Mass Is What a Closed Loop of Frames Refuses to Forget

Registered research scaffold — no numerical result

Brecht Corbeel
Absolute Digital Publishers
absolutedigitalpublishers.com

September 20, 2026

Abstract

A sequence of frame changes that returns every classical parameter to zero — translate, boost, translate back, boost back — is the identity element of the underlying symmetry group. Acting on a quantum state, the same sequence need not return the identity operator: it can return a phase. For the Galilei group that phase is fixed exactly, to all orders, by mass, because mass enters the boost–momentum commutator as a central charge and every higher Baker–Campbell–Hausdorff term vanishes identically. For the Poincaré group no continuous central charge exists; the corresponding commutator is proportional to the energy operator rather than the identity, so the analogous loop phase is state dependent, and mass instead enters as the ordinary Casimir invariant. This draft registers, without resolving against experiment, one operational estimator, $m_{\mathrm{op}}[\mathcal{L}] = \hbar\,\Phi[\mathcal{L}]/\mathcal{A}[\mathcal{L}]$, built from a closed loop's dimensionless phase debt Φ divided by its specific-action invariant $\mathcal{A}$ (units $\mathrm{m}^2\,\mathrm{s}^{-1}$), so that $\hbar/\mathcal{A}$ is dimensionally a mass. It derives the Galilei-loop phase from the centrally extended commutator by an exact, non-perturbative Baker–Campbell–Hausdorff truncation; derives the Poincaré loop's failure to centralize from the same calculation performed on the ordinary relativistic algebra; and recovers the Galilei phase as the explicit $c \to \infty$ contraction of the Poincaré one. It states, but does not compute, the composite-clock instance of the same estimator, in which a system's internal energy contributes to its rest mass and a proper-time loop between two worldlines makes that contribution phase-visible. The construction is then confronted with a genuine, decade-long dispute in atom-interferometer physics: the standard light-pulse gravimeter phase is mass-independent by the weak equivalence principle, so the composite-clock phase this paper isolates must be a logically separate signature of local position invariance, engineered to survive after recoil phase and laser phase — both several orders of magnitude larger — are subtracted by a differential protocol. Two exact, explicitly labeled illustrative evaluations of the derived formulas are reported for idealized, unmeasured parameter choices; neither is a simulation, a measurement, or a claimed experimental result. The construction is rejected if the phase-inferred mass disagrees with the independently known mass or internal-energy gap once recoil, laser, path, and internal-energy terms are fully accounted for by a reviewed differential protocol, if the Galilei loop's phase fails the exact recoil-formula identity it must reproduce, or if the Poincaré loop's state-dependent phase fails to converge to the Galilei central charge under the stated contraction. No differential-interferometer simulation, fitted curve, numerical mass value, or external review is reported.

Keywords: mass superselection; Bargmann central extension; Galilei group; Poincaré group; group contraction; atom interferometry; gravitational redshift; proper time; equivalence principle.
Series: Revolutionary Life 75, Series D, D03.
Status: Registered scaffold. The full recoil/laser-phase interferometer accounting, the differential quantum-clock-interferometry implementation, the 60-source priority review, and external review are incomplete; Section 11 states what each one blocks.

1 The question is what a closed loop is allowed to remember

Take any sequence of symmetry operations — translations, velocity boosts, elapsed proper time — whose net effect, read off the group's own classical parameters, is the identity. Call such a sequence a *frame loop* $\mathcal{L}$. A ruler, a stopwatch, and a velocity meter all report that nothing happened: the net translation is zero, the net boost is zero, the net elapsed coordinate time matches whatever a stationary observer would have recorded.

The question this scaffold registers is narrow and entirely quantum: does the state on which $\mathcal{L}$ acts also return to itself, and if not, what invariant does the leftover phase measure?

Two loops carry this paper, chosen because each can be built from ordinary interferometer hardware without inventing new physics. The first, $\mathcal{L}_{\mathrm{G}}$, lives in the nonrelativistic (Galilei) symmetry group: translate the apparatus by a fixed vector $\mathbf{b}$, apply a velocity boost $\mathbf{v}$ to the particle, translate by $-\mathbf{b}$, boost by $-\mathbf{v}$. As abstract group elements this composes to the identity, because ordinary translations and ordinary Galilean boosts, treated as instantaneous operations, generate an abelian subgroup. The second, $\mathcal{L}_{\mathrm{P}}$, lives in spacetime itself: send a system from event A to event B along one worldline and compare it against the reverse of a different worldline that also connects A to B. The loop closes in space and time, but the two worldlines can carry different proper time, because proper time is path dependent once gravity or acceleration is present.

It is already known, and not this scaffold's contribution, that a quantum representation of the Galilei group requires a centrally extended algebra in which mass appears as a fixed number multiplying the identity operator, and that no analogous continuous central charge exists for the semisimple Lorentz factor of the Poincaré group [1, 2]. It is also established that a composite system's internal energy contributes to its rest mass, and that two branches of a matter-wave interferometer that traverse different proper time acquire a relative phase set by that contribution [10]. What is not established, and what this scaffold treats as an open construction rather than a discovery, is that both phases are usefully packaged as one operational estimator, $m_{\mathrm{op}}[\mathcal{L}] = \hbar\, \Phi[\mathcal{L}]/\mathcal{A}[\mathcal{L}]$, extracted as a slope across swept loop parameters rather than read from a single phase, and that this estimator can be run into a real, unresolved dispute over what an atom interferometer is allowed to call a measurement of mass [6–8].

This draft does not resolve that dispute. It derives the two phases from their respective commutators, connects them through the ordinary $c \to \infty$ contraction, states the exact arithmetic form the dispute's resolution must take, and reports two illustrative closed-form evaluations, explicitly labeled as neither simulations nor measurements. No differential-interferometer result, fitted visibility curve, or numerical mass value is reported anywhere in this scaffold.

2 What Bargmann and the central charge already settled

Bargmann's theorem on unitary ray representations of continuous groups establishes exactly when a projective representation of a Lie group can, and cannot, be lifted to an ordinary unitary representation of a central extension, and identifies the Galilei group's mass parameter as precisely such an unremovable central charge [1]. Lévy-Leblond works out the resulting nonrelativistic algebra explicitly, including the centrally extended boost–momentum commutator this paper uses without rederiving [2]. Giulini gives the companion statement on the algebra's physical content: taken as an exact, literal statement about all admissible states, the same central extension forbids coherent superpositions of different mass eigenvalues, the Bargmann superselection rule, and clarifies the rule's status once embedded in a genuinely relativistic theory [4]. Zych and Greenberger address directly the puzzle this scaffold inherits rather than resolves: ordinary atomic clocks routinely superpose internal states of different rest mass, apparently violating the rule, and the resolution offered in that literature is that the rule survives only as a nonrelativistic limit rather than as an exact relativistic law [5]. None of these four results is rederived here; they are the fixed mathematical floor the constructions of Sections 4–5 stand on.

Inönü and Wigner's contraction procedure is the second inherited tool: a prescription for obtaining one Lie group and its representations as a controlled limit of another, holding specified physical quantities fixed while a parameter — here $1/c$ — is sent to zero [3]. This paper applies that general procedure to one explicit commutator rather than invoking it as a slogan; the general theorem is inherited, the specific limit computed in Section 5 is this paper's application of it.

On the interferometry side, the ordinary relativistic free-particle phase $\Delta\phi = -mc^2\Delta\tau/\hbar$ and its observable consequence for two branches with different proper time were proposed as an interferometric witness of general-relativistic proper time, with the internal-energy-dependent term isolated by subtracting the shared free-particle phase [10], and were reframed explicitly as a quantum analogue of the twin paradox with a fully worked photon-clock treatment [11]. Gravitationally induced matter-wave interference itself long predates either proposal, established by the first neutron-interferometer demonstration [15], and the standard three-pulse light-pulse atom interferometer used throughout this paper's confound analysis is the stimulated-Raman configuration in continuous use since the early 1990s [16], with its phase computed by the classical-action (ABCD) method formalized as a general path-integral technique for atom interferometry [14].

Finally, the objection this paper spends its second half addressing is itself inherited, not invented for this scaffold. Müller, Peters, and Chu proposed reading an ordinary atom gravimeter's phase as an implicit measurement of gravitational redshift at the atom's Compton frequency [6]; Wolf and coauthors computed the standard phase explicitly and showed it carries no mass dependence [7]; Sinha and Samuel sharpened the resulting objection into the claim that the atom in that configuration is not functioning as a clock at all [8]. Roura's quantum-clock-interferometry analysis and Di Pumpo and collaborators' general redshift-test accounting are the literature's answer, isolating the conditions under which a genuine internal-energy signature survives [9, 13], and Overstreet and collaborators' dual-species equivalence-principle test is cited here only as evidence about the size of the confounding terms this construction must engineer around, not as a test of this paper's estimator [12]. No claim in this paper's remaining sections may present any of the results in this section as a finding of this project.

3 A frame loop is defined by what it does not change classically

Definition 3.1 (Frame loop)**.** *A frame loop $\mathcal{L}$ is a finite, ordered sequence of symmetry operations drawn from a declared kinematical group — spatial translations, velocity boosts, and elapsed proper time among them — such that the composition of the operations' classical parameters, read off the group itself, is the identity element.*

A frame loop is a statement about the group's own bookkeeping, not about any particular representation. Represented on a Hilbert space by unitaries $U_1, \ldots, U_n$ corresponding to the loop's operations in order, the operator product $U_n \cdots U_1$ need not equal the identity operator even though the classical composition is trivial, because the generators of noncommuting operations can carry a representation-dependent correction that classical bookkeeping cannot see. When it does not, write

$$U_n \cdots U_1 = e^{i\Phi[\mathcal{L}]}\,\mathbb{1}, \tag{1}$$

and call $\Phi[\mathcal{L}] \in \mathbb{R}$ the loop's **phase debt**: the correction the representation charges for a sequence the classical group considered free. Equation (1) is a statement about a specific loop and a specific representation; it is not assumed to hold for every loop or every representation, and Sections 4–5 exhibit one case in which it holds exactly and one in which the left-hand side is not proportional to the identity at all.

Definition 3.2 (Specific-action invariant)**.** *For a loop whose phase debt is proportional to a single loop-dependent kinematic invariant, write that invariant $\mathcal{A}[\mathcal{L}]$ and require it to carry units of $\mathrm{m}^2\,\mathrm{s}^{-1}$, the units of $\hbar/\mathrm{kg}$. Two instances used in this paper are $\mathcal{A}[\mathcal{L}_\mathrm{P}] = c^2\Delta\tau$ for a proper-time loop and $\mathcal{A}[\mathcal{L}_\mathrm{G}] = \mathbf{v}\cdot\mathbf{b}$ for a boost–translation loop.*

Definition 3.3 (Operational mass estimator)**.**

$$m_\mathrm{op}[\mathcal{L}] = \frac{\hbar\,\Phi[\mathcal{L}]}{\mathcal{A}[\mathcal{L}]}. \tag{2}$$

Equation (2) is dimensionally a mass: $[\hbar] = \mathrm{kg}\,\mathrm{m}^2\,\mathrm{s}^{-1}$, $[\mathcal{A}] = \mathrm{m}^2\,\mathrm{s}^{-1}$, and Φ is dimensionless as a phase must be, so $[\hbar/\mathcal{A}] = \mathrm{kg}$. Appendix B lists this as an exact, symbolic exponent-arithmetic check performed for both instances of $\mathcal{A}$ listed above; it is a unit check, not a claim that $m_\mathrm{op}[\mathcal{L}]$ equals any particular mass for any particular apparatus. Definition 3.3 is this paper's proposed packaging. It has scientific content beyond notation only if, after the calculations of Sections 4–8 are carried through and confronted with a reviewed differential-interferometer accounting, $m_\mathrm{op}[\mathcal{L}]$ recovers an independently known mass or internal-energy gap within a stated, non-post-hoc error budget for at least one physically realizable loop; Section 8 states this as an arithmetic kill criterion rather than a hope.

A single measured phase at one fixed value of $\mathcal{A}$ cannot, by itself, distinguish m_op from an unknown additive phase offset — an uncalibrated pulse timing, an unaccounted path-length difference, or a stray field looks identical, in a single reading, to a rescaling of $\mathcal{A}$. Section 6 therefore defines the estimator actually proposed for use as a derivative across swept loop parameters, not as Eq. (2) evaluated once.

4 The Galilei loop's phase debt is exact, not approximate

Represent $\mathcal{L}_\mathrm{G}$ — translate by $\mathbf{b}$, boost by $\mathbf{v}$, translate by $-\mathbf{b}$, boost by $-\mathbf{v}$ — on a Hilbert space carrying a unitary representation of the Galilei group. A spatial translation is $T(\mathbf{b}) = \exp(-i\mathbf{b}\cdot\mathbf{P}/\hbar)$, built from the momentum operator $\mathbf{P}$; a boost is $G(\mathbf{v}) = \exp(-i\mathbf{v}\cdot\mathbf{K}/\hbar)$, built from the boost generator $\mathbf{K}$. As group elements, $T(\mathbf{b})G(\mathbf{v})T(-\mathbf{b})G(-\mathbf{v}) = \mathbb{1}$, since ordinary translations and Galilean boosts commute as classical parameters. The corresponding operator product need not be the identity, because $[\mathbf{K}, \mathbf{P}]$ need not vanish.

It does not vanish. Bargmann's theorem, applied to the Galilei group by Lévy-Leblond, shows that a faithful, genuinely quantum-mechanical (projective) representation of the Galilei group requires

$$[K_i, P_j] = i\hbar\, m\, \delta_{ij}\,\mathbb{1}, \tag{3}$$

with m a fixed real number — the representation's mass label — multiplying the identity operator across the entire representation [1, 2]. "Central" means exactly this: the commutator commutes with every generator in the algebra, including $\mathbf{K}$ and $\mathbf{P}$ themselves. Equation (3) is inherited; nothing in this section modifies it.

Proposition 4.1 (Exact Galilei loop phase)**.** *Let* $X = -i\mathbf{b}\cdot\mathbf{P}/\hbar$ *and* $Y = -i\mathbf{v}\cdot\mathbf{K}/\hbar$*. Under Eq. (3),*

$$T(\mathbf{b})\,G(\mathbf{v})\,T(-\mathbf{b})\,G(-\mathbf{v}) = \exp\!\left(\frac{im}{\hbar}\,\mathbf{b}\cdot\mathbf{v}\right)\mathbb{1}, \tag{4}$$

with no small-loop or leading-order approximation.

Derivation. Substituting Eq. (3),

$$[X, Y] = \frac{im}{\hbar}\,\mathbf{b}\cdot\mathbf{v}\,\mathbb{1}. \tag{5}$$

Because $[X, Y]$ is itself proportional to the identity, it commutes with both X and Y: $[X, [X, Y]] = [Y, [X, Y]] = 0$. Every Baker–Campbell–Hausdorff term beyond the first commutator involves a nested commutator of X or Y with $[X, Y]$, and every such term vanishes identically. The Zassenhaus/BCH expansion of $e^X e^Y e^{-X} e^{-Y}$ therefore truncates exactly at

$$e^X e^Y e^{-X} e^{-Y} = e^{[X,Y]}, \tag{6}$$

with no residual series to bound or discard, which is Eq. (4). □

Proposition 4.1 is the same structural fact that gives an Aharonov–Bohm phase or a Berry-phase holonomy: a closed path in a classical parameter space acquires a phase fixed by a curvature living on that space, even though the endpoint and the start are classically identical. Here the "curvature" is the commutator in Eq. (3) and the "charge" coupling to it is mass; the analogy is a guide to intuition, not a claim that a background gauge field exists on the group manifold, and the phase's independence from the loop's interior path — so long as the declared net translation and net boost match — follows from Eq. (4) being built only from the loop's endpoints in group-parameter space, exactly as for a Wilson loop.

Comparing Eq. (4) with Definition 3.3 identifies

$$\Phi[\mathcal{L}_\mathrm{G}] = \frac{m\,\mathbf{b}\cdot\mathbf{v}}{\hbar}, \qquad \mathcal{A}[\mathcal{L}_\mathrm{G}] = \mathbf{v}\cdot\mathbf{b}, \tag{7}$$

so that $m_\mathrm{op}[\mathcal{L}_\mathrm{G}] = \hbar\Phi[\mathcal{L}_\mathrm{G}]/\mathcal{A}[\mathcal{L}_\mathrm{G}] = m$ by construction: Eq. (7) is definitionally circular in m, and Section 6 explains why that circularity is the correct starting point rather than a defect, and how the estimator actually proposed for use breaks it.

Because m enters only through the fixed central charge, and $\mathbf{b}, \mathbf{v}$ are the loop's own declared parameters, $\Phi[\mathcal{L}_\mathrm{G}]$ does not depend on which external inertial frame is used to describe the apparatus: boosting the whole laboratory relabels positions and velocities for an outside observer but does not change which group elements were composed, nor the central charge attached to the representation. Nothing about this phase disturbs conservation of energy or momentum: $\Phi[\mathcal{L}_\mathrm{G}]$ is not a new Hamiltonian term but a bookkeeping consequence of how already-unitary operators compose, and the product of unitaries in Eq. (4) is unitary regardless of whether the factors commute.

5 The Poincaré loop has no central charge to owe

Run the identical construction on the Poincaré group. The relativistic boost–momentum commutator is

$$[K_i, P_j] = \frac{i\hbar}{c^2}\,\delta_{ij}\,H, \tag{8}$$

with H the total energy operator [2]. Equation (8) is not proportional to the identity; it is proportional to H, an operator with a spectrum that differs from state to state. Repeating the derivation of Proposition 4.1 with $X' = -i\mathbf{b}\cdot\mathbf{P}/\hbar$, $Y' = -i\mathbf{v}\cdot\mathbf{K}/\hbar$, and Eq. (8) in place of Eq. (3) gives, for a state with sharp energy expectation $\langle H\rangle$,

$$[X', Y'] = \frac{i\langle H\rangle}{\hbar c^2}\,\mathbf{b}\cdot\mathbf{v}\,\mathbb{1} \tag{9}$$

on that state only; on the algebra itself, $[X', Y']$ is proportional to H, not to the identity, so it does not commute with $\mathbf{K}$ or $\mathbf{P}$ in general and the BCH truncation used in Proposition 4.1 does not apply representation-wide. A central extension, by definition, assigns the same phase to every vector in the representation; a phase that varies with $\langle H\rangle$ is an ordinary consequence of ordinary dynamics, removable in principle by working throughout with true, non-projective unitary representations of the Poincaré group, not evidence of a hidden extension.

This is Bargmann's own cohomology result applied to a semisimple factor: the Lorentz group admits no continuous central charge, so relativistic quantum mechanics carries no analogue of the Galilei mass superselection rule [1, 4]. Mass in special relativity is instead the ordinary Casimir invariant $P^\mu P_\mu = m^2c^2$, diagonalized on states, not a phase a loop leaves behind. This is a restriction inherited from Bargmann's theorem, not a new no-go result derived in this paper.

The Galilei phase is the Poincaré non-extension's $c \to \infty$ shadow

The two pictures are not independent theories sharing a symbol; the Galilei phase is the explicit low-velocity residue of the Poincaré non-extension. Write the rest-frame energy operator as $H = Mc^2 + H_\mathrm{NR}$, separating rest energy from whatever nonrelativistic energy — kinetic, internal — sits on top of it. Equation (9)'s state-dependent phase becomes

$$\frac{\langle H\rangle}{\hbar c^2}\,\mathbf{b}\cdot\mathbf{v} = \frac{M}{\hbar}\,\mathbf{b}\cdot\mathbf{v} + \frac{\langle H_\mathrm{NR}\rangle}{\hbar c^2}\,\mathbf{b}\cdot\mathbf{v}. \tag{10}$$

Proposition 5.1 (Contraction to the Galilei central charge)**.** *Holding* $\mathbf{b}$*,* $\mathbf{v}$*,* M*, and* $\langle H_\mathrm{NR}\rangle$ *fixed and sending* $c \to \infty$ *in Eq. (10), the second term vanishes as* c^{-2}

and the first survives unchanged, so the state-dependent relativistic phase converges to $M\mathbf{b}\cdot\mathbf{v}/\hbar$: a fixed number multiplying every vector in the fixed-M sector, identical to Eq. (7) with $m \to M$.

Proposition 5.1 is this paper's application of the general Inönü–Wigner contraction procedure [3] to the one commutator in Eq. (8), rather than an assertion by slogan; Appendix B lists a numerical convergence check over a grid of increasing c as a self-check on the algebra above, not as a physical prediction about any real apparatus's velocity regime. The reading is that the nonrelativistic mass superselection rule is the shadow a fully relativistic, non-superselected world casts onto its own $c \to \infty$ limit, a point developed independently in work reconciling the rule with ordinary relativistic quantum mechanics [4]. Section 7 returns to the practical consequence: every real atomic clock superposes internal states of strictly different rest mass, which is direct evidence that nature is Poincaré symmetric rather than exactly Galilei symmetric, with the superselection rule surviving only as the unreachable limit of Proposition 5.1 [5].

6 A mass gauge is a slope, not a snapshot

Section 3 already flagged the trap: Eq. (2) evaluated at one fixed $\mathcal{A}$ cannot separate m from an unknown constant phase offset, since any uncalibrated timing error, unaccounted path length, or stray field contributes a phase indistinguishable, in a single reading, from rescaling $\mathcal{A}$. The estimator this paper actually proposes for use is a derivative across a swept family of loops sharing the same classical identity element,

$$m_{\mathrm{op}} = \frac{\hbar}{c^2}\,\frac{\partial \Delta\phi}{\partial \Delta\tau}, \qquad m_{\mathrm{op}}[\mathcal{L}_{\mathrm{G}}] = \hbar\,\frac{\partial \Phi}{\partial(\mathbf{v}\cdot\mathbf{b})}, \tag{11}$$

with the sign of the proper-time form fixed by the free relativistic phase $\Delta\phi = -mc^2\Delta\tau/\hbar$, so that $m_{\mathrm{op}} = -(\hbar/c^2)\,\partial\Delta\phi/\partial\Delta\tau$ recovers m identically whenever $\Delta\phi$ has no other dependence on $\Delta\tau$. Both forms in Eq. (11) require sweeping the loop's own parameter — hold time, translation distance, boost velocity — across at least two settings and extracting a slope, exactly as an energy gap is extracted from a swept detuning rather than a single photon count; a single unswept measurement reports a phase, not a mass gauge.

Two degenerate limits confirm the derivative form is doing the intended job rather than concealing a division problem. As $m \to 0$, $\Delta\phi \to 0$ at every fixed $\Delta\tau$ in the proper-time form, so the slope — and therefore m_{op} — vanishes identically: a massless worldline has no proper time to speak of, and the formula correctly returns nothing to divide rather than an indeterminate ratio. If instead the two branches are held at equal proper time, $\Delta\tau = 0$, then $\Delta\phi = 0$ regardless of m, and a single measurement at that one setting is a 0/0 statement about mass, not a small number; both limits are trivial once stated and are exactly why Eq. (11) is defined as a slope over a family rather than a ratio at one point in it.

Proposition 6.1 (Recoil-phase identity)**.** *For the Galilei loop realized by a two-photon stimulated-Raman recoil kick, with effective wavevector k_{eff}, recoil velocity $v_r = \hbar k_{\mathrm{eff}}/m$ by definition of a momentum-conserving kick, and translation leg $b = v_r T$ over hold time T,*

$$\Phi[\mathcal{L}_{\mathrm{G}}] = \frac{m\,v_r b}{\hbar} = \frac{\hbar k_{\mathrm{eff}}^2 T}{m}, \tag{12}$$

identically, for every m, k_{eff}, and T.

Equation (12) is an algebraic consequence of substituting $v_r = \hbar k_{\mathrm{eff}}/m$ into the Galilei loop's own phase debt, Eq. (7); the right-hand side is the ordinary recoil phase already used throughout precision atom interferometry to extract $\hbar/m$ via the path-integral formalism for atomic interferometry [14, 16]. Proposition 6.1 is therefore a consistency check that the loop construction of Section 4 reproduces a quantity the field already measures by an independent route, not a new physical prediction; Appendix B lists its verification over a numerical parameter grid as a self-check, and Appendix A reports one fully worked, explicitly labeled illustrative evaluation of Eq. (12) for stated idealized sodium-interferometer parameters. Because $\Phi[\mathcal{L}_{\mathrm{G}}]$ was constructed from m in the first place, inverting Eq. (2) on this loop alone recovers m by definition; this circularity is exactly why Section 7 moves to a case — the composite-clock proper-time loop — where the quantity being estimated is not already fixed by the construction, and why Section 8 states the actual test as a differential-protocol accounting rather than as this loop's self-consistent recoil arithmetic.

7 A composite clock's mass depends on what it is doing inside, and the standard interferometer already forbids reading that off

A composite system with an internal excitation of energy ΔE above its ground configuration carries additional rest mass $\Delta E/c^2$ by ordinary mass–energy equivalence, so that the total rest energy is $Mc^2 = M_0c^2 + H_{\mathrm{int}}$ for internal-energy operator H_{int}. Sending such a system along two paths that differ in proper time by $\Delta\tau$, with the internal state differing between paths — ground on one, excited on the other — gives branch phases differing by

$$\Delta\phi_{\text{clock}} = -\frac{\Delta E}{\hbar}\,\Delta\tau, \tag{13}$$

after the shared $M_0 c^2 \Delta\tau/\hbar$ term, common to both arms and therefore unobservable in any interference signal, is dropped. Equation (13) is the interferometric proper-time witness proposed by Zych, Costa, Pikovski, and Brukner [10] and reframed as a quantum twin paradox with a fully worked optical-clock treatment by Loriani and collaborators [11]; it is the $\mathcal{L}_{\text{P}}$ instance of Definition 3.3 with $\mathcal{A}[\mathcal{L}_{\text{P}}] = c^2\Delta\tau$ and $\Phi[\mathcal{L}_{\text{P}}] = \Delta\phi_{\text{clock}} = -\Delta E\,\Delta\tau/\hbar$ (dimensionless, since $[\Delta E][\Delta\tau]/[\hbar]$ is a pure number), so that $m_{\text{op}}[\mathcal{L}_{\text{P}}] = \hbar\Phi[\mathcal{L}_{\text{P}}]/(c^2\Delta\tau) = -\Delta E/c^2$ up to the sign convention fixed in Eq. (11): the composite-clock estimator targets the internal-energy mass excess, not the bare mass M_0, and unlike Proposition 6.1 this quantity is not fixed by the construction itself — it is an independent physical fact about the clock transition that the loop is being asked to report.

This is the instance of the estimator with genuine content, and it is also the instance with a strong, decade-old rival account. In 2010, Müller, Peters, and Chu proposed that an ordinary light-pulse atom gravimeter's phase already realizes a measurement of the gravitational redshift at the atom's Compton frequency $mc^2/\hbar$, reading a standard interferometer as an implicit implementation of Eq. (13) [6]. Wolf, Blanchet, Bordé, Reynaud, Salomon, and Cohen-Tannoudji computed the standard three-pulse Mach–Zehnder phase by the ordinary classical-action method and found $\phi \approx k_{\text{eff}}\,g\,T^2$: built entirely from the laser wavevector, gravitational acceleration, and pulse timing, with no dependence on the atom's mass [7]. Sinha and Samuel sharpened this into a decisive claim: the atom in that configuration is not a clock ticking at its Compton frequency, because nothing in the standard phase formula reads that frequency out [8].

This objection is close to a theorem, not merely an engineering complaint. The weak equivalence principle — tested to high precision by exactly the dual-species interferometry technique cited below [12] — guarantees that a falling object's center-of-mass trajectory cannot depend on its mass or internal composition. Since the leading term of the standard gravimeter phase is built entirely from that trajectory and the laser field, mass is *required* to cancel from it; a surviving mass-dependence in that leading term would itself be evidence against the weak equivalence principle, not evidence for a working mass gauge. If $m_{\text{op}}[\mathcal{L}_{\text{P}}]$ is to mean anything operationally, it cannot live in the part of an interferometer phase the weak equivalence principle already forbids from carrying mass information.

Remark 7.1 (Two clauses, not one). The weak equivalence principle protects the laser/center-of-mass phase term, forcing it to be mass-independent; no version of the construction in this paper challenges that. Local position invariance — the universality of clock rates across spacetime position — is the logically separate clause that Eq. (13)'s internal-energy term actually probes, untouched by the weak equivalence principle's guarantee. Confusing the two clauses is, on this reading, the exact error the Müller–Peters–Chu proposal and the Wolf et al./Sinha–Samuel objection were arguing about.

8 What must be subtracted before $m_{\text{op}}[\mathcal{L}_{\text{P}}]$ means anything

Roura's quantum-clock-interferometry analysis identifies which terms in a full interferometer phase carry redshift information and under what conditions they can be isolated from the laser-driven center-of-mass phase the weak equivalence principle protects [9]. Di Pumpo and collaborators extend the bookkeeping to a general accounting of atomic-clock and atom-interferometric redshift tests, making explicit that the redshift signature in Eq. (13) requires a genuine internal-state superposition correlated with the spatial paths — not merely two branches of the same internal state — together with a symmetric, single-photon pulse design that keeps recoil and laser-phase contributions common to both branches so that they cancel rather than swamp the signal [13]. Neither reference computes or reports a value of $m_{\text{op}}[\mathcal{L}_{\text{P}}]$; both are cited here as the literature's stated path to isolating the term this paper's estimator targets, not as a result of this project.

The scale of what a differential protocol must cancel is large. For illustrative parameters worked in Appendix A — a single-photon clock transition near 698 nm, a strontium-mass atom, one second of hold time, and one meter of height separation — the laser/center-of-mass term scales as $k_{\text{eff}} g T^2$, some eight orders of magnitude above the clock term of Eq. (13); the recoil term of Proposition 6.1 scales roughly five orders of magnitude above it. Both figures are read directly off the formulas already derived, using the same illustrative constants as Appendix A; no interferometer has been run to produce them.

Definition 8.1 (Kill criterion, stated arithmetically). *Let $\widehat{m_{\text{op}}[\mathcal{L}_{\text{P}}]}$ denote the value a fully implemented, reviewed differential protocol — including the recoil term (Proposition 6.1), the laser/center-of-mass term, all path-dependent terms, and the internal-energy term of Eq. (13) — extracts by the slope method of Eq. (11), together with the protocol's own preregistered control budget ε. The construction is rejected as a measurement* *concept if*

$$\left|\widehat{m_{\mathrm{op}}[\mathcal{L}_{\mathrm{P}}]} - \frac{\Delta E}{c^2}\right| > \varepsilon \tag{14}$$

for the internal-energy gap ΔE independently known from clock spectroscopy, regardless of whether the underlying phase formula, Eq. (13), remains theoretically correct.

Definition 8.1 is deliberately arithmetic rather than qualitative: a formula that cannot be isolated from its own confounds at the precision a differential protocol claims is not an operational estimator, it is an equation. No value of $\widehat{m_{\mathrm{op}}[\mathcal{L}_{\mathrm{P}}]}$, no differential-protocol simulation, and no control budget ε has been computed, estimated, or bounded anywhere in this scaffold; both are explicitly registered required outputs, not numbers already in hand.

9 Degenerate limits where the operational mass must refuse to answer

At $m \to 0$ or $\Delta\tau \to 0$, Section 6's degenerate limits require m_{op} to vanish or become an explicit $0/0$ statement rather than a small finite number; an implementation that instead returns a finite value in either limit has a defect, not a discovery. Under the $c \to \infty$ contraction of Proposition 5.1, the state-dependent Poincaré-loop phase must converge to the fixed Galilei central charge and to nothing else; no additional finite term surviving the limit is compatible with Lévy-Leblond's algebra.

The weak-equivalence-protected-term countermodel. The standard atom-gravimeter phase already fully explains observed fringe shifts using only the laser field and free-fall trajectory, with mass provably absent from the leading term [7, 8]. This scaffold accepts the countermodel for the standard configuration outright: no version of $m_{\mathrm{op}}[\mathcal{L}_{\mathrm{P}}]$ is claimed to live there, and the decisive comparison is whether a differential protocol of the kind proposed by Roura and by Di Pumpo and collaborators can isolate a distinct, mass-dependent residual outside that protected term [9, 13].

The relabeled-redshift countermodel. A rival reading holds that Eq. (13) adds no operational content beyond the already-established, non-interferometric gravitational redshift measured by direct clock comparison across a height difference: the interferometric packaging changes the apparatus, not the physics being reported. The decisive comparison is whether the differential protocol's control budget ε in Definition 8.1 is actually smaller, for some accessible regime, than what non-interferometric clock comparison already achieves; if it is never smaller, the interferometric construction is at best a re-derivation, not a new operational gauge.

The superselection-taken-literally countermodel. Read as an exact statement about every physical state, Eq. (3) forbids any coherent superposition of different mass eigenvalues outright [1]. Every atomic clock ever operated superposes exactly such states, since an excited internal configuration carries strictly more rest mass than its ground configuration by $\Delta E/c^2$, and Ramsey-type coherence between them is measured routinely. This is not a contradiction but direct, everyday evidence that nature is Poincaré symmetric rather than exactly Galilei symmetric, with the rule surviving only as the unreachable $c \to \infty$ limit of Proposition 5.1 [4, 5]. The decisive comparison this paper owes, and has not yet run, is a stated bound on how nonrelativistic an apparatus must be before treating $\mathcal{A}[\mathcal{L}_{\mathrm{G}}]$'s central charge as exactly central, rather than approximately so, is numerically justified.

The confound-dominance countermodel. Even granting every derivation above, recoil and laser phase are several orders of magnitude larger than the target signal (Section 8), so a differential protocol's residual systematic error may remain at or above the signal's own size at any presently achievable precision. The decisive comparison is Definition 8.1 itself, evaluated against a concretely engineered protocol rather than against the idealized, confound-free evaluations of Appendix A.

The estimator-degeneracy countermodel. A single, unswept phase reading cannot distinguish m_{op} from a constant offset (Section 3). If the swept-parameter slope of Eq. (11) cannot be executed with enough independent settings at the precision Definition 8.1 requires, the estimator remains formally derivable but operationally untestable with current apparatus, and the paper must report that distinction rather than blur it.

Each countermodel above requires a stated calculation, not a rhetorical concession, before this construction may claim to have survived it.

10 How this could be wrong

The construction is rejected outright if Definition 8.1 is violated by a fully implemented, reviewed differential protocol, if Proposition 6.1's identity fails a symbolic or numerical check on any admissible parameter grid, or if Proposition 5.1's limit fails to converge to the Galilei central charge. It is withdrawn as non-novel if the completed 60-source priority search finds this specific operational packaging, $m_{\mathrm{op}}[\mathcal{L}] = \hbar\Phi[\mathcal{L}]/\mathcal{A}[\mathcal{L}]$ extracted as a swept-parameter slope, already published under a different name, or if every preregistered comparison in Section 9 reduces to a relabeling of quantities the field already reports without this estimator.

More local failures matter and each has a planned diagnostic. Treating Eq. (2) as a snapshot rather than a slope would reintroduce the offset-degeneracy Section 3 exists to block. Applying the exact BCH truncation of Proposition 4.1 to the Poincaré commutator, where it does not hold representation-wide, would manufacture a spurious central charge; Section 5 states explicitly why the truncation is state-dependent there. Quoting the Holevo-style confidence a differential protocol claims without an independently benchmarked control budget ε would make Definition 8.1 untestable by construction rather than by result. Presenting either illustrative evaluation in Appendix A as a measured quantity, rather than as an exact evaluation of a derived formula at stated idealized parameters, would misrepresent an arithmetic exercise as an experimental claim; the honesty audit in Appendix B exists specifically to catch this class of error before it reaches prose.

The most serious conceptual failure available to this paper is presenting the weak-equivalence-protected laser/center-of-mass term and the local-position-invariance-sensitive internal-energy term as if they were the same quantity at different precision, rather than as two logically distinct clauses of the equivalence principle (Section 7). That conflation is the exact error the Müller–Peters–Chu proposal and the Wolf et al./Sinha–Samuel objection spent a decade disputing, and repeating it here, even inadvertently, would forfeit the paper's only claim to having learned anything from that dispute.

11 What the kill criterion is still waiting for

Definition 8.1 is written so that the construction can be destroyed by an experiment. That is only meaningful once there is something to evaluate it against, and at present there is not. Four items are outstanding.

First, the priority review. The sixty-source hard floor and the title, abstract and equation-level priority search are incomplete. Until they close, no sentence in Section 1 moves from proposal to result language, and the operational mass of Definition 3.3 is withdrawn as non-novel if the search finds the same closed-loop quantity already established.

Second, the decomposition. Section 8 requires a boundary-resolved recoil, laser-phase and path decomposition, together with a finite-precision control-budget estimate, for at least one concrete transition and one concrete geometry. None exists. Definition 8.1 therefore cannot be evaluated against anything — it is a well-formed criterion with an empty domain, and saying so plainly is the difference between a registered protocol and a claim.

Third, the failure record. Section 10 sets out how the argument could break. It is rewritten only after an implemented differential protocol, its control budget and its residuals exist. Every kill criterion and countermodel test in Section 9 and Definition 8.1 survives that rewrite unchanged; substituting a post hoc justification for a preregistered criterion would leave the paper looking identical while removing the thing that makes it falsifiable.

Fourth, the ledger. Appendix B lists what the accompanying program verifies today, which is the exact algebraic identities of Appendix A and nothing dynamical. Solver, configuration, output and figure-generation hashes are appended when those artefacts exist.

The central-extension algebra of Section 4 and the contraction argument of Section 5 are complete as derivations and are not waiting on any of this. What is missing is the bridge from those identities to a measurement, and that bridge is the second item above.

12 A packaging has been registered, not a measurement performed

This scaffold defines an operational mass estimator built from a closed frame loop's phase debt divided by its specific-action invariant, and derives, rather than assumes, both instances that make the packaging worth naming: the Galilei loop's exact, non-perturbative phase from the centrally extended boost–momentum commutator, and the Poincaré loop's failure to centralize, together with its explicit $c \to \infty$ contraction back to the Galilei charge. Those two derivations, and the recoil-phase identity that checks the Galilei case against an independently known formula, stand regardless of what any interferometer eventually reports.

What remains genuinely open is whether the composite-clock instance of the estimator survives contact with a real interferometer once recoil phase, laser phase, and every other confound the field already knows how to compute are included in full through a reviewed differential protocol. This paper does not resolve the thirty-year dispute between Müller, Peters, and Chu on one side and Wolf and collaborators and Sinha and Samuel on the other; it restates that dispute as an arithmetic kill criterion, Definition 8.1, with a stated control budget rather than a rhetorical verdict. If a fully implemented protocol delivers $m_{\mathrm{op}}[\mathcal{L}_{\mathrm{P}}]$ within that budget of the independently known internal-energy gap, the packaging has scientific content beyond notation. If it does not, the honest conclusion is not that the algebra of Sections 4–5 was wrong; it is that a closed loop of frames can owe a Hilbert space a phase without any apparatus built so far being able to collect that debt cleanly enough to call it a reading of mass.

No redundancy value, capacity map, receipt time, or analogous scaling result exists in this paper's companion scaffolds; correspondingly, no differential-protocol simulation, fitted visibility curve, numerical mass value, control budget, or external review exists in this one. The scientific deliverable at this stage is a falsifiable registration of the packaging, together with a functioning refusal to report a number that has not been calculated.

A Exact illustrative evaluations (not measurements)

Every number in this appendix is an exact evaluation of a formula already derived in the body of this paper, for stated, idealized parameter choices. None is a simulation, a measurement, or the output of any code run to produce a physical result; the arithmetic is re-verified as a self-check in `simulate_frame_loop_mass.py` (Appendix B) and every value below agrees with that self-check to the tolerance stated there. No experiment corresponding to either evaluation has been performed.

A.1 The Galilei loop's recoil-identity evaluation

Take a sodium atom, $m = 3.8175 \times 10^{-26}\,\mathrm{kg}$, driven by a two-photon stimulated-Raman transition on the 589 nm line, effective wavevector $k_{\mathrm{eff}} = 2(2\pi/589\,\mathrm{nm}) = 2.1335 \times 10^{7}\,\mathrm{m}^{-1}$, comparable to the configuration of the earliest stimulated-Raman atom interferometer [16]. The recoil velocity is $v_r = \hbar k_{\mathrm{eff}}/m = 5.894\,\mathrm{cm\,s^{-1}}$. Holding the loop open for $T = 50\,\mathrm{ms}$ gives translation leg $b = v_r T = 2.947\,\mathrm{mm}$, loop area $\mathcal{A}[\mathcal{L}_{\mathrm{G}}] = v_r b = v_r^2 T = 1.737 \times 10^{-4}\,\mathrm{m^2\,s^{-1}}$, and

$$\Phi[\mathcal{L}_{\mathrm{G}}] = \frac{m\, v_r b}{\hbar} \approx 6.287 \times 10^{4}\ \mathrm{rad}. \tag{15}$$

By Proposition 6.1 this equals $\hbar k_{\mathrm{eff}}^2 T/m$ exactly, to the floating-point precision reported in Appendix B; inverting Eq. (2) therefore returns $m_{\mathrm{op}}[\mathcal{L}_{\mathrm{G}}] = m$ exactly, because $\Phi[\mathcal{L}_{\mathrm{G}}]$ was built from m in the first place (Section 6). This evaluation isolates what a confound-free reading of the Galilei loop looks like; it is not a claim that a real interferometer reports this number free of recoil, laser, or systematic phase, and no such claim is made anywhere in this paper.

A.2 The composite-clock loop's illustrative evaluation and its confounds

Take the weak-field proper-time rate $d\tau/dt \approx 1 + gz/c^2$, so that two branches held at height difference $\Delta z = 1\,\mathrm{m}$ for coordinate time $T = 1\,\mathrm{s}$ acquire $\Delta\tau \approx g\,\Delta z\,T/c^2 \approx 1.091 \times 10^{-16}\,\mathrm{s}$. Paired with a single-photon optical clock transition near 698 nm, comparable to a strontium-lattice-clock transition proposed for this kind of test [11, 13], $\Delta\nu \approx 4.295 \times 10^{14}\,\mathrm{Hz}$, $\Delta E = h\Delta\nu \approx 2.846 \times 10^{-19}\,\mathrm{J}$, giving an internal-energy mass excess $\Delta E/c^2 \approx 3.167 \times 10^{-36}\,\mathrm{kg}$ and, by Eq. (13),

$$\Delta\phi_{\mathrm{clock}} = 2\pi\,\Delta\nu\,\Delta\tau \approx 0.294\ \mathrm{rad}. \tag{16}$$

No experiment at this scale has been performed; the cited proposals argue it is near the edge of feasibility with current optical-clock coherence and large-scale atom-fountain technology, not a regime already demonstrated.

Using the same illustrative parameters, the laser/center-of-mass confound for the same single-photon transition ($k_{\mathrm{eff}} \approx 9.00 \times 10^{6}\,\mathrm{m}^{-1}$) scales as $k_{\mathrm{eff}}\, gT^2 \approx 8.8 \times 10^{7}$ rad, and the recoil confound for a strontium-mass atom ($m \approx 1.443 \times 10^{-25}$ kg) scales as $\hbar k_{\mathrm{eff}}^2 T/m \approx 5.9 \times 10^{4}$ rad: roughly eight and five orders of magnitude above the clock term, respectively. Both figures are exact evaluations of formulas already given in Section 8, for the same stated idealized parameters; neither is a simulation of an actual differential protocol, and Definition 8.1 is not evaluated by either of them.

B Computation and test ledger

The scaffold program `simulate_frame_loop_mass.py` implements no interferometer model and computes no physical mass. In self-check mode it performs the following exact/numerical sanity tests, all standard-library Python:

1. symbolic exponent-arithmetic dimensional checks that $\mathcal{A} = c^2\Delta\tau$ and $\mathcal{A} = \mathbf{v}\cdot\mathbf{b}$ both carry units $\mathrm{m^2\,s^{-1}}$, and that $\hbar/\mathcal{A}$ carries units kg for both loop types;

2. the recoil-phase identity of Proposition 6.1, $m\,v_r b/\hbar = \hbar k_{\mathrm{eff}}^2 T/m$ with $v_r = \hbar k_{\mathrm{eff}}/m$, verified over a preregistered grid of (m, k_{eff}, T) values;

3. the $c \to \infty$ contraction of Proposition 5.1, verified as a monotone, quadratically convergent numerical limit over a preregistered grid of increasing c;

4. invariance of $m_{\mathrm{op}}[\mathcal{L}_{\mathrm{G}}]$ under a stated rescaling $\mathbf{b} \to \lambda\mathbf{b}$, $\mathbf{v} \to \mathbf{v}/\lambda$ that holds $\mathcal{A}[\mathcal{L}_{\mathrm{G}}] = \mathbf{v}\cdot\mathbf{b}$ fixed, over a preregistered grid of λ;

5. exact round-trip recovery of an assumed mass m from $\Phi[\mathcal{L}_{\mathrm{G}}]$ and, separately, from the swept-slope form of Eq. (11) applied to the proper-time phase, both to machine precision;

6. an honesty audit of the configured phase-source registry, failing if any entry claims a measured or project-computed positive value rather than `not_yet_measured`, `not_yet_calculated`, or an explicitly labeled, non-measurement `exact_illustrative_evaluation` with a resolvable citation key.

The saved output labels these as scaffold validation. The program returns a nonzero exit status and a blocking message for `-simulate`, naming explicitly the full recoil/laser-phase interferometer accounting, the differential quantum-clock-interferometry implementation, the Poincaré-loop state-dependence treatment beyond the contraction check above, the solver dependency lock, the 60-source review gate, and external review as the incomplete prerequisites. A later implementation must not silently reuse the program name while relaxing these gates: it must add interferometer-model, state-validity, and control-budget tests appropriate to an actual differential-protocol calculation, write a versioned output schema, and update the dependency lock and environment record before any phase source may be promoted to a measured or project-computed status.

References

[1] V. Bargmann (1954). On unitary ray representations of continuous groups. *Annals of Mathematics* 59, 1–46. https://doi.org/10.2307/1969831.

[2] J.-M. Lévy-Leblond (1963). Galilei group and nonrelativistic quantum mechanics. *Journal of Mathematical Physics* 4, 776–788. https://doi.org/10.1063/1.1724319.

[3] E. Inönü and E. P. Wigner (1953). On the contraction of groups and their representations. *Proceedings of the National Academy of Sciences* 39, 510–524. https://doi.org/10.1073/pnas.39.6.510.

[4] D. Giulini (1996). On Galilei invariance in quantum mechanics and the Bargmann superselection rule. *Annals of Physics* 249, 222–235. https://arxiv.org/abs/quant-ph/9508002. DOI https://doi.org/10.1006/aphy.1996.0069.

[5] M. Zych and D. M. Greenberger (2019). Puzzling out the mass-superselection rule. *arXiv preprint*. https://arxiv.org/abs/1906.03725. DOI https://doi.org/10.48550/arXiv.1906.03725.

[6] H. Müller, A. Peters, and S. Chu (2010). A precision measurement of the gravitational redshift by the interference of matter waves. *Nature* 463, 926–929. https://doi.org/10.1038/nature08776.

[7] P. Wolf, L. Blanchet, C. J. Bordé, S. Reynaud, C. Salomon, and C. Cohen-Tannoudji (2010). Atom gravimeters and gravitational redshift. *Nature* 467, E1. https://doi.org/10.1038/nature09340.

[8] S. Sinha and J. Samuel (2011). Atom interferometry and the gravitational redshift. *Classical and Quantum Gravity* 28, 145018. https://doi.org/10.1088/0264-9381/28/14/145018.

[9] A. Roura (2020). Gravitational redshift in quantum-clock interferometry. *Physical Review X* 10, 021014. https://doi.org/10.1103/PhysRevX.10.021014.

[10] M. Zych, F. Costa, I. Pikovski, and Č. Brukner (2011). Quantum interferometric visibility as a witness of general relativistic proper time. *Nature Communications* 2, 505. https://doi.org/10.1038/ncomms1498.

[11] S. Loriani, A. Friedrich, C. Ufrecht, F. Di Pumpo, S. Kleinert, S. Abend, N. Gaaloul, C. Meiners, C. Schubert, D. Tell, É. Wodey, M. Zych, W. Ertmer, A. Roura, D. Schlippert, W. P. Schleich, E. M. Rasel, and E. Giese (2019). Interference of clocks: a quantum twin paradox. *Science Advances* 5, eaax8966. https://doi.org/10.1126/sciadv.aax8966.

[12] C. Overstreet, P. Asenbaum, T. Kovachy, R. Notermans, J. M. Hogan, and M. A. Kasevich (2018). Effective inertial frame in an atom interferometric test of the equivalence principle. *Physical Review Letters* 120, 183604. https://doi.org/10.1103/PhysRevLett.120.183604.

[13] F. Di Pumpo, C. Ufrecht, A. Friedrich, E. Giese, W. P. Schleich, and W. G. Unruh (2021). Gravitational redshift tests with atomic clocks and atom interferometers. *PRX Quantum* 2, 040333. https://doi.org/10.1103/PRXQuantum.2.040333.

[14] P. Storey and C. Cohen-Tannoudji (1994). The Feynman path integral approach to atomic interferometry: a tutorial. *Journal de Physique II* 4, 1999–2027. https://doi.org/10.1051/jp2:1994103.

[15] R. Colella, A. W. Overhauser, and S. A. Werner (1975). Observation of gravitationally induced quantum interference. *Physical Review Letters* 34, 1472–1474. https://doi.org/10.1103/PhysRevLett.34.1472.

[16] M. Kasevich and S. Chu (1991). Atomic interferometry using stimulated Raman transitions. *Physical Review Letters* 67, 181–184. https://doi.org/10.1103/PhysRevLett.67.181.