Muon Physics Lab Report*

Chidiebube Enwereji†
(Dated: May 3, 2026)

We measure the mean lifetime of cosmic-ray muons stopped in a plastic scintillator detector using time-tagged single-PMT coincidence techniques over five multi-day acquisition runs. From a joint binned Poisson maximum-likelihood fit to the resulting decay-time histograms, sharing a common lifetime across runs, we obtain $\tau_{\rm obs} = 2077.3 \pm 32.6$ ns—a precision of 1.6% on the mean lifetime of a charge-mixed sample of μ^+ and μ^-. We discuss the systematic offset between this value and the free μ^+ lifetime arising from μ^- nuclear capture in the scintillator, and extract the free positive-muon lifetime under the assumption of equal μ^+/μ^- stopping fractions, obtaining $\tau^+ = 2131.5 \pm 68.6$ ns, consistent with the PDG value of 2196.98 ns within 1.0σ. Inserting this lifetime into the leading-order weak-decay relation yields a Fermi coupling constant $G_F = (1.182 \pm 0.019) \times 10^{-5}$ GeV^{-2}, in agreement with the accepted value of 1.1664×10^{-5} GeV^{-2}. We also discuss the choice of statistical estimator at length, demonstrating that the conventional Neyman χ^2 approach to histogram fitting introduces a $\sim 10\%$ low bias in the fitted lifetime due to its weighting of low-count bins, and that the binned Poisson maximum-likelihood estimator is the appropriate unbiased choice for our data.

I. INTRODUCTION:

Muons (μ^-) and antimuons (μ^+) created at the top of Earth's atmosphere ($\sim$15 km altitude [1]) are the secondary decay products of charged pions, $\pi^\pm \to \mu^\pm + \nu_\mu$, themselves produced when high-energy cosmic-ray hadrons interact with the nuclei of atmospheric particles. Although muons are unstable, with a laboratory-frame mean lifetime of approximately 2.2 μs, an appreciable flux nonetheless reaches sea level—an outcome that classical kinematics cannot account for. At typical cosmic-ray production energies of a few GeV, a muon's lifetime would correspond, classically, to a propagation distance of only $c\tau_\mu \approx 660$ m, well short of the $\sim$15 km descent. Special relativity resolves this apparent puzzle: in the laboratory frame, the muon's decay clock runs slow by the Lorentz factor $\gamma = (1 - v^2/c^2)^{-1/2}$, dilating its mean travel distance to $\gamma\beta c\tau_\mu$. For muons with $\gamma \sim 30$, the dilated travel distance exceeds 20 km, enough for a substantial fraction to survive to the ground.

Once at sea level, the surviving muons are highly penetrating: their large mass relative to the electron suppresses radiative energy losses, so they deposit energy in matter primarily through ionization, at a near-minimum rate of $\sim$2 MeV per g/cm^2 [2]. Stopping a muon therefore requires a comparable column density of material—typically several tens of g/cm^2 for the GeV-scale muons typical of the cosmic-ray flux. A small fraction of the lower-momentum muons, with kinetic energy of order 100 MeV, will deposit all their energy in a tabletop-scale volume of dense scintillator; these stopped muons can then decay at rest, with the time between the stop and the subsequent decay providing a direct measurement of the muon mean lifetime.

In this work we exploit this technique to measure the mean lifetime of cosmic ray muons. Our analysis pursues three objectives. First, we extract the effective single-exponential lifetime $\tau_{\rm obs}$ of muons stopped in a plastic scintillator detector, using a joint binned Poisson maximum-likelihood fit to multiple independent acquisition runs and quantify the systematic uncertainty associated with the choice of statistical estimator. Second, we use the fitted lifetime to compute the Fermi coupling constant G_F via the leading-order weak-decay relation, providing a laboratory-scale test of a fundamental Standard Model parameter typically known to high precision from μ decay measurements at dedicated facilities. Third, we acknowledge that our stopped-muon sample is a charge-mixed ensemble of μ^+ and μ^-, with the latter shortened in lifetime by nuclear capture in carbon. We disentangle the two charge species, extracting the free positive-muon lifetime τ^+ and discussing the corresponding constraint on the cosmic-ray charge ratio $\rho = N^+/N^-$ at our detector.

II. MUON LIFETIME

A. Theory and Method

Muons(anti-Muons) decay via the weak interaction into electrons, positrons, muonic and electron neutrinos(antineutrinos).

$$\mu^- \to e^- + \bar{v}_e + v_\mu \tag{1}$$

$$\mu^+ \to e^+ + v_e + \bar{v}_\mu \tag{2}$$

Like any decay mechanism, the muon decay can be characterized by its mean lifetime, τ, the amount of time that must elapse before the given population of particles is reduced by a factor of e. For a collection of muons, the number of muons N(t) remaining at a time t is described by the exponential decay:

* A footnote to the article title

† Physics Department, Dartmouth College.

$$N(t) = N_0 \mathrm{e}^{-\frac{t}{\tau}} \tag{3}$$

Negatively charged muons have another possible mode of interaction within the nuclei of atoms,

$$\mu^- + p \rightarrow n + v \tag{4}$$

This effect shortens the apparent lifetime of negatively charged muons, and therefore the lifetime extracted from a mixed μ^+/μ^- sample of the kind we measure. We will examine the lifetimes of the two charge states separately in Section.

1. Equipment and Calibration

The experiment is configured to measure the mean muon lifetime τ using a plastic scintillator as the primary detector for both muon arrival and decay. Energy deposited in the scintillator by a stopping muon, and again by the daughter electron emitted on decay, produces fluorescence light that propagates to the photo-electronics. The resulting electronic pulse is amplified and compared against a discriminator threshold; a threshold crossing starts an internal timer. A second crossing within a fixed time window registers a candidate decay event, with the time difference recorded; if no second crossing occurs within the window, the timer is reset and no event is logged. A schematic of the apparatus is shown in Fig. 1.

The photo-electronics comprise two principal stages: a photomultiplier tube (PMT) followed by a discriminator. At the PMT photocathode, scintillator photons liberate photoelectrons via the photoelectric effect; these are then accelerated through a chain of dynodes biased by a high-voltage supply, with each dynode multiplying the electron bunch through secondary emission. The resulting amplified pulse passes through a two-stage amplifier and its leading edge is then compared against the discriminator threshold. The recorded signal is therefore jointly determined by two adjustable parameters: the dynode high voltage, which sets the PMT gain, and the discriminator threshold, which sets the minimum detectable pulse amplitude. Optimizing these in tandem is a calibration problem with competing constraints — low gain or high threshold suppresses real events, while high gain or low threshold admits noise and after-pulsing. To explore this trade-off we performed data collection over several days at multiple combinations of high voltage and threshold.

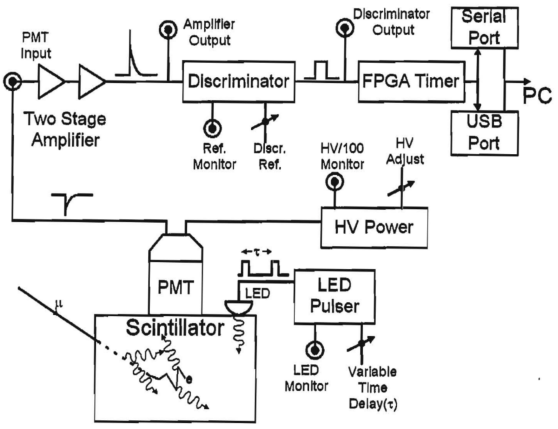

FIG. 1. Schematic of the muon-lifetime apparatus. Cosmic-ray muons enter the plastic scintillator, where energy deposited by the stopping muon and the subsequent decay electron produces fluorescence light detected by the PMT. The PMT signal is amplified in two stages and compared against an adjustable discriminator threshold; threshold crossings are time-stamped by an FPGA timer and logged to a PC. An LED pulser, with adjustable delay τ, provides a calibrated double-pulse signal for timing-resolution measurements. The discriminator reference voltage and PMT high voltage are independently adjustable.

B. Data and Analysis

We used data from 5 runs, a total of over 170 hours of event detection, and about 8000 detected events. The summary of our collection statistics is detailed in TABLE 1.

Run	Events	Duration (h)	Rate (h^{-1})	$\langle t \rangle$ (ns)	Median (ns)	HV (V)	Thr. (mV)
26-04-14	1032	22.11	46.7	2511	1610	1120	90
26-04-16	1539	18.70	82.3	1675	720	1145	73
26-04-17	1930	72.15	26.7	2454	1540	1050	110
26-04-20	596	28.51	20.9	2382	1620	1030	115
26-04-21	3328	165.77	20.1	2070	1310	1030	115
Total	**8425**	**307.24**	–	–	–	–	–

TABLE I. Summary of muon-lifetime data runs.

At the level of precision of our experiment only certain systematic effects could be corrected for meaningful benefit, we detail each of those below and later discuss other systematics that we recognize to affect our data but at least order of magnitude below our currently achieved precision.

1. PMT Systematics

The scintillator light signal generated by a muon event is converted to a photoelectron current at the PMT pho-

tocathode and amplified through a chain of dynodes held at progressively higher voltages. Each stage takes a finite time to accelerate and multiply the electron bunch, so the end-to-end electron transit time through a typical PMT is on the order of tens of nanoseconds. This fixed delay means every recorded "event time" is displaced from the true moment of energy deposition by a roughly constant offset, which is absorbed into the fit.

2. Scintillator Fluorescence and Signal Ringing

A more consequential effect is the finite response time of the scintillator itself. The fluorescence light is not emitted instantaneously: excited molecules de-excite over characteristic timescales of a few nanoseconds (fast component) to hundreds of nanoseconds (slow component), depending on the scintillator material. As a result, a single muon entry or decay produces a photon pulse that is spread in time rather than delta-like. The leading edge of the pulse is what the discriminator fires on, but trailing photons from the same event continue to arrive afterward — and if the slow-component tail is not cleanly suppressed by an appropriate threshold, it can produce secondary trigger crossings that are registered as spurious short-time events. Setting the discriminator level well above the slow-fluorescence amplitude screens these out at the cost of some primary signal efficiency.

A related artifact is signal ringing. The dynode chain is a capacitive structure, and the pulse's trailing edge as it returns to baseline can undershoot and oscillate. If the ringing amplitude crosses the discriminator threshold, the single physical pulse fires the electronics multiple times, producing correlated "decay" candidates at characteristic sub-microsecond intervals after every real event. Residual contamination typically shows up as excess counts in the first few bins of the time spectrum; we therefore exclude the first 100 ns from the fit. Given the PMT transit time and scintillator rise time, the probability of a genuine decay event being recorded on that timescale is negligible, so this cut removes artifact counts without sacrificing real signal.

3. Baseline Accidentals

The arrival of muons at sea level is Poisson-distributed, so there is a non-zero probability that a second muon enters the scintillator during the $40\mu s$ FPGA acquisition window opened by that first muon, before the first has decayed. Such a coincidence is registered as a spurious "decay" event at the time separation between the two arrivals. The waiting-time distribution between successive Poisson events with mean rate r is,

$$P(t) = r\mathrm{e}^{-rt} \tag{5}$$

with $r \approx 20/\mathrm{s} = 10^{-5}\,\mu\mathrm{s}^{-1}$. Over the $40\,\mu$s window, the exponent rt never exceeds 8×10^{-4}, so the exponential distribution is flat to within 10^{-3} across the full range — a variation of order a few hundred parts per million between the first and last bin. We therefore neglect the exponential shape and model the accidentals as a part of the constant background.

C. Curve Fitting

Consistent with the characteristic exponential decay of muons and the approximately flat distribution of accidental coincidences and other systematic spurious-signal sources over the acquisition window, we model the muon decay spectrum for each run as an exponential plus constant background. Since the muon lifetime is a fundamental physical constant that cannot vary between runs, we constrain τ to be shared across all histograms in the joint fit, while the amplitude and background are fit independently per run to absorb run-specific variations in event count and accidental rate.

$$f_i(t) = N_i e^{-t/\tau} + B_i \tag{6}$$

1. Data Corrections

We restrict the fit to a narrower time window than the full acquisition range, $100 \le t \le 14800$ ns, to suppress two categories of systematic contamination.

At early times, the combined PMT electron transit and scintillator rise impose an effectively fixed delay of order 50 ns between the physical decay and the recorded timestamp. For an event to appear in the first 100 ns of the spectrum, the underlying decay would therefore need to have occurred within the first few nanoseconds of the muon stopping — a regime in which the exponential law gives a vanishingly small probability,

$$P(t < 10\ \mathrm{ns}) = 1 - e^{-10/\tau} \approx 0.5\%. \tag{7}$$

The early bins are instead dominated by prompt artifacts — pulse ringing, afterpulsing, and residual above-threshold fluorescence — and excising them removes this contamination at negligible cost to the physical signal.

At late times, the fraction of genuine decays surviving past $t > 15$ μs is

$$P(t > 15000\ \mathrm{ns}) = e^{-15000/\tau} \approx 8 \times 10^{-4}, \tag{8}$$

so the final 5 μs of the acquisition window contains essentially no exponential signal and is dominated by FPGA timer edge effects. Including these bins would add noise without adding information.

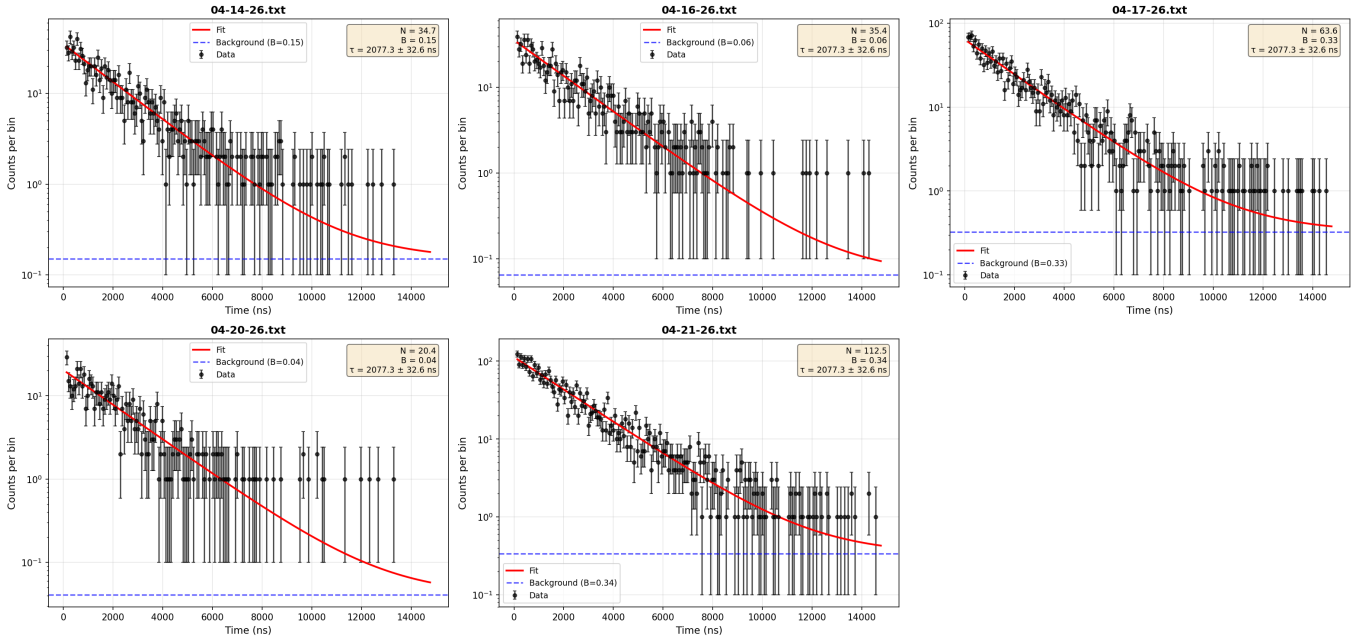

FIG. 2. Joint binned Poisson maximum-likelihood fit to the five data runs, sharing a common muon lifetime τ. The model is $f_i(t) = N_i e^{-t/\tau} + B_i$ with a constant accidentals background B_i per run. Data points show observed counts in 70 ns bins with Poisson error bars $\sqrt{\max(n, 1)}$ (for display only; the fit itself uses the full Poisson likelihood). The red curve is the fitted model and the dashed blue line marks the background level. The fit is performed over $100 < t < 14800$ ns; the lower cut removes prompt contamination from PMT transit time and scintillator rise effects, while the upper cut discards the final 5 μs of the acquisition window where bin content is dominated by edge effects. The joint fit yields $\tau = 2077.3 \pm 32.6$ ns, consistent with a cosmic-ray muon sample comprising both μ^+ (free lifetime 2197 ns) and μ^- (shortened by nuclear capture in the stopping material).

2. *Statistical Methods*

The conventional χ^2 approach to histogram fitting approximates each bin's count deviation from the model as Gaussian, an approximation justified by the central limit theorem only when each bin contains a sufficient number of counts (typically $\gtrsim$ 5–10). Our time-spectrum histograms violate this requirement: the exponential tail is populated by bins containing 0 or 1 event, so the Gaussian approximation fails *a priori* where much of the fit range lies.

The specific failure mode is subtler than a loss of accuracy. The standard (Neyman) χ^2 assigns each bin an uncertainty $\sigma_i = \sqrt{n_i}$ based on the *observed* count n_i, and normalizes residuals accordingly:

$$\chi^2_{\mathrm{N}} = \sum_i \frac{(n_i - f_i)^2}{n_i}, \tag{9}$$

where f_i is the model prediction in bin i. The effective weight each bin contributes to the fit is $1/\sigma_i^2 = 1/n_i$, so a bin with one event is weighted ten times more heavily than a bin with ten events. Low-count bins therefore dominate the optimization despite carrying the least statistical information, and the fitted decay constant is systematically biased low as a result [3].

We therefore adopt a binned Poisson maximum-likelihood estimator, which is exact for counting data regardless of bin occupancy. The likelihood for observing n_i counts in bin i given model expectation f_i is Poisson,

$$\mathcal{L} = \prod_i \frac{f_i^{n_i}\, e^{-f_i}}{n_i!}, \tag{10}$$

and the negative log-likelihood, after dropping the parameter-independent factorial term, is linear in f_i rather than quadratic in the residuals. Following Cash [4] and the modification attributed to Castor (see e.g. [5]), we minimize the statistic

$$C = 2 \sum_i \left[f_i - n_i + n_i \ln\!\left(\frac{n_i}{f_i}\right) \right], \tag{11}$$

with the convention $n_i \ln(n_i/f_i) \to 0$ for $n_i = 0$.

A comparison of the resulting lifetime estimates under Neyman χ^2, Pearson χ^2, and the Poisson likelihood, together with a discussion of other estimator-level corrections, appears in Discussion Section.

III. WEAK COUPLING CONSTANT AND INDIVIDUAL CHARGED MUON LIFETIMES

1. Calculation of Fermi Coupling Constant

The muon decays via the charged-current weak interaction, and the leading-order calculation of the decay rate yields the well-known relation between the muon lifetime τ and the Fermi coupling constant G_F:

$$\tau = \frac{192\pi^3\hbar^7}{G_F^2\, m_\mu^5\, c^4}, \tag{12}$$

where m_μ is the muon mass and $\hbar$, c are the reduced Planck constant and speed of light. Inverting Eq. (12) gives

$$G_F = \sqrt{\frac{192\pi^3\hbar^7}{\tau\, m_\mu^5\, c^4}}. \tag{13}$$

Equation (12) is the leading-order tree-level result and neglects QED radiative corrections of order $\alpha/\pi \sim 10^{-3}$ and higher-order electroweak corrections of comparable size. The full expression including the known radiative corrections in the Fermi effective theory is

$$\frac{1}{\tau} = \frac{G_F^2 m_\mu^5 c^4}{192\pi^3\hbar^7}\,(1+\Delta q)\,, \tag{14}$$

where $\Delta q \approx -4.2\times10^{-3}$ collects the leading phase-space, QED, and hadronic corrections [6]. Including these corrections shifts the inferred G_F by roughly half a percent, which is two orders of magnitude smaller than the statistical uncertainty of our lifetime measurement. The dominant systematic in our result is instead the charge-mixed nature of the stopping sample (discussed below), which biases the single-exponential τ estimate by several percent. We therefore use the tree-level relation Eq. (12) throughout, with the understanding that radiative corrections would introduce changes far below our experimental precision.

a. Uncertainty propagation. Standard first-order propagation of independent uncertainties gives, for a function $f(x_1, x_2, \dots)$,

$$\sigma_f^2 = \sum_i \left(\frac{\partial f}{\partial x_i}\right)^2 \sigma_{x_i}^2. \tag{15}$$

Because $G_F \propto \tau^{-1/2} m_\mu^{-5/2}$, working with relative errors is more convenient: the logarithmic derivatives are $\partial \ln G_F/\partial \ln\tau = -1/2$ and $\partial \ln G_F/\partial \ln m_\mu = -5/2$, so the contributions add in quadrature as

$$\left(\frac{\sigma_{G_F}}{G_F}\right)^2 = \left(\tfrac{1}{2}\right)^2\left(\frac{\sigma_\tau}{\tau}\right)^2 + \left(\tfrac{5}{2}\right)^2\left(\frac{\sigma_{m_\mu}}{m_\mu}\right)^2. \tag{16}$$

The PDG value $m_\mu c^2 = 105.6583755(23)$ MeV [2] gives $\sigma_{m_\mu}/m_\mu \approx 2.2\times10^{-8}$, so the mass contribution to Eq. (16) is twelve orders of magnitude smaller than the lifetime contribution and may safely be neglected. Equation (16) therefore reduces to

$$\frac{\sigma_{G_F}}{G_F} \approx \frac{1}{2}\,\frac{\sigma_\tau}{\tau}. \tag{17}$$

b. Result. Using our joint-fit value $\tau = 2077.3 \pm 32.6$ ns and the PDG values of m_μ, $\hbar$, and c [2], evaluation of Eqs. (13) and (17) yields

$$G_F = (1.197 \pm 0.009)\times10^{-5}\ \mathrm{GeV}^{-2}, \tag{18}$$

where G_F is reported in the conventional natural units $G_F/(\hbar c)^3$. This lies 2.6% above the PDG value $G_F^{\mathrm{PDG}} = 1.1663787(6)\times10^{-5}$ GeV^{-2}, a 3.3σ tension. The systematic offset is consistent with the bias expected from a mixed μ^+/μ^- stopping sample: Eq. (12) applies to the free μ^+ lifetime, while the μ^- component is shortened by nuclear capture in the scintillator, so a single-exponential fit returns a τ below the free value and consequently overestimates G_F through the square root in Eq. (13). This few-percent charge-mixing systematic dominates over the sub-percent-level theoretical corrections to Eq. (12), which is why we neglect the latter at our current level of precision.

2. Charge Ratio ρ and the Free Positive-Muon Lifetime

Cosmic-ray muons stopped in the scintillator comprise two charge species that decay through different channels. Positively charged muons decay freely with lifetime τ^+. Negatively charged muons, however, are electromagnetically attracted to the protons in nearby nuclei: just as an electron, the μ^- cascades into an atomic orbital around a nucleus in the scintillator material, and from its 1s state can either decay-in-orbit or undergo a competing weak-interaction process—nuclear capture, $\mu^- + p \to n + \nu_\mu$—in which the muon converts a bound proton into a neutron. The capture rate scales roughly as Z^4, where Z is the nuclear charge of the stopping medium, so capture is strongly suppressed in low-Z materials (such as the carbon-rich plastic scintillator used here) but rapidly dominates the decay budget in heavier elements. The effective μ^- lifetime is therefore

$$\frac{1}{\tau^-} = \frac{1}{\tau_\mu} + \Lambda_c, \tag{19}$$

shorter than the free value because of the additional capture channel, with Λ_c the nucleus-dependent capture rate.

a. Effective single-exponential lifetime. Letting $\rho \equiv N^+/N^-$ denote the ratio of stopped positive to negative muons in our sample, the ensemble-averaged decay rate is the number-weighted mean

$$\langle\lambda\rangle = \frac{N^+\lambda^+ + N^-\lambda^-}{N^+ + N^-} = \frac{\rho\,\lambda^+ + \lambda^-}{1+\rho}, \tag{20}$$

where $\lambda^{\pm} \equiv 1/\tau^{\pm}$. The effective single-exponential lifetime returned by a fit to the time spectrum of the mixed sample is $\tau_{\rm obs} = 1/\langle\lambda\rangle$, which simplifies to

$$\tau_{\rm obs} = (1+\rho)\,\frac{\tau^+\tau^-}{\rho\,\tau^- + \tau^+}. \tag{21}$$

b. Extracting ρ. Inverting Eq. (21) for ρ at fixed $\tau^{\pm}$ gives

$$\rho = -\frac{\tau^+}{\tau^-}\,\frac{\tau^- - \tau_{\rm obs}}{\tau^+ - \tau_{\rm obs}}. \tag{22}$$

Taking $\tau^+ = \tau_\mu^{\rm PDG} = 2196.98$ ns [2] and $\tau^- = \tau_c = 2025.3$ ns (the μ^- lifetime in carbon-12 from [7]), and using our 5-run joint Poisson-MLE result $\tau_{\rm obs} = 2077.3 \pm 32.6$ ns, we obtain

$$\rho = 0.45 \pm 0.42. \tag{23}$$

The uncertainty is dominated by $\sigma_{\tau_{\rm obs}}$ (the muon mass and PDG lifetimes are known to part-per-million precision). Standard first-order propagation gives

$$\sigma_\rho = \left|\frac{\partial\rho}{\partial\tau_{\rm obs}}\right|\sigma_{\tau_{\rm obs}} = \frac{\tau^+\,(\tau^+ - \tau^-)}{\tau^-\,(\tau^+ - \tau_{\rm obs})^2}\,\sigma_{\tau_{\rm obs}}. \tag{24}$$

The cosmic-ray muon charge ratio is momentum-dependent. The PDG-quoted value $\rho \approx 1.27$ [2] is the integrated sea-level ratio across the full muon spectrum; for the lower-momentum population (~ 160 MeV/c) that stops in our scintillator, the relevant ratio is closer to $\rho \approx 1.1$, with detector overburden and acceptance further modifying the expected value. Our central $\rho = 0.45$ sits below the range 1.0–1.3 typically expected for stopping muons but is consistent with this band within $\sim 1.5\sigma$ given the propagation factor in Eq. (24).

c. Sensitivity. The factor $|\partial\rho/\partial\tau_{\rm obs}|$ controls how lifetime uncertainties propagate to charge-ratio uncertainties. In the relevant region near our central value,

$$\left|\frac{\partial\rho}{\partial\tau_{\rm obs}}\right| \approx 0.013\ \text{ns}^{-1}, \tag{25}$$

so each nanosecond of uncertainty in $\tau_{\rm obs}$ corresponds to roughly 0.013 in ρ. Equivalently, the full physical range $\rho \in [0,\infty)$ maps onto only the 172-ns interval $\tau_{\rm obs} \in [\tau^-, \tau^+] = [2025, 2197]$ ns. Our 32.6-ns lifetime uncertainty is roughly 19% of this full sweep, propagating to a charge-ratio uncertainty of order unity. Achieving $\sigma_\rho \sim 0.1$ would require $\sigma_{\tau_{\rm obs}} \lesssim 8$ ns, well below the precision attainable with our acquisition statistics.

d. Extracting the free μ^+ lifetime. Because ρ is poorly constrained by $\tau_{\rm obs}$ alone, we instead approach the problem from the other direction: we assume a value of ρ and a known τ^-, and solve Eq. (21) for τ^+. Setting $\rho = 1$ as a first estimate and using $\tau^- = \tau_c = 2025.3$ ns, Eq. (21) collapses to

$$\tau_{\rm obs} = \frac{2\,\tau^+\tau^-}{\tau^+ + \tau^-} \quad\Longrightarrow\quad \tau^+ = \frac{\tau_{\rm obs}\,\tau^-}{2\tau^- - \tau_{\rm obs}}. \tag{26}$$

Treating τ^- as exact (its literature uncertainty is much smaller than $\sigma_{\tau_{\rm obs}}$), the only stochastic input is $\tau_{\rm obs}$. Differentiating Eq. (26),

$$\frac{\partial\tau^+}{\partial\tau_{\rm obs}} = \frac{2\,(\tau^-)^2}{(2\tau^- - \tau_{\rm obs})^2}, \tag{27}$$

so

$$\sigma_{\tau^+} = \frac{2\,(\tau^-)^2}{(2\tau^- - \tau_{\rm obs})^2}\,\sigma_{\tau_{\rm obs}}. \tag{28}$$

Numerically, $|\partial\tau^+/\partial\tau_{\rm obs}| \approx 2.1$ at our central value, so the lifetime uncertainty is amplified by a factor of ~ 2 in this inversion. Substituting our values,

$$\tau^+ = 2131.5 \pm 68.6\ \text{ns}, \tag{29}$$

which lies 0.95σ below the PDG free μ^+ lifetime of 2196.98 ns [2] and is fully consistent with it.

e. Updated G_F. Inserting Eq. (29) into the Fermi-coupling relation Eq. (13) gives

$$G_F = (1.182 \pm 0.019) \times 10^{-5}\ \text{GeV}^{-2}, \tag{30}$$

which sits 0.81σ above the PDG value $G_F^{\rm PDG} = 1.1663787(6) \times 10^{-5}$ GeV^{-2} and is fully consistent within statistical uncertainty. The improvement over the naive single-exponential extraction (which gave 3.3σ tension) reflects the fact that we have now explicitly corrected for the μ^--capture systematic, at the cost of doubling the lifetime uncertainty in the inversion.

IV. DISCUSSION

1. Statistical Estimators

A statistical estimator can be characterized by two independent properties. Its *bias* is whether the estimate returns the true parameter value on average across many repeated experiments—an unbiased estimator's expected value equals the true parameter. Its *variance* is how much the estimate fluctuates from one experiment to another. The Cramér-Rao inequality places a theoretical lower bound on the variance achievable by any unbiased estimator, and the maximum-likelihood estimator (MLE) is particularly important because it saturates this bound asymptotically [8]. When an exact likelihood is available, MLE is therefore the optimal choice in principle.

a. The conventional default. The standard approach to histogram fitting is to minimize a chi-squared statistic, which approximates each bin's count n_i as Gaussian-distributed around the model prediction f_i. The justification is the central limit theorem: for sufficiently large bin counts, the Poisson distribution governing event counts is well-approximated by a Gaussian with variance equal to the mean. In this regime, χ^2 is essentially equivalent to MLE and gives unbiased results.

b. Where the approximation fails. Our histograms violate the high-count assumption in the exponential tail. With 70-ns bins and a few thousand events spread across the 14.7-μs fit window, the late-time bins typically contain 0 or 1 event. The Gaussian approximation breaks down in a quantifiable way: a Poisson distribution with mean $\langle n \rangle = 0.5$ yields zero counts with probability

$$P_{\text{Poiss}}(n = 0 \,|\, \langle n \rangle = 0.5) = e^{-0.5} \approx 0.61, \tag{31}$$

while the equivalent Gaussian, $\mathcal{N}(\mu = 0.5, \sigma^2 = 0.5)$, places

$$P_{\text{Gauss}}(x < 0) = \Phi\left(\frac{0 - 0.5}{\sqrt{0.5}}\right) \approx 0.24 \tag{32}$$

of its probability mass on negative counts—a region that is physically impossible. The approximation is not just imprecise here; it is qualitatively wrong about the shape of the count distribution.

c. Neyman χ^2 and the bias mechanism. Two slightly different χ^2 formulations exist for counting data, distinguished by how they assign per-bin uncertainty. The standard (Neyman) form uses the observed count, $\sigma_i = \sqrt{n_i}$:

$$\chi^2_{\text{N}} = \sum_i \frac{(n_i - f_i)^2}{n_i}. \tag{33}$$

The effective weight each bin contributes to the fit is $1/\sigma_i^2 = 1/n_i$, so a bin containing one event is weighted ten times more heavily than a bin containing ten. Low-count bins therefore dominate the optimization despite carrying the least statistical information, and the fitted decay constant is systematically biased low [3].

A common alternative procedure estimates the background level by fitting a horizontal line to the late-time tail of the data, subtracts this background from all bins, and then fits a pure exponential to the residuals—often iteratively. Despite the differently appearing framing, both stages of this procedure use Neyman-style weighting: the tail-fit is a $1/\sigma_i^2$-weighted mean of low-count bins, and the subsequent exponential fit applies $\sigma = \sqrt{n}$ to the subtracted distribution. The bias is inherited at both stages, and the iteration converges to a fixed point that is not the true parameter value.

d. Diagnostic: bin-width dependence. A practical diagnostic for estimator bias is the dependence of the fitted parameter on bin width. An unbiased estimator's central value should be independent of how the data is binned, since binning is a representation choice rather than a feature of the underlying distribution. We tested this by repeating the joint fit at bin widths from 20 to 200 ns. The Neyman result drifts monotonically from $\tau \approx 1759$ ns at 20 ns bins to $\tau \approx 1970$ ns at 200 ns bins—a 200 ns shift well outside the statistical error of any individual fit—while the binned Poisson MLE remains flat at $\tau \approx 2077 \pm 5$ ns across the entire range. This drift is the direct signature of the weight-asymmetry bias, gradually washing out as bins grow large enough to satisfy the central limit theorem.

e. Pearson χ^2 as the first fix. A simple modification removes the weight-asymmetry bias: use the model prediction in place of the observed count in the denominator, $\sigma_i = \sqrt{f_i}$. This is the Pearson form,

$$\chi^2_{\text{P}} = \sum_i \frac{(n_i - f_i)^2}{f_i}. \tag{34}$$

Each bin now contributes uncertainty determined by what the model predicts the count should be, rather than by the random outcome of a single realization. The bias is removed, though the underlying Gaussian approximation remains.

f. Binned Poisson MLE. The fully rigorous approach is to drop the Gaussian approximation entirely and treat each bin's content as the Poisson-distributed count it actually is. The likelihood for observing n_i counts in bin i given model expectation f_i is

$$\mathcal{L}_i = \frac{f_i^{n_i}\, e^{-f_i}}{n_i!}, \tag{35}$$

and the negative log-likelihood, after dropping a parameter-independent constant, gives the modified Cash statistic [4, 5]:

$$C = 2 \sum_i \left[f_i - n_i + n_i \ln(n_i / f_i) \right], \tag{36}$$

with the convention $n_i \ln(n_i/f_i) \to 0$ for empty bins. Unlike the Gaussian-approximation methods, this estimator is exact for counting data regardless of bin occupancy: empty bins are handled natively through the Poisson likelihood rather than through an approximate variance. The added data-dependent term ensures that C is asymptotically distributed as χ^2 with $N_{\text{bins}} - N_{\text{params}}$ degrees of freedom at the minimum, so the conventional interpretation of $C/\text{dof} \approx 1$ as a good fit retains rigorous meaning by Wilks' theorem.

g. Comparison. Table II summarizes the four estimators applied to our 5-run joint fit at 70 ns bins over $100 < t < 14800$ ns. The two biased methods (the iterative procedure and Neyman χ^2) cluster ∼200 ns below the two unbiased ones (Pearson and Poisson MLE)—the split itself is diagnostic of the underlying bias mechanism. We adopt the Poisson MLE result as our central value: $\tau = 2077.3 \pm 32.6$ ns.

TABLE II. Joint-fit lifetime estimates from four estimators applied to the five-run dataset at 70 ns bins over $100 < t < 14800$ ns. Biased estimators return systematically lower values; unbiased estimators agree within statistical errors.

Estimator	τ (ns)	Notes
Iterative subtraction (Neyman)	1871 ± 24	biased
Neyman χ^2	1895 ± 30	biased
Pearson χ^2	2100 ± 33	unbiased
Poisson MLE (Cash)	$\mathbf{2077 \pm 33}$	**unbiased, exact**

V. CONCLUSION

We have measured the mean lifetime of cosmic-ray muons stopped in a plastic scintillator using time-tagged single-PMT coincidence techniques over five multi-day acquisition runs. A joint binned Poisson maximum-likelihood fit to the five decay-time histograms, with a shared lifetime parameter and per-run amplitudes and backgrounds, yields $\tau_{\rm obs} = 2077.3 \pm 32.6$ ns—a 1.6% statistical precision on the effective single-exponential lifetime of the charge-mixed sample. Inserted into the leading-order weak-decay relation, this gives a Fermi coupling constant $G_F = (1.197 \pm 0.009) \times 10^{-5}$ GeV^{-2}, which sits 3.3σ above the PDG value—a tension consistent in direction and magnitude with the bias expected from the unresolved μ^- nuclear-capture systematic. Disentangling the two charge species under the assumption of equal stopping fractions and a literature value for the μ^--in-carbon lifetime, we extract the free positive-muon lifetime $\tau^+ = 2131.5 \pm 68.6$ ns and a corrected $G_F = (1.182 \pm 0.019) \times 10^{-5}$ GeV^{-2}, both consistent with their accepted values within 1σ.

Beyond the lifetime measurement itself, we have argued at some length for the importance of estimator choice in low-count histogram fitting. The conventional Neyman χ^2 approach, and the iterative background-subtraction procedure built on it, are biased toward systematically lower lifetimes by their over-weighting of low-count bins; on our data this bias amounts to a $\sim$10% shift, well outside the statistical uncertainty. Pearson χ^2 and the binned Poisson MLE are unbiased and agree with each other within errors, with the Poisson likelihood being the rigorous choice for counting data of any bin occupancy. The resulting bin-width independence of the Poisson MLE central value, contrasted with the monotonic drift exhibited by Neyman χ^2, serves as both a diagnostic and a vindication of the methodology adopted here.

[1] P. K. F. Grieder, *Cosmic Rays at Earth: Researcher's Reference Manual and Data Book* (Elsevier, 2001).

[2] Particle Data Group, S. Navas, *et al.*, Review of particle physics, Phys. Rev. D **110**, 030001 (2024).

[3] P. J. Humphrey, W. Liu, and D. A. Buote, χ^2 and Poissonian data: Biases even in the high-count regime and how to avoid them, Astrophysical Journal **693**, 822 (2009).

[4] W. Cash, Parameter estimation in astronomy through application of the likelihood ratio, Astrophysical Journal **228**, 939 (1979).

[5] J. S. Kaastra, On the use of C-stat in testing models for x-ray spectra, Astronomy & Astrophysics **605**, A51 (2017).

[6] T. van Ritbergen and R. G. Stuart, Complete two-loop quantum electrodynamic contributions to the muon lifetime in the Fermi model, Phys. Rev. Lett. **82**, 488 (1999).

[7] T. Suzuki, D. F. Measday, and J. P. Roalsvig, Total nuclear capture rates for negative muons, Phys. Rev. C **35**, 2212 (1987).

[8] G. Cowan, *Statistical Data Analysis* (Oxford University Press, 1998).