Measuring the altitude and latitude dependence of muon counts on a plane

Lily Sun*
MIT Department of Physics
(Dated: May 13, 2026)

We measure cosmic-ray muon counts on flights from Boston to Miami to Rio de Janeiro using a handheld CosmicWatch detector in coincidence mode. Combining CosmicWatch data with flight logs, we resolve both the altitude and geomagnetic latitude dependence of the count rate. The rate increases sharply with altitude and is well described by a log-quadratic model, while the latitude dependence shows a clear suppression near the geomagnetic equator due to magnetic shielding of primary cosmic rays. A decomposed model $R(h, \phi) = R(h)r(\phi)$ captures the main trends, though residual systematics show that altitude and latitude effects are not fully independent.

I. INTRODUCTION

When cosmic rays hit Earth's atmosphere, they produce showers of high-energy secondary particles including muons. Muons were first identified in cosmic-ray experiments by Neddermeyer and Anderson [1] as charged particles less massive than protons but more penetrating than electrons. Because they are produced high in the atmosphere but can survive to reach the ground, these muons provide a probe of both cosmic-ray air showers and relativistic propagation through the atmosphere.

In this experiment, we use the handheld CosmicWatch detector [2] to measure muon counts during flights from Boston to Miami to Rio de Janeiro. The flight environment provides a wide range of altitudes and geomagnetic locations, allowing us to observe how the detected count rate changes along the route. By combining detector timestamps with flight-tracking data, we extract empirical models for both effects.

II. BACKGROUND

Primary cosmic rays produce muons in the atmosphere. Primary cosmic rays are mostly protons (74%) and helium nuclei (18%) from space. They hit the Earth's atmosphere and, via a cascade of reactions and intermediates (e.g. pions, kaons), eventually produce a shower of muons in the upper atmosphere (at ≈15 km). Muons are highly penetrating: as leptons they do not experience the strong force, and their large mass suppresses bremsstrahlung relative to electrons. Muons with enough energy $\gtrsim 2.4$ GeV are sufficiently time-dilated to reach the surface before decaying.

Muon rates depend on altitude and latitude. We expect muon counts to increase with altitude, as muons have to travel a shorter distance, meaning fewer of them lose energy or decay before detection. Muon counts also depend on latitude because *primary* counts (which produce muons) depend on latitude—primary cosmic rays are deflected by Earth's magnetic field which is strongest at the geomagnetic equator.

CosmicWatch detects muons via coincidence events. The CosmicWatch contains a plastic scintillator, which is excited by charged particles and releases light that is multiplied through a silicon photomultiplier and signal chain and recorded as an event. This process detects any charged particle; however, when we connect two scintillators together in "coincidence mode", we can find events that occur in both scintillators within $\approx 3\,\mu$s (coincidence events) which are highly likely to be muons as other particles do not have the energy to trigger both. We thus use this "coincidence mode" and arrange the scintillators close together like in Figure 1 to maximize surface area for coincidence events.

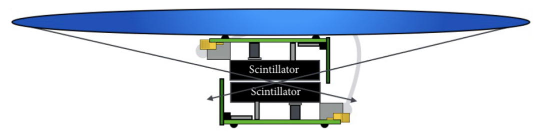

FIG. 1. CosmicWatch setup to maximize muon counts.

Poisson statistics. Let the CosmicWatch record N events over some time period Δt. Muon counts are a Poisson process, thus, the error in counts is $\sqrt{N}$ and the fractional error is $1/\sqrt{N}$. To obtain a rate, we correct for the deadtime t_{dead} of the detector (where the detector is inactive for a short time after each event) and obtain $R = \frac{N}{\Delta t - t_{\text{dead}}}$ with the corresponding fractional error $1/\sqrt{N}$.

III. RESULTS AND ANALYSIS

Experimental procedure. We bring the CosmicWatch on the following flights: 1. BOS-MIA (AA1397, April 22), 2. MIA-GIG (AA905, April 22) and 3. GIG-MIA (AA992, April 27).[1] We extract flight data (altitude h, geographic latitude λ and longitude ψ)

[1] Unfortunately, the CosmicWatch stopped working during the layover before the last MIA-BOS flight.

* xqsun@mit.edu

from FlightAware. The CosmicWatch was started in coincidence mode before boarding each flight, and the real time was recorded simultaneously on a phone as the CosmicWatch's clocks are only relative to its start time. The CosmicWatch was powered with a powerbank throughout the flight and placed upright on the plane floor.

Figure 2 shows the count rate (blue) in bins of $\Delta t = 50s$ (chosen for reasonable uncertainty) and altitude (orange) for each flight.[2] We observe that count rates rise sharply during takeoff and drop during landing, confirming the altitude effect. During the cruise of Flight 2 and 3, we see a gentle dip in rates even as altitude is constant or stepping up, suggesting a latitude effect as the MIA-GIG path crosses a wide range of latitudes. We aim to find $R(h, \phi)$ where h is altitude and ϕ is geomagnetic latitude. We model this as $R(h, \phi) = R(h) \cdot r(\phi)$ so we can investigate each dependence separately (we discuss joint effects in III.3).

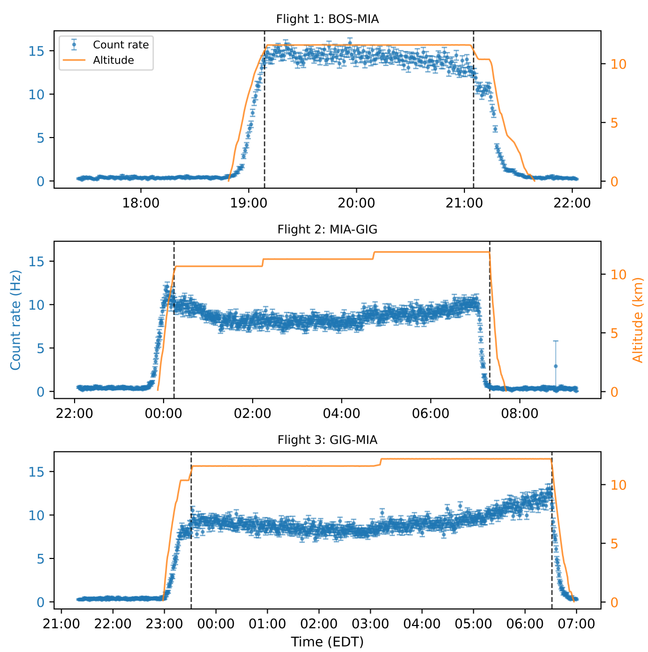

FIG. 2. Count rates and altitude against time for all flights. Clear altitude and latitude dependences are visible.

III.1. Altitude Dependence $R(h)$

To obtain $R(h)$ we use only data during takeoff and landing[3] phases as that is when altitude varies significantly. To obtain $R(h)$ we bin the data in altitude bins of width $\Delta h = 0.5$ km rather than time bins. That is, we count how many events $N(h)$ occurred and the total live time $T(h)$ while the plane is in each altitude bin, then $R(h) = N(h)/T(h)$.

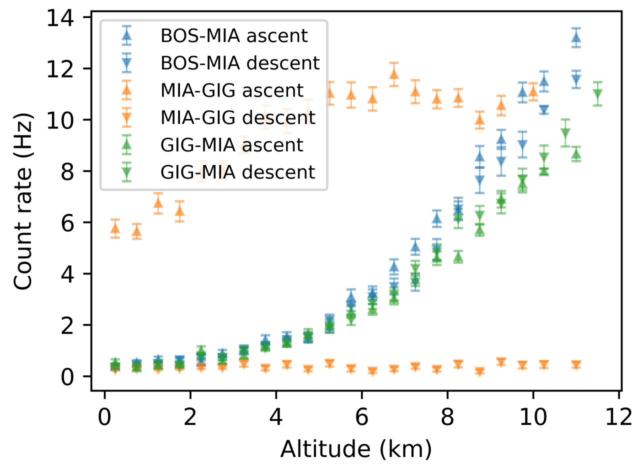

FIG. 3. Altitude dependence $R(h)$ derived from each flight's takeoff and landing independently. Flight 2's (MIA-GIG) CosmicWatch data is offset in time causing it to not match the right altitudes and must be corrected.

Correction for Flight 2. We find $R(h)$ for each flight's takeoff and landing independently (Figure 3) and notice that Flight 2 MIA-GIG is anomalous, where the count rates are much higher than expected during takeoff and lower during landing. This suggests that the start time of CosmicWatch data was recorded to be *earlier* than it actually was due to human error. Therefore, we must apply a correction τ to all times in Flight 2's CosmicWatch data by matching it with Flight 3 assuming the latter is trusted. We shift the Flight 2 data by each trial τ to obtain $R_{2,h}(h)$ and minimize the chi-squared distance $\chi^2(\tau) = \sum_h \frac{[R_{2,\tau}(h)-R_3(h)]^2}{\sigma^2_{2,\tau}(h)+\sigma^2_3(h)}$. This pins Flight 2's data to Flight 3's data, meaning they are no longer fully independent and any errors between them could be underestimated, but is necessary to be able to use the Flight 2 data. We can do this by matching only each of MIA and GIG's $R(h)$, using both cities' $R(h)$ or pooling all data to obtain one $R(h)$ per flight and matching that. All give similar offsets (Figure 4) but we chose to match both cities' $R(h)$ and sum their χ^2 to minimize as that is the most principled (we assume $R(h)$ depends on city).

Now that the data is fixed, we can get a $R(h)$ per city by pooling all data at the same city when finding $R(h) = N(h)/T(h)$ i.e. $R_{\mathrm{BOS}}(h)$ from Flight 1's takeoff, $R_{\mathrm{MIA}}(h)$ from Flight 2's takeoff and Flight 3's landing, and $R_{\mathrm{GIG}}(h)$ from Flight 2's landing and Flight

[2] Note that the two detectors record a slightly different number of coincidence events (maximum difference 6/235,732 or 0.002% on Flight 3) and their clocks drift slightly (1.3 s/11.2 h or 0.003% on Flight 2) so we safely pick an arbitrary detector for all analyses.

[3] Takeoff: from start of altitude tracking until the plane reaches 95% of its cruise altitude 1/4 into the flight. Landing: from when the plane reaches 95% of its maximum altitude until end of altitude tracking.

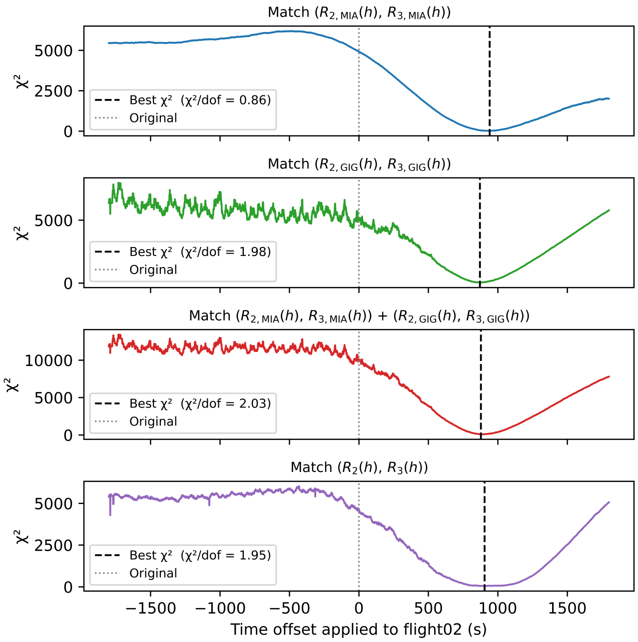

FIG. 4. Fixing Flight 2's offset by minimizing the χ^2 of its $R(h)$ with Flight 3. We match both MIA and GIG's $R(h)$ (third panel).

3's takeoff. This assumes $R(h)$ differs by city but takeoff/landing does not matter. We indeed see a difference per city caused by latitude differences (Figure 5). We fit both a log-linear $\log R = a + bh$ and log-quadratic $\log R = a + bh + ch^2$ model to data, finding that the log-quadratic model fits better particularly at high h.

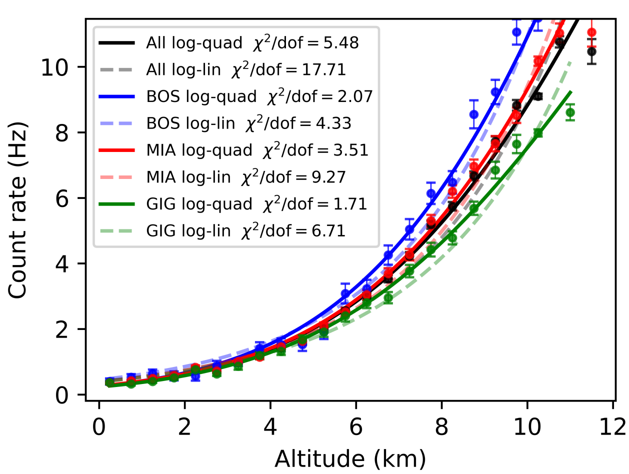

FIG. 5. $R(h)$ data and fitted log-linear/log-quadratic models, for each city pooled, and for all data pooled. There is a dependence of $R(h)$ on city, but we use the pooled $R(h)$ (black) for our decomposed model.

Our model $R(h,\phi) = R(h) \cdot r(\phi)$ requires one $R(h)$ independent of city/latitude. To do that, we pool *all* data when finding $R(h) = N(h)/T(h)$ and fit $\log R = (-0.92 \pm 0.10) + (0.31 \pm 0.01)h + (-0.012 \pm 0.004)h^2$ (note the negative coefficient on h^2). The city-dependent curves and worse fit for pooled data in Figure 5 show that this pooling introduces systematic error, discussed later.

III.2. Latitude Dependence $r(\phi)$

To isolate the latitude modulation, we use cruise data and divide the observed rate by the fitted altitude model. In each geomagnetic-latitude bin $\Delta\phi = 0.2°$, $r(\phi) = \frac{N(\phi)/T(\phi)}{\hat{R}(\bar{h}(\phi))}$ where $\bar{h}(\phi)$ is the mean altitude in the latitude bin and ϕ is computed from geographic coordinates using `apexpy` [3].

We do this separately per flight, obtaining Figure 6. The error bars include both the statistical Poisson error in $N(\phi)$, and also any error from the $R(h)$ fit which is a mix of statistical errors from $N(h)$ and systematic error (due to the assumption of latitude-independent $R(h)$). There is a systematic deviation at $\phi \approx 30°$ (near MIA) between Flight 2 and 3 data. We think this is because Flight 2 was at 10.7 km (soon after takeoff) while Flight 3 was at 12.2 km (shortly before landing) near MIA—since the $R(h)$ equation was fitted on data < 12.2 km, it was extrapolated for Flight 3's cruise and there may be a systematic error in $\hat{R}(h)$ leading to a systematic offset.

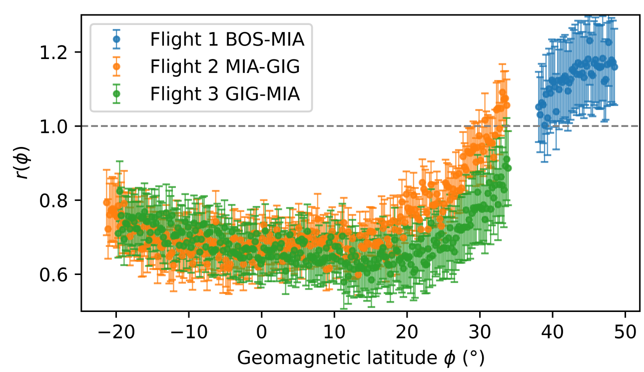

FIG. 6. Geomagnetic latitude dependence $r(\phi)$ for each flight. Counts are lowest near geomagnetic equator.

We see that $r(\phi)$ is lowest near the equator (skewed slightly North) and rises towards higher ϕ. The equatorial dip is due to deflection of *primary* cosmic rays by Earth's magnetic field being stronger near the equator, and fewer primaries produce fewer muons. To model this, we consider two approaches.

First approach: Power law in rigidity. A primary cosmic ray of momentum p and charge q has rigidity $\mathcal{R} = \frac{pc}{q}$ which quantifies its resistance to deflection by magnetic fields. Störmer [4] assumed Earth's magnetic field is a dipole and derived that only primaries with $\mathcal{R} > \mathcal{R}_c$ can pass through, and this cutoff depends on geomagnetic latitude as $\mathcal{R}_c = 14.9\cos^4\phi$ GV.

To find the total flux Φ of primary rays we then integrate their spectrum $\frac{\mathrm{d}\Phi}{\mathrm{d}\mathcal{R}}$ above $\mathcal{R}_c$. A common form for the spectrum is $\frac{\mathrm{d}\Phi}{\mathrm{d}\mathcal{R}} \propto \mathcal{R}^{-\gamma}$ with $\gamma \approx 2.7$ [5]. Therefore,

$$\Phi = \int_{\mathcal{R}_c}^{\infty} \frac{\mathrm{d}\Phi}{\mathrm{d}\mathcal{R}} \mathrm{d}\mathcal{R} \propto \mathcal{R}_c^{-(\gamma-1)} \propto \cos^{-4(\gamma-1)} \phi$$

We first assume that the rate of muons detected is proportional to the flux of primaries $r \propto \Phi \propto \cos^{-4(\gamma-1)} \phi = \cos^{-6.8} \phi$. Thus, we fit to our data a model $r = A \cos^n \phi$ ($\log r = \log A + n \log \cos \phi$), obtaining Figure 7. We find $n = -1.80$ much smaller than the expected -6.8, implying that the drop in rates towards the equator is much gentler than theory suggests.

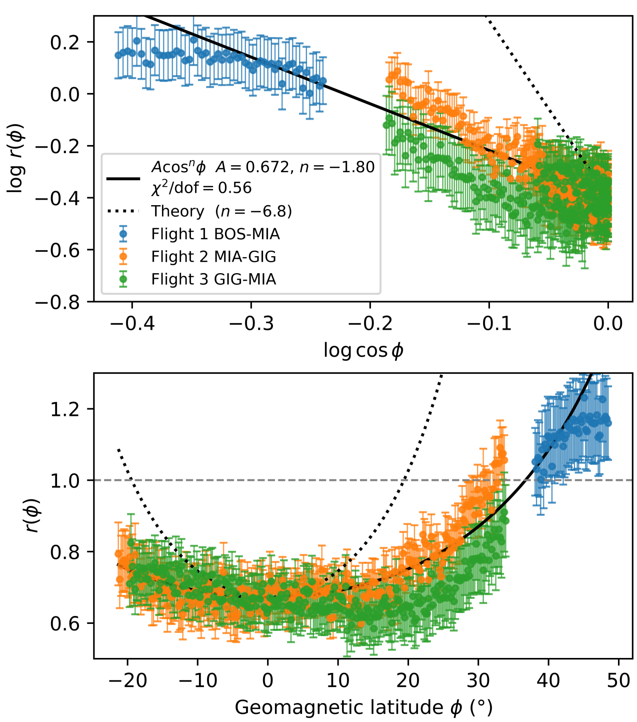

FIG. 7. Fitted power-law-based latitude dependence (top: log scale). We obtain a much smaller n than expected, and the model shows a systematic deviation from data.

This is because the rate of muons detected is not just proportional to the flux of primaries, but higher rigidity (energy) primaries produce higher energy muons and thus more detected muons. Instead, if we assume a power-law yield of detected muons per primary of $Y(\mathcal{R}) \propto \mathcal{R}^{\alpha}$, the integral gives $r \propto \cos^{-4(\gamma-\alpha-1)} \phi$. Our fit implies $\alpha \approx 1.25$ so higher rigidity primaries do indeed produce more muons.

Second approach: Empirical shape of $r(\mathcal{R}_c)$. However, the systematic deviation of data from even the fitted model in Figure 7 suggests that a power law in rigidity may not be the right relationship. We plot $r(\mathcal{R}_c)$ directly via Stormer's transformation $\mathcal{R}_c = 14.9 \cos^4 \phi$ GV in Figure 8 and find that while the data follows a plausibly exponential or power law for low rigidities,

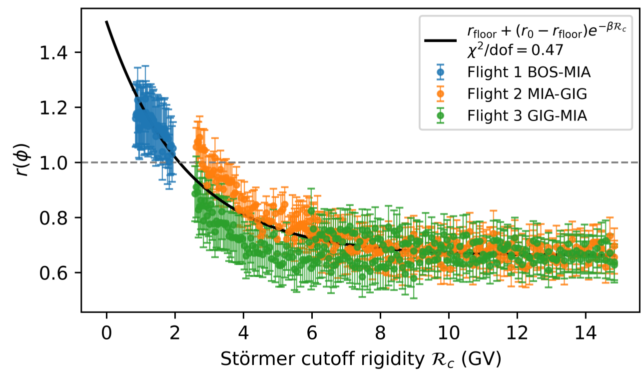

FIG. 8. Empirical shape of $r(\mathcal{R}_c)$ shows a non-zero floor of muons even at high cutoff rigidities.

for high rigidities $\gtrsim 8$ GV, the muon rate plateaus. We think this can be explained as follows: the muons that make it to and are detected by the CosmicWatch are mostly muons from high rigidity primaries $\mathcal{R} \gg \mathcal{R}_c$, and these primaries always pass through the magnetic field even at high $\mathcal{R}_c$ (low ϕ) and produce a baseline rate of muons that account for $\approx 66\%$ of detected muons. This, combined with the previous fact that higher rigidity primaries produce more muons, results in a gentler drop and plateau in muon count near the equator. We thus fit a relationship $r(\mathcal{R}_c) = r_{\text{floor}} + (r_0 - r_{\text{floor}})e^{-\beta\mathcal{R}_c}$ to get $r_{\text{floor}} = 0.661 \pm 0.005$, $r_0 = 1.51 \pm 0.04$ and $\beta = 0.43 \pm 0.02$.

III.3. Putting It Together

We modeled count rates as having independent altitude and latitude dependences $R(h, \phi) = R(h) \cdot r(\phi)$. From Section III.1 we fitted $\log R(h) = a + bh + ch^2$ and from section III.2 we found $r(\phi) = r_{\text{floor}} + (r_0 - r_{\text{floor}})e^{-\beta\mathcal{R}_c}$. Putting these together, we can see how well our model fits our observed $R(t)$ throughout all 3 flights. To the overall $R(t)$ data we also fit the Axani [6] model $R = N \exp\left[\alpha \left(\sin(\phi+\theta)^2 + \beta\right) h\right] - b$. We find that our method has a lower χ^2 despite the altitude and latitude components being fitted independently, implying that this $R(h, \phi) = R(h) r(\phi)$ decomposition and our models for each $R(h)$ and $r(\phi)$ are empirically valid.

In Section III.2 our χ^2 values are < 1 implying some overfitting, likely because of the large errors in $\hat{R}(h)$ propagated which are systematic in nature rather than purely statistical errors (removing $R(h)$ error leads to larger $\chi^2 > 1$). The largest systematic error comes from the assumption of constant $R(h)$ in obtaining $r(\phi)$. A more expressive model would couple altitude and latitude dependence (such as the interaction term $\alpha \sin(\phi+\theta)^2 h$ term in Axani [6]'s model), however we chose the $R(h, \phi) = R(h) \cdot r(\phi)$ to be able to analyze the effects independently and find that this is a good enough approximation.

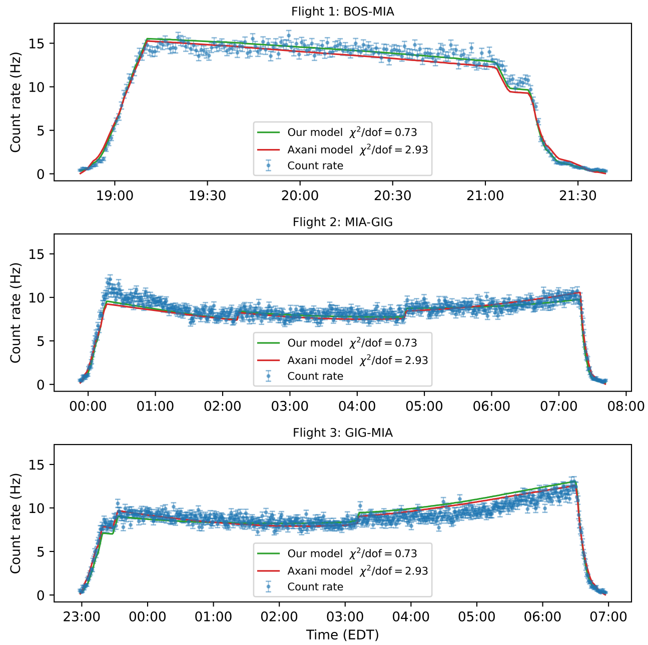

FIG. 9. Our model vs. Axani [6] model compared to our $R(t)$ data ($\Delta t = 50$ s). Our model gives a good prediction despite $R(h)$ and $r(\phi)$ being fitted individually.

IV. CONCLUSION

We find that a simple decomposed model $R(h, \phi) = R(h) \cdot r(\phi)$ captures the observed muon count rate across all three flights, even though the altitude and latitude terms are fitted independently. This shows that the dominant atmospheric and geomagnetic effects can be described separately in this dataset. The latitude modulation is qualitatively consistent with geomagnetic rigidity cutoff physics: rates decrease near the geomagnetic equator, where low-rigidity primaries are more strongly excluded. A naive Størmer-cutoff power law predicts this trend but gives too steep a drop, while an empirical cutoff-rigidity model with a nonzero floor better matches the data. Overall, these flight measurements show that a compact scintillator detector can recover both the altitude dependence of atmospheric muons and the geomagnetic suppression of their primary cosmic rays.

ACKNOWLEDGMENTS

We thank Liam Sheldon for helping set up the CosmicWatch detector in the given configuration, and the 8.13 TAs and instructors for their guidance.

[1] S. H. Neddermeyer and C. D. Anderson, Phys. Rev. **51**, 884 (1937).
[2] S. Axani, K. Frankiewicz, and J. Conrad, Journal of Instrumentation **13**, P03019–P03019 (2018).
[3] C. van der Meeren, K. M. Laundal, A. G. Burrell, L. Lamarche, G. Starr, A. Reimer, A. Morschhauser, and I. Michaelis, "Apexpy vx.y.z," (2024).
[4] C. Störmer, *The Polar Aurora* (Oxford University Press, 1955).
[5] S. Navas *et al.* (Particle Data Group), Physical Review D **110**, 030001 (2024).
[6] S. N. Axani, "The physics behind the cosmicwatch desktop muon detectors," (2019), arXiv:1908.00146 [physics.ins-det].