorem forces and what it leaves alone, since the two are easy to blur.

The mechanism is an analyticity argument, not a positivity coincidence. If a system evolves under a Hamiltonian whose energy spectrum is bounded below, the vector $\psi(t) = e^{-iHt/\hbar}\psi(0)$ extends, as a Hilbert-space-valued function, to a function holomorphic in a half-plane of complex t ; matrix elements built from $\psi(t)$ are then boundary values of holomorphic functions and are consequently real-analytic in t along the real axis. A real-analytic function that equals a fixed value throughout any open interval, however short, is fixed at that same value everywhere its analytic continuation reaches — in particular, for a free particle, for all real t . Hegerfeldt used this to show that a wavefunction exactly zero outside a bounded spatial region at one instant, evolving freely with positive-energy dynamics, develops a strictly positive probability of being found

arbitrarily far outside that region after any elapsed time greater than zero, however small [6]; a state held exactly inside the region for any stretch of time, rather than only at one instant, would by the argument above have had to be found there always, including arbitrarily far in the past, which is not the generic behavior of a genuinely evolving free state. Hegerfeldt later weakened the hypothesis from exact confinement to merely exponentially bounded spatial tails, reaching the same instantaneous-spreading conclusion [7]. Perez and Wilde reached an equivalent conclusion by a different, more axiomatic route, showing causality, translation covariance, and a positive-energy spectrum are jointly incompatible with strict localization in any bounded spacetime region at all [8], and Malament later formalized a stronger version of the same incompatibility directly from locality and covariance axioms, without reference to a specific free-particle dispersion relation, arguing that no genuinely localizable relativistic particle theory can be built at all [9]; Halvorson and Clifton examined and defended that argument against several proposed loopholes [10].

Three things this result does not say deserve equal emphasis. First, the argument never uses the position operator's commutator structure, only positivity of the energy and the existence of an exactly zero exterior amplitude; it applies to $\mathbf{X}_{\text{NW}}$ exactly as it applies to $\hat{\mathbf{x}}$ or to the covariant centroid of Section 4, and commuting components buy no protection against it. Second, it is not by itself a statement that usable superluminal signals propagate in a full interacting field theory; it is a statement about a single-particle wavefunction's amplitude under free time evolution, and the stronger causality property actually required of a field theory — that observables localized at spacelike separation commute — is a separate property of the field algebra, not something this argument settles either way. Third, and most directly relevant to Sections 7–8, it says nothing at all about an object that never claims exact confinement in the first place; a probability measure over detector outcomes is simply not the kind of object the argument constrains.

A simpler, purely kinematic fact shares the same analyticity mechanism and is fully elementary to state and check, without being Hegerfeldt's own dynamical theorem: a nonzero wavefunction cannot be compactly supported in both position and momentum. If $\psi(x, 0)$ vanishes identically outside a bounded interval, its Fourier transform $\hat{\psi}(p, 0)$ extends to an entire function of a complexified momentum variable, by the Paley–Wiener representation of a compactly supported function's transform as a convergent power series in that variable; a nonzero entire function cannot vanish on any open interval of the real line without vanishing identically, by the identity theorem for analytic functions. So $\hat{\psi}(p, 0)$ cannot itself vanish on any open momentum interval,

and a state exactly confined in position at $t = 0$ already has momentum-space support spread over the entire real line before any time evolution is applied at all. This static fact is not offered here as a substitute for Hegerfeldt's dynamical result — it says nothing about what happens as t increases, only that the momentum-space ingredient the dynamical argument needs is already present at $t = 0$ — and it is stated and used only at this level, not extended into a re-derivation of the theorem itself.

Quantum field theory sharpens the same conclusion in language that never mentions a wavepacket. The Reeh–Schlieder theorem states that the vacuum is cyclic for the algebra of operators localized in any bounded spacetime region, however small: acting on the vacuum with operators supported only in that region can approximate any state in the full Hilbert space to arbitrary precision, so no nontrivial local operator annihilates the vacuum and no state is exactly and only present in one bounded region with nothing correlated to it anywhere else [12]. Sharp localization is foreclosed here at the level of the algebra the charts of Sections 3–5 are ultimately built out of, not merely difficult to arrange with a cleverer choice of wavefunction.

This is the atlas's genuine boundary, not a gap awaiting a better construction. Every chart from $\mathbf{R}_N$ through $\mathbf{X}_{\text{NW}}$ is a legitimate, checkable answer to "where, approximately, and by which convention," and each stops being exactly true, in the strict sharp-operator sense, the moment it is pushed to describe a state confined with certainty to a bounded region while evolving under causal free-particle dynamics. The next two charts exist because of that limit, not despite it.

7 A screen that never promised a point

The response to Hegerfeldt's obstruction that has actually held up is to stop demanding a sharp position operator at all. Werner's screen observables replace the eigenvalue-and-projector structure of $\hat{\mathbf{x}}$ or $\mathbf{X}_{\text{NW}}$ with a positive-operator-valued measure, a family of positive operators assigned to regions of a detection screen, summing to the identity, each reporting the probability of "a click registered in this region" without asserting the particle possessed a definite position immediately beforehand [11]. A POVM chart never claims exact confinement to begin with, so Section 6's argument has no state to act on that it can call "exactly here": the resolution of the measurement is set by the physical smearing of the detection scheme, not driven to zero by fiat.

This chart is not held in reserve for elegance; for one important class of system it is the only chart on offer. Wightman's 1962 analysis of localizability by ir-

reducible representation of the Poincaré group established that a genuine, covariant, Newton–Wigner-type position operator exists for every massive representation of any spin, but fails to exist for massless representations of helicity magnitude one or greater [5]. The photon is exactly such a representation, with no rest frame in which to define a Newtonian centroid and no covariant operator of the kind Section 4 discussed. Sipe’s photon wave function formalism and Peres and Terno’s review of relativistic quantum information both work through the operational consequences of this gap directly rather than around it [13, 14]: what remains for a photon is a detection screen, a POVM, and a click, not an operator eigenvalue.

A single explicit construction demonstrates that a POVM is not merely a name for "whatever avoids the theorem," but an object whose defining properties, positivity and normalization, can be exhibited and checked. Take, on a system for which a self-adjoint coordinate $\hat{X}$ is well defined, a family of Gaussian-smeared elements

$$E(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\hat{X} - x)^2}{2\sigma^2}\right), \quad (11)$$

parametrized by a fixed resolution width $\sigma > 0$ representing the physical blur of the detection scheme, the floor no shrinking of the apparatus removes. $E(x)$ is a smooth, non-negative real function of the self-adjoint operator $\hat{X} - x\mathbb{I}$; by the spectral theorem, a non-negative real function of a self-adjoint operator is a positive operator, so $E(x) \geq 0$ for every x without further argument. Resolution of the identity follows from the same spectral representation: on any eigenvector of $\hat{X}$ with eigenvalue x_0 , $E(x)$ acts as multiplication by the ordinary Gaussian density $\exp(-(x_0 - x)^2/2\sigma^2)/(\sqrt{2\pi}\sigma)$, and

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x_0 - x)^2}{2\sigma^2}\right) dx = 1 \quad (12)$$

exactly, for every real x_0 and every $\sigma > 0$, by the standard Gaussian integral. Since $\int dx E(x)$ therefore acts as the identity on every eigenvector of $\hat{X}$, it equals the identity operator. The accompanying self-check program confirms Eq. (12) by quadrature for several (x_0, σ) pairs spanning four orders of magnitude in both parameters, to a residual below 10^{-13} in every case, and confirms the kernel never samples negative on a dense grid, both properties Rule 3 of this paper’s construction protocol requires be checked rather than assumed for any POVM it writes down. Taking $\sigma \rightarrow 0$ formally recovers the sharp projector-valued measure of Section 6, the limit in which the theorem’s obstruction reappears; an isotropic three-dimensional version follows from the identical one-dimensional construction applied independently to each Cartesian direction, since the resulting product kernel is again non-negative and integrates to

one by Fubini’s theorem applied to Eq. (12) three times over.

Eq. (11) is offered as one explicit, checkable instance satisfying the defining properties of a localization POVM, not as the unique or canonical detector response; Bliokh, Rodríguez-Fortuño, Nori, and Zayats’s review of spin–orbit interactions of light surveys the range of operational localization schemes actually used for photons, none of which is this paper’s business to adjudicate [17]. What matters for the atlas registered in Section 9 is only that this chart is of a different mathematical type from every chart in Sections 3–5: a measure-valued object whose output on a given trial is a probability distribution over an actual instrument’s actual geometry, not a number or an operator eigenvalue.

8 An estimator that files no report with the disk

There is a fifth kind of answer to "where," and its relationship to the four charts above has to be stated precisely before it is used. Prepare a system in $|\psi\rangle$, couple it weakly to a pointer, and postselect on a final state $|f\rangle$ measured after the coupling. The weak value of an observable $\hat{A}$ conditioned on that postselection,

$$A_w = \frac{\langle f|\hat{A}|\psi\rangle}{\langle f|\psi\rangle}, \quad (13)$$

is the quantity a weakly coupled pointer’s mean shift actually reports, to leading order in the coupling strength [15]. A_w is not an eigenvalue of $\hat{A}$, not an ordinary expectation value, and not guaranteed to be real; it is an amplitude ratio over a sub-ensemble selected by an event that has not yet occurred at the moment the weak coupling acts.

A closed-form family, worked out in full in Appendix A, shows exactly how far this departs from an operator-eigenvalue answer. Preselect a spin- $\frac{1}{2}$ system along x , $|\psi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)$, and postselect on $|f(\theta)\rangle = \cos(\theta/2)|\uparrow\rangle + \sin(\theta/2)|\downarrow\rangle$, a real superposition at angle θ . The weak value of σ_z works out, by the tangent-subtraction identity, to

$$\langle\sigma_z\rangle_w(\theta) = \tan\left(\frac{\pi}{4} - \frac{\theta}{2}\right). \quad (14)$$

The accompanying self-check program confirms Eq. (14) against a direct bra-ket computation built independently from the same two states, to a residual below 10^{-9} at every sampled angle. σ_z has eigenvalues ± 1 only, yet Eq. (14) exceeds 1 in magnitude for $\theta < 0$ and diverges as $\theta \rightarrow -\pi/2$, exactly where $|f(\theta)\rangle$ approaches orthogonality to $|\psi\rangle$ and the denominator of Eq. (13) vanishes: the self-check program confirms the

magnitude grows from about 96 at $\theta = -1.55$ to over 2×10^4 at $\theta = -1.5707$, a monotone divergence toward the singular angle, not a bounded quantity that merely happens to look large. Allowing a complex postselection phase, $|f(\alpha)\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + e^{i\alpha}|\downarrow\rangle)$, produces weak values off the real axis entirely; at $\alpha = \pi/2$, the postselection Aharonov, Albert, and Vaidman originally used, the construction gives $\langle\sigma_z\rangle_w = i$, purely imaginary, confirmed directly by the self-check program's independent bra-ket evaluation [15]. Neither divergence nor a departure from the real line is an error in the estimator; both are exact, closed-form consequences of Eq. (13) applied to a postselection chosen close to or exactly at orthogonality.

Applied to position, $x_w = \langle f|\hat{x}|\psi\rangle/\langle f|\psi\rangle$ is this paper's fifth chart, and it answers a different question from every chart in Sections 3–7: not what a position operator returns or what a detector reports, but what a weakly coupled pointer's mean reading becomes once conditioned on an event that filters the ensemble after the fact. Because postselection can make $\langle f|\psi\rangle$ arbitrarily small, x_w can be pushed arbitrarily far from anything the Møller disk of Section 3 constrains; it is not bounded by ρ_M , and it does not need to be, because it was never built from the observer-transition protocol that bound was derived for. An anomalous weak value lying outside that disk is not a contradiction of Proposition 3.1; treating the two as commensurable would be exactly the mismatched-protocol error the kill criteria of Section 9 are built to catch. For that reason this paper computes no disagreement number between x_w and $\mathbf{X}_{\text{NW}}$ or $X^\mu(u, \Sigma)$: the comparison is registered as undefined, not evaluated and reported as large.

The one place any of this touches an actually measured number is optical, and it arrives through this chart rather than through an operator eigenvalue. Light crossing an interface between media of different refractive index acquires a transverse, polarization-dependent displacement, the spin Hall effect of light, far smaller than the beam's own wavelength under direct detection. Hosten and Kwiat used exactly the weak-value amplification structure above, preselecting polarization and postselecting on a near-orthogonal state, to enhance the effective displacement by close to four orders of magnitude and resolve a shift at the angstrom scale [16]. What was measured is a pointer's mean shift, the quantity Eq. (13) defines, not a photon position eigenvalue — consistent with Section 7's point that no covariant position operator exists for a photon to have an eigenvalue of in the first place. Harte and Oancea showed the equations governing that spin-dependent trajectory in a medium of varying refractive index are a special case of the Mathisson–Papapetrou equations governing a classical spinning body's centroid motion, the same equations behind the Møller disk of Section 3 [21]. The spin Hall effect of light is, in this sense,

a Møller-type centroid shift for a massless spinning ray, measured in a regime with no covariant Newton–Wigner-type operator to underlie it, extracted through the one chart, the postselected weak value, that does not require one. This measurement is cited as an existing result, not reproduced, simulated, or extended here.

9 A bound proved, a grid registered, not yet run

Three charts built in this paper, $\mathbf{R}_N$, $X^\mu(u, \Sigma)$, and $\mathbf{X}_{\text{NW}}$, form a genuine, well-behaved sub-atlas: each is a candidate answer to the same question, "where is the centroid," evaluated under a stated protocol, and each converges to the same point in the nonrelativistic, zero-spin limit. Define, for any two of them evaluated on the same hypersurface Σ ,

$$\delta_{ij}(u, \Sigma) = \frac{|X_i^\mu(u, \Sigma) - X_j^\mu(u, \Sigma)|_\Sigma}{\rho_M}, \quad (15)$$

the proper spatial separation within Σ between the two charts' answers, normalized by the spin-to-mass length fixed in Section 2.

Proposition 9.1 (Properties of δ_{ij} on the well-behaved sub-atlas). *δ_{ij} is non-negative and symmetric in i, j by the definition of a norm; it equals zero if and only if the two charts return the identical point on Σ ; for the observer-transition comparison between the rest-frame centroid and $X^\mu(u, \Sigma)$ at fixed supplementary condition, it is bounded above by exactly 1, since Proposition 3.1 showed $|\Delta \mathbf{x}| \leq \rho_M$ saturates rather than merely approaches that value; and it vanishes in both directions this construction is required to pass, as $v/c \rightarrow 0$ for any spin (Section 3) and as $S \rightarrow 0$ for any observer or supplementary condition (Section 4).*

Non-negativity and symmetry are immediate from Eq. (15), since $|\cdot|_\Sigma$ is a norm on the tangent space of Σ . The zero-iff-coincidence property, the boundedness, and the two limiting recoveries are exactly the results already proved in Sections 3 and 4, restated here as properties of δ_{ij} rather than re-derived; Proposition 9.1 packages existing results into the form the atlas below is required to check, and adds nothing beyond that packaging.

δ_{ij} is not extended to the fourth and fifth charts of Sections 7 and 8. The weak value x_w answers a postselection-conditioned question with no protocol-matched notion of "the same hypersurface" to compare against, and Section 8 already registered that comparison as undefined rather than small; the POVM chart of Section 7 returns a probability distribution, not a point, and $|\cdot|_\Sigma$ in Eq. (15) has no argument to act

on. Extending a single functional across all five charts, rather than stopping at the three that share a common type, would itself be the kind of protocol mismatch this paper is built to flag rather than commit.

Definition 9.1 (Kill criteria). δ_{ij} is rejected as a physically meaningful chart-disagreement functional if, once implemented over an actual grid of boosts, spins, and hypersurface tilts, either of the following holds: its value fails to vanish in the agreed Newtonian limit, $v/c \rightarrow 0$ and $S \rightarrow 0$ jointly or separately, where Proposition 9.1 requires exact vanishing rather than a small residual; or it is applied to compare quantities drawn from incompatible protocols, such as an operator expectation value against a classical centroid, or a postselected estimator against either, as though the comparison were defined.

The first clause is a numerical-implementation check on an analytic result already proved above: a correct implementation of $X^\mu(u, \Sigma)$ and $\mathbf{X}_{\text{NW}}$ on an actual discretized wavepacket must reproduce Proposition 9.1's exact zero to within a stated quadrature tolerance, and a failure to do so indicates an implementation error, not a new physical effect. The second clause is the discipline Sections 7–8 exercised explicitly by declining to compute a number: applying δ_{ij} -style reasoning across a type mismatch would manufacture an apparent disagreement out of a category error rather than a trajectory or state difference.

The numerical program this construction is built toward, and which this paper registers without running, compares $\mathbf{R}_N$, $X^\mu(u, \Sigma)$, and $\mathbf{X}_{\text{NW}}$ over a grid of boost velocity, spin magnitude, hypersurface tilt relative to the boost, and wavepacket width, for an explicit boosted Gaussian wavepacket carrying spin, together with one worked numerical example applying the same bookkeeping to a spin-Hall-of-light beam centroid built from an implemented weak-value measurement model rather than only the closed-form estimator of Section 8. Every saved grid point would need to record the protocol it corresponds to in full, a quadrature error estimate, and a reparametrization-invariance residual obtained by recomputing the same point under an alternative hypersurface parametrization; no such grid point exists yet, and Section 10 states why. A later curved-spacetime extension of this bookkeeping is not attempted here either; Janson and Lopp's recent extension of hypersurface-centroid bookkeeping to a composite quantum particle's internal structure on a Schwarzschild background is the specific published result any such extension would first have to reproduce before a new curved-spacetime claim could be built on it [19].

Two readings of a completed atlas would leave δ_{ij} without independent content, and each has a decisive comparison rather than a rhetorical dismissal. An *apparatus-dominance* reading holds that a nonzero grid

value is shaped more strongly by the choice of hypersurface parametrization or wavepacket discretization than by the physical boost or spin; it is checked by the reparametrization residual named above, required to be consistent with zero within quadrature error at every grid point. A *redundancy* reading holds that δ_{ij} , once computed, adds nothing beyond Pryce's own 1948 classification of which constructions coincide and which do not; it is checked by confirming the atlas exposes a genuine tradeoff, such as a specific combination of boost and spin at which $X^\mu(u, \Sigma)$ and $\mathbf{X}_{\text{NW}}$ disagree by an amount not already visible from the exact bound of Proposition 3.1 alone. Neither reading is addressed here; both are registered as the standard a completed atlas would have to meet.

10 What keeps this atlas from being read

Everything computed in this draft is an exact identity, a closed-form bound, or a convergent numerical check of a closed-form claim. Nothing in it is the atlas Section 9 registers, and this section names, one at a time, what stands between the two and which sentence already written above each gap keeps from being true.

The first gap is a boosted, spinful wavepacket implementation. Section 9 states that "the numerical program this construction is built toward... compares $\mathbf{R}_N$, $X^\mu(u, \Sigma)$, and $\mathbf{X}_{\text{NW}}$ over a grid of boost velocity, spin magnitude, hypersurface tilt... and wavepacket width, for an explicit boosted Gaussian wavepacket carrying spin"; no such wavepacket has been built, discretized, or evaluated on any hypersurface. Until it exists, that sentence describes a program, not a result, and no numerical value of δ_{ij} from Eq. (15) can be reported for any boost other than the exact analytic limits Sections 3 and 4 already cover in closed form.

The second gap is the spin-Hall beam-centroid worked example named in the same paragraph of Section 9, "one worked numerical example applying the same bookkeeping to a spin-Hall-of-light beam centroid built from an implemented weak-value measurement model rather than only the closed-form estimator of Section 8." Section 8 cites Hosten and Kwiat's measured shift as an existing result and explicitly declines to reproduce or extend it; no measurement model connecting that citation to this paper's own bookkeeping has been implemented, so the dossier's own stated numerical target for this entry remains a description of what would be built, not a value this paper supplies.

The third gap is the curved-spacetime extension. Section 9 names Janson and Lopp's published hypersurface-centroid bookkeeping on a Schwarzschild background as "the specific published result any such extension would first have to reproduce before a new

curved-spacetime claim could be built on it"; that reproduction has not been attempted, so every claim in this paper stays confined to flat spacetime, as Section 1 states at the outset, and no sentence anywhere in this draft should be read as extending past that boundary.

The fourth gap is the quadrature and reparametrization error budget Definition 9.1 requires before its first clause can be tested at all: "a correct implementation of $X^\mu(u, \Sigma)$ and $\mathbf{X}_{\text{NW}}$ on an actual discretized wavepacket must reproduce Proposition 9.1's exact zero to within a stated quadrature tolerance." No discretized implementation and no such tolerance have been produced, so the kill criteria of Definition 9.1 are registered as a standard to meet, not yet applied to anything.

The fifth gap is the redundancy comparison Section 9 poses against Pryce's own 1948 classification: confirming that a computed atlas "exposes a genuine tradeoff... not already visible from the exact bound of Proposition 3.1 alone." Without a computed grid to inspect, that comparison cannot be run, and the honest position is that this paper has not yet shown δ_{ij} does independent work beyond restating the exact bounds Sections 3 and 4 already established analytically.

The sixth gap is a named external reviewer with expertise in relativistic localization theory and spinning-body dynamics, who has not yet been approached. Concretely, before any numerical value from this construction is reported as a finding rather than a registered protocol, the following must all be true: a validated wavepacket implementation of the three-chart sub-atlas; an implemented weak-value measurement model connecting this paper's bookkeeping to the Hosten–Kwiat measurement; a curved-spacetime extension checked against Janson and Lopp's published result; a computed quadrature and reparametrization error budget; a completed redundancy comparison against Pryce's classification; and an external review of all of the above. As of this draft, all six remain open, and any one of them being incomplete is sufficient reason to withhold a numerical claim.

11 A radius, an algebra, and a proven boundary

This paper delivers three checked results and one registered protocol built on top of them. It derives, from the elementary bound on a cross product rather than by assumption, that the observer-transition shift between a body's rest-frame centroid and its hypersurface energy centroid saturates exactly at the spin-to-mass length ρ_M when spin and boost are transverse, and confirms the same bound and its exact zero-boost limit numerically. It recomputes, rather than quotes, the commuting components of the Newton–Wigner momentum-space representation from Clairaut's theorem on mixed

partial derivatives, and the spin rotation subalgebra from direct matrix multiplication on a spin- $\frac{1}{2}$ representation, tying the same length scale to the zitterbewegung frequency of the naive Dirac coordinate along the way. It states precisely, rather than loosely, what the Hegerfeldt and Reeh–Schlieder theorems and Malament's axiomatic argument forbid — a sharp, causally propagating, exactly localized position operator, for any construction built by the methods of Sections 3–5 — and what they leave alone: an operational, detector-tied probability measure, exhibited here as an explicit Gaussian POVM checked for positivity and normalization by spectral construction and by quadrature, and a postselected weak-value estimator, exhibited here as a closed-form family that diverges exactly as its postselection approaches orthogonality. It defines a normalized disagreement functional over the three charts that share a common type, proves it non-negative, symmetric, exactly zero on coincidence, bounded by one on the observer-transition comparison, and exactly vanishing in the nonrelativistic and zero-spin limits.

None of this is the atlas. No boosted, spinful wavepacket has been discretized and evaluated; no spin-Hall-of-light beam centroid has been computed from an implemented weak-value model; no curved-spacetime extension has been checked against Janson and Lopp's published result; no quadrature or reparametrization error budget exists; no redundancy comparison against Pryce's own classification has been run; and no external reviewer has examined any of it. Section 10 states this without softening it, because a construction that reports exactly 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 survives the kill criteria of Definition 9.1, the normalized disagreement functional would offer a single bounded number, with declared units, a proved zero-limit, and a stated failure mode, in place of a qualitative list of charts that happen to disagree. If it fails — if a computed grid turns out to be dominated by discretization artifacts rather than by physical boost or spin, or to add no classification beyond Pryce's original one — the results already proved in Sections 3 and 4 are untouched, because they were derived and checked independently of any atlas ever being run. The plurality of relativistic position was never a defect waiting on a cleverer single operator to resolve it; it is a small number of exactly comparable charts, a length scale that bounds one pair of them exactly, an algebra recomputed rather than assumed, and a proven boundary past which the word "position" stops referring to a sharp point at all — and this draft's contribution stops at registering that structure precisely enough for the atlas, whenever it is built, to be judged against it.

A Full derivations: the commuting chart and the estimator's divergence

This appendix carries out, in full, the two calculations stated in Sections 4 and 8.

A.1 The mixed partial derivative behind Eq. (6)

Take the representative test function used by the accompanying self-check program,

$$f(p_x, p_y, p_z) = \sin(p_x p_y) + \cos(p_y p_z) + e^{-(p_x^2 + p_z^2)/4} + 0.3 p_x p_z^2,$$

chosen with no special symmetry between its three arguments. Differentiating first in p_x ,

$$\frac{\partial f}{\partial p_x} = p_y \cos(p_x p_y) - \frac{p_x}{2} e^{-(p_x^2 + p_z^2)/4} + 0.3 p_z^2,$$

and then in p_y ,

$$\frac{\partial^2 f}{\partial p_y \partial p_x} = \cos(p_x p_y) - p_x p_y \sin(p_x p_y),$$

since the second and third terms of $\partial f / \partial p_x$ carry no p_y dependence. Differentiating in the opposite order, first in p_y ,

$$\frac{\partial f}{\partial p_y} = p_x \cos(p_x p_y) - p_z \sin(p_y p_z),$$

and then in p_x ,

$$\frac{\partial^2 f}{\partial p_x \partial p_y} = \cos(p_x p_y) - p_x p_y \sin(p_x p_y),$$

the identical expression. At the test point $(p_x, p_y, p_z) = (0.7, -1.3, 0.4)$ used by the self-check program, $p_x p_y = -0.91$, and both expressions evaluate to $\cos(-0.91) - (-0.91) \sin(-0.91) \approx -0.104703$, matching the value the self-check program reports before any finite-difference estimate is compared against it. This is Clairaut's theorem confirmed by direct hand calculation for this specific function, not asserted from its name; the self-check program's central finite-difference estimate of the same mixed partial then converges to -0.104703 at the $O(h^2)$ rate reported in Section 4, closing the loop between the analytic argument and the numerical check.

A.2 The postselected weak-value family, Eq. (14)

Preselect $|\psi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)$ and postselect $|f(\theta)\rangle = \cos(\theta/2)|\uparrow\rangle + \sin(\theta/2)|\downarrow\rangle$, both real, so $\langle f(\theta)|$ has the same components. Then

$$\langle f|\psi\rangle = \frac{1}{\sqrt{2}} \left(\cos \frac{\theta}{2} + \sin \frac{\theta}{2} \right),$$

and, since $\sigma_z|\psi\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle - |\downarrow\rangle)$,

$$\langle f|\sigma_z|\psi\rangle = \frac{1}{\sqrt{2}} \left(\cos \frac{\theta}{2} - \sin \frac{\theta}{2} \right).$$

Dividing,

$$\langle \sigma_z \rangle_w(\theta) = \frac{\cos(\theta/2) - \sin(\theta/2)}{\cos(\theta/2) + \sin(\theta/2)} = \frac{1 - \tan(\theta/2)}{1 + \tan(\theta/2)},$$

after dividing numerator and denominator by $\cos(\theta/2)$. The tangent-subtraction identity, $\tan(A - B) = (\tan A - \tan B)/(1 + \tan A \tan B)$ with $A = \pi/4$ and $\tan(\pi/4) = 1$, gives exactly this expression, so

$$\langle \sigma_z \rangle_w(\theta) = \tan \left(\frac{\pi}{4} - \frac{\theta}{2} \right),$$

Eq. (14). The denominator $\cos(\theta/2) + \sin(\theta/2)$ vanishes when $\tan(\theta/2) = -1$, i.e. $\theta = -\pi/2 \pmod{2\pi}$, exactly where $|f(\theta)\rangle$ becomes orthogonal to $|\psi\rangle$; the closed form's divergence at that angle is therefore the same orthogonality that makes $\langle f|\psi\rangle \rightarrow 0$ in Eq. (13), not a separate singularity introduced by the trigonometric simplification.

Allowing a complex postselection phase, $|f(\alpha)\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + e^{i\alpha}|\downarrow\rangle)$ with bra $\langle f(\alpha)| = \frac{1}{\sqrt{2}}(\langle\uparrow| + e^{-i\alpha}\langle\downarrow|)$,

$$\langle f|\psi\rangle = \frac{1 + e^{-i\alpha}}{2}, \quad \langle f|\sigma_z|\psi\rangle = \frac{1 - e^{-i\alpha}}{2},$$

so

$$\langle \sigma_z \rangle_w(\alpha) = \frac{1 - e^{-i\alpha}}{1 + e^{-i\alpha}}.$$

Multiplying numerator and denominator by $e^{i\alpha/2}$,

$$\langle \sigma_z \rangle_w(\alpha) = \frac{e^{i\alpha/2} - e^{-i\alpha/2}}{e^{i\alpha/2} + e^{-i\alpha/2}} = \frac{2i \sin(\alpha/2)}{2 \cos(\alpha/2)} = i \tan \left(\frac{\alpha}{2} \right),$$

purely imaginary for every real $\alpha \neq \pi$, and reducing to $i \tan(\pi/4) = i$ at $\alpha = \pi/2$, the postselection Aharonov, Albert, and Vaidman originally used [15]. The accompanying self-check program confirms this closed form against an independent bra-ket evaluation at five values of α , with residuals at the level of floating-point roundoff.

B What the self-check program actually verifies

The scaffold program `check_position_atlas.py` performs no wavepacket simulation, no implemented weak-value measurement model, and no curved-spacetime evaluation. In its default `-self-check` mode it runs the following exact and closed-form analytic checks, each corresponding to a claim made in the sections above:

1. the exact analytic mixed partial of a representative smooth momentum-space function, worked out by hand in Appendix A.1, compared against a central finite-difference estimate at six shrinking step sizes, confirming $O(h^2)$ convergence to the exact value and licensing $[\mathbf{X}_{\text{NW}}^i, \mathbf{X}_{\text{NW}}^j] = 0$ (Eq. (6));
2. the spin rotation subalgebra, $[S^i, S^j] = i\hbar\epsilon^{ijk}S^k$ (Eq. (8)), verified exactly by direct 2×2 complex matrix multiplication on a spin- $\frac{1}{2}$ representation for all six ordered index pairs;
3. the cross-product bound behind Proposition 3.1, checked on twenty thousand random spin and velocity vectors plus the exact perpendicular case, together with the electron-scale numerical evaluations of ρ_M and the observer-transition shift at $\beta = 0.6$ quoted in Eq. (5);
4. the exact algebraic identity $\rho_M = \bar{\lambda}_C/2$ (Eq. (10)) and the zitterbewegung frequency and period of Eq. (9), evaluated with CODATA physical constants for the electron;
5. positivity, by pointwise non-negativity of the kernel, and normalization, by quadrature to a residual below 10^{-13} , of the Gaussian localization POVM of Eq. (11), at four widely separated (x_0, σ) pairs;
6. the closed-form weak-value family of Eq. (14), checked against an independent bra-ket evaluation at six real postselection angles, its divergence toward the orthogonal postselection at $\theta = -\pi/2$, and the complex-postselection closed form $i \tan(\alpha/2)$ derived in Appendix A.2, checked against an independent bra-ket evaluation at five values of α ;
7. the exact nonrelativistic and zero-spin recoveries of Sections 3 and 4: $\Delta \mathbf{x}$ exactly linear in β with exact zero at $\beta = 0$, and ρ_M exactly zero at $S = 0$.

The saved output, `scaffold_validation_output.txt`, is the program's real captured standard output from a `-self-check` run, and records `"scientific_result": false` and `"atlas_computed": false`; no field in it reports a value for δ_{ij} over any nontrivial grid point. The program's `-simulate` option exits with status 2 and prints the five-item list of unimplemented atlas components named in Section 10; it never proceeds to produce a number. A later implementation must not reuse the name `check_position_atlas.py` while silently relaxing these seven checks; it must add the wavepacket, beam-centroid, 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] M. H. L. Pryce (1948). The mass-centre in the restricted theory of relativity and its connexion with the quantum theory of elementary particles. *Proceedings of the Royal Society of London A* 195, 62–81. <https://doi.org/10.1098/rspa.1948.0103>.
- [2] T. D. Newton and E. P. Wigner (1949). Localized states for elementary systems. *Reviews of Modern Physics* 21, 400–406. <https://doi.org/10.1103/RevModPhys.21.400>.
- [3] L. L. Foldy and S. A. Wouthuysen (1950). On the Dirac theory of spin 1/2 particles and its non-relativistic limit. *Physical Review* 78, 29–36. <https://doi.org/10.1103/PhysRev.78.29>.
- [4] G. N. Fleming (1965). Covariant position operators, spin, and locality. *Physical Review* 137, B188–B197. <https://doi.org/10.1103/PhysRev.137.B188>.
- [5] A. S. Wightman (1962). On the localizability of quantum mechanical systems. *Reviews of Modern Physics* 34, 845–872. <https://doi.org/10.1103/RevModPhys.34.845>.
- [6] G. C. Hegerfeldt (1974). Remark on causality and particle localization. *Physical Review D* 10, 3320–3321. <https://doi.org/10.1103/PhysRevD.10.3320>.
- [7] G. C. Hegerfeldt (1985). Violation of causality in relativistic quantum theory? *Physical Review Letters* 54, 2395–2398. <https://doi.org/10.1103/PhysRevLett.54.2395>.
- [8] J. F. Perez and I. F. Wilde (1977). Localization and causality in relativistic quantum mechanics. *Physical Review D* 16, 315–317. <https://doi.org/10.1103/PhysRevD.16.315>.
- [9] D. B. Malament (1996). In defense of dogma: why there cannot be a relativistic quantum mechanics of (localizable) particles. In R. Clifton (ed.), *Perspectives on Quantum Reality*, 1–10. Springer Netherlands. https://doi.org/10.1007/978-94-015-8656-6_1.
- [10] H. Halvorson and R. Clifton (2002). No place for particles in relativistic quantum theories? *Philosophy of Science* 69, 1–28. <https://doi.org/10.1086/338939>.
- [11] R. F. Werner (1986). Screen observables in relativistic and nonrelativistic quantum mechanics. *Journal of Mathematical Physics* 27, 793–803. <https://doi.org/10.1063/1.527184>.

- [12] E. Witten (2018). APS Medal for Exceptional Achievement in Research: Invited article on entanglement properties of quantum field theory. *Reviews of Modern Physics* 90, 045003. <https://doi.org/10.1103/RevModPhys.90.045003>.
- [13] A. Peres and D. R. Terno (2004). Quantum information and relativity theory. *Reviews of Modern Physics* 76, 93–123. <https://doi.org/10.1103/RevModPhys.76.93>.
- [14] J. E. Sipe (1995). Photon wave functions. *Physical Review A* 52, 1875–1883. <https://doi.org/10.1103/PhysRevA.52.1875>.
- [15] Y. Aharonov, D. Z. Albert, and L. Vaidman (1988). How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100. *Physical Review Letters* 60, 1351–1354. <https://doi.org/10.1103/PhysRevLett.60.1351>.
- [16] O. Hosten and P. Kwiat (2008). Observation of the spin Hall effect of light via weak measurements. *Science* 319, 787–790. <https://doi.org/10.1126/science.1152697>.
- [17] K. Y. Bliokh, F. J. Rodríguez-Fortuño, F. Nori, and A. V. Zayats (2015). Spin-orbit interactions of light. *Nature Photonics* 9, 796–808. <https://doi.org/10.1038/nphoton.2015.201>.
- [18] L. F. O. Costa and J. Natário (2015). Center of mass, spin supplementary conditions, and the momentum of spinning particles. In *Fundamental Theories of Physics*, 215–258. Springer. https://doi.org/10.1007/978-3-319-18335-0_6.
- [19] G. Janson and R. Lopp (2025). General relativistic center-of-mass coordinates for composite quantum particles. *Physical Review D* 111, 064005. <https://doi.org/10.1103/PhysRevD.111.064005>.
- [20] R. Gerritsma, G. Kirchmair, F. Zähringer, E. Solano, R. Blatt, and C. F. Roos (2010). Quantum simulation of the Dirac equation. *Nature* 463, 68–71. <https://doi.org/10.1038/nature08688>.
- [21] A. I. Harte and M. A. Oancea (2022). Spin Hall effects and the localization of massless spinning particles. *Physical Review D* 105, 104061. <https://doi.org/10.1103/PhysRevD.105.104061>.