Probing Moving Gamma-Ray Source Populations through Spatiotemporal Clustering of Fermi-LAT Photons

Student: Dou Lai
Walter Payton College Preparatory High School, Chicago, IL, USA

Instructor: James Unwin
Department of Physics, University of Illinois Chicago
Chicago, IL, USA

September 14, 2026

Probing Moving Gamma-Ray Source Populations through Spatiotemporal Clustering of Fermi-LAT Photons

Dou Lai

ABSTRACT

Nearby compact gamma-ray emitters can leave their accumulated photons distributed along trajectories that traditional stationary-source searches miss. We conduct a blind, full-sky photon-level search of ten years of 1–100 GeV Fermi-LAT data and propose a novel algorithm that clusters photons within annual time bins and uses dynamic programming to link these clusters. Our evaluation of moving-source candidates produced by this algorithm finds no significant moving-source excess relative to identically analyzed time-scrambled skies, with an empirical upper-tail probability of 0.066. Using artificial source injections to measure detection efficiency, we derive an upper limit of approximately 122 steady sources over the full sky at 90% confidence, at a photon flux of $7.08 \times 10^{-10}\,\mathrm{ph\,cm^{-2}\,s^{-1}}$ and an angular speed of $0.5^\circ\,\mathrm{yr}^{-1}$. We also convert this source-recovery probability into effective volumes and local abundance limits, illustrated with dressed and quasi-steady evaporating primordial black holes (PBHs). Under the adopted dressed-PBH model with a surviving annihilation halo and the E^{-2} recovery response, the central model constrains PBHs of mass $5.79 \times 10^{-10} M_\odot$ to contribute at most 0.849% of the local dark matter, at 90% confidence.

Keywords: Fermi-LAT, moving point sources, proper motion, primordial black holes, gamma-ray astronomy

ACKNOWLEDGMENTS

I am extremely grateful to Professor James Unwin for his outstanding mentorship throughout this research project. His deep physical expertise was invaluable to this project. I also thank the PRIMES program for the opportunity to undertake this research.

I used OpenAI models as tools to assist in the preparation and revision of this paper. Model sessions involved:

- Generating Figure 1 used in this paper (Sep 2026, one session, GPT-6 Astra).
- Improving clarity, grammar, and spelling (Sep 2026, two sessions, GPT-5.6 Sol).
- Debugging a stationary source detection prototype and an early version of a moving-source candidate-finding pipeline (May 2026, one session, GPT-5.5).
- Debugging and reviewing an early statistical validation implementation for the moving-source pipeline (June 2026, one session, GPT-5.5).

Commitments on Academic Honesty and Integrity

We hereby declare that we

1. are fully committed to the principle of honesty, integrity and fair play throughout the competition.
2. actually perform the research work ourselves and thus truly understand the content of the work.
3. observe the common standard of academic integrity adopted by most journals and degree theses.
4. have declared all the assistance and contribution we have received from any personnel, agency, institution, etc. for the research work.
5. undertake to avoid getting in touch with assessment panel members in a way that may lead to direct or indirect conflict of interest.
6. undertake to avoid any interaction with assessment panel members that would undermine the neutrality of the panel member and fairness of the assessment process.
7. observe the safety regulations of the laboratory(ies) where we conduct the experiment(s), if applicable.
8. observe all rules and regulations of the competition.
9. agree that the decision of YHSA is final in all matters related to the competition.

We understand and agree that failure to honour the above commitments may lead to disqualification from the competition and/or removal of reward, if applicable; that any unethical deeds, if found, will be disclosed to the school principal of team member(s) and relevant parties if deemed necessary; and that the decision of YHSA is final and no appeal will be accepted.

(Signatures of full team below)

X *Dou Lai*

X *James Unwin* Prof. James Unwin, University of Illinois, Mentor.

Date: 13 September 2026

Contents

1. Introduction 5
   1.1. Background and Motivation 5
   1.2. Previous Searches and Their Limitations 5
   1.3. Our Approach and Main Findings 6
2. Flux and Proper Motion of Nearby Sources 6
3. Search Algorithm for Finding Moving-Source Candidates 8
   3.1. Overview 8
   3.2. Sky Division 9
   3.3. Cluster Criteria 9
   3.4. Cluster Linking 10
   3.5. Removing Duplicates 12
4. Calibration and Validation 12
   4.1. Estimating False-Positive Detections with Time Scrambling 12
   4.2. Statistical Validation of the Search Results 13
   4.3. Detection Efficiency of the Moving-Source Algorithm 15
   4.4. Detection Efficiency as a Function of Source Flux 18
   4.5. Injection Validation and Calibration Boundaries 19
5. Results 19
   5.1. Upper Limit on Detected Tracks 20
   5.2. Limits on the Full-Sky Source Population 21
      5.2.1. Effective Survey Volume 22
      5.2.2. Local Number Density and Dark Matter Abundance 23
   5.3. Primordial Black Hole Benchmarks 24
      5.3.1. Dressed Primordial Black Holes 24
      5.3.2. Quasi-Steady Evaporating Primordial Black Holes 25
6. Discussion 27
   6.1. Background-Corrected Source Limits 27
   6.2. Effects of Source Luminosity and Motion 27
7. Future Work 28
8. Conclusion 29
A. Highest-Ranked Candidates 30
B. Derivations of Source Motion, Detection Efficiency, and Population Limits 30
   B.1. Effects of Motion on Source Shape and Trajectory 30
   B.2. Poisson Count Fluctuations and Detection Efficiency 31
   B.3. Existence and Uniqueness of the Signal Upper Limit 31
   B.4. Source Velocities and Effective Survey Volume 32
   B.5. Annihilation Integral for the Capped Density Spike 33

1. INTRODUCTION

1.1. *Background and Motivation*

Despite its essential role in the Universe, the physical nature of dark matter remains largely unknown. Dark matter makes up roughly 85% of all matter and about 26% of the Universe's total mass-energy density [1], and its gravity has played a central role in forming galaxies, galaxy clusters, and the large-scale cosmic web. Although its presence is revealed through galaxy rotation, gravitational lensing, the cosmic microwave background, and the growth of cosmic structure [2], these observations tell us primarily what dark matter does, not what it actually is. No dark matter particle or compact-object population has yet been conclusively identified at the time of writing [3, 4]. Among many possibilities, dark matter could consist of a new elementary particle, a population of compact objects, or some mixture of the two. Determining its physical nature therefore remains fundamental to our understanding of the contents and formation of the Universe.

One dark matter scenario is particularly relevant to searches for nearby gamma-ray sources. Primordial black holes (PBHs), black holes that could have formed in the early universe, might constitute some fraction of dark matter [4]. Studying how common these objects are can therefore tell us about the nature of dark matter and about conditions in the early Universe. These objects are difficult to detect directly because dark matter does not normally produce much observable radiation. However, PBHs can, under certain models, produce gamma rays. Low-mass PBHs can emit Hawking radiation [5], while dense clumps of particle dark matter [6, 7] can produce gamma rays if the dark matter particles annihilate [8]. Heavier PBHs may also be surrounded by dense dark matter halos whose particles annihilate and emit gamma rays [2, 8, 9]. These possibilities motivate a search for compact gamma-ray sources that could reveal otherwise elusive dark matter populations.

Gamma-ray searches typically assume that sources have fixed celestial coordinates. This approximation is excellent for distant stars, galaxies, and most Galactic objects, but does not necessarily hold for a nearby compact source. A comparison to visible-spectrum astronomical measurements of Solar System objects makes this clear. A source that moves appreciably during a multiyear survey will deposit its photons along a path rather than at a single point, leading to a curve in the spatiotemporal data. Such a moving source will likely be missed or discarded by a traditional stationary point-source analysis.

The Fermi Large Area Telescope (Fermi-LAT) is particularly well suited for observing moving gamma-ray sources because it repeatedly surveys the whole sky and records crucial information such as the direction, energy, and arrival time of each detected photon [10, 11]. Its long time baseline allows a faint source's position to be compared across several years. If such a source is nearby, its apparent position can shift significantly across the sky during Fermi-LAT's observing period. That proper motion provides an additional way to distinguish nearby compact sources from distant, effectively stationary gamma-ray sources. This motivates searching not only for gamma-ray emission, but specifically for gamma-ray sources whose positions change along with time.

1.2. *Previous Searches and Their Limitations*

Previous studies have considered the possibility that nearby compact gamma-ray sources could exhibit measurable proper motion. Earlier work forecast proper-motion signatures from dark matter microhalos and developed statistical methods for detecting unresolved moving sources in diffuse gamma-ray backgrounds [12, 13]. More recently, a 2018 Fermi-LAT search for primordial black holes tested 318 preselected, unassociated high-latitude sources from the `3FGL` catalogue and found no likely PBH candidate [14]. That study placed an upper bound on the local PBH evaporation rate, although its targets were selected from sources that had already been detected as approximately

stationary.

This catalogue-based strategy leaves open an important limitation. It can test proper motion only for sources that first pass a stationary detection threshold and enter a source catalogue [15, 16]. A moving source that fails this prerequisite is absent from the candidate set, even if its photons form a coherent trajectory in the underlying data.

The unresolved problem is therefore not proper-motion fitting itself, but the absence of an all-sky search that begins directly from the photon data. In this work, we address this limitation by starting from the full Fermi-LAT photon field rather than from preselected catalogue sources. Existing catalogue sources are used only to veto known stationary emitters and to control local backgrounds, rather than to define the initial candidate list. As such, we present the first blind, full-sky search for moving gamma-ray sources at the photon level to date.

1.3. *Our Approach and Main Findings*

We reconstruct candidate trajectories by linking photon clusters across successive years, then use time-scrambled skies to assess whether the observed tracks exceed chance expectations. We also use synthetic moving-source injections to establish the search's sensitivity. Unlike the 2018 analysis, which seeded candidates from the existing Fermi-LAT `3FGL` catalogue [14], our pipeline starts from the full photon field and uses the existing `FL16Y` catalogue [17] solely to veto known stationary sources and mask local backgrounds.

We identify 577 candidate tracks, compared with a mean of 532.8 in 60 matched scrambled skies. This corresponds to a p-value of 0.066, providing no statistically significant evidence for a moving-source population. We therefore derive a 90% upper limit of 80.8 detected signal tracks, corresponding at maximum sensitivity to approximately 122 steady full-sky sources.

We use the effective survey volume to translate this null result into limits on local source density and dark matter abundance. The dressed-PBH model gives its strongest central limit of $f_{\mathrm{PBH},90} = 8.49 \times 10^{-3}$ based on the adopted annihilation-halo model and E^{-2} spectral response. For quasi-steady evaporating PBHs, the adopted Hawking-emission model and the same spectral response give their strongest central number-density limit of $n_{90} = 1.58 \times 10^{8}\,\mathrm{pc}^{-3}$. Because our response is expressed independently of a particular source model, it can be used to place limits on local source density under alternative PBH models and for other moving gamma-ray populations, such as microhalos.

The remainder of this paper is structured as follows. Section 2 develops the moving-source physical model, and Sections 3–4 describe the search methodology and calibration. Section 5 presents the results and population constraints, while Sections 6–8 discuss their implications, outline future work, and conclude.

2. FLUX AND PROPER MOTION OF NEARBY SOURCES

We consider nearby gamma-ray emitters that appear point-like to the Fermi-LAT telescope. At fixed intrinsic luminosity and transverse speed, their distance from the telescope determines both their observed flux and proper motion. These two effects jointly determine what sources the search can detect, while motion during each interval spreads their photons along the travel direction.

For isotropic photon luminosity L_γ, heliocentric distance d, and transverse speed $v_\perp$, we can compute the observed gamma-ray photon flux F_γ and the source's angular proper motion μ across the sky as

$$(F_\gamma, \mu) = \left(\frac{L_\gamma}{4\pi d^2}, \frac{v_\perp}{d} \right).$$

Because $F_\gamma \propto d^{-2}$ while $\mu \propto d^{-1}$, increasing distance suppresses the observed photon flux more

rapidly than the apparent angular motion. Distant sources can be too faint and appear nearly stationary, while sufficiently nearby sources can exceed our calibrated speed range. Together, these effects limit the search to a finite distance range. We calibrate the search over $0.25^\circ\,\mathrm{yr}^{-1} \leq \mu \leq 0.8^\circ\,\mathrm{yr}^{-1}$. Slower sources are harder to distinguish from stationary sources, while faster sources are harder to link across annual time bins. To translate this interval into an illustrative distance range, we take $238\ \mathrm{km\,s}^{-1}$, the local-standard-of-rest speed recommended for the Standard Halo Model [18], as the reference transverse speed. Converting $\mu = v_\perp/d$ to degrees per year then gives

$$\mu = 0.014^\circ\,\mathrm{yr}^{-1} \left(\frac{v_\perp}{238\,\mathrm{km\,s}^{-1}}\right)\left(\frac{d}{1\,\mathrm{pc}}\right)^{-1}.$$

At this reference speed, the proper-motion interval corresponds to a distance range of $0.017\,\mathrm{pc} \lesssim d \lesssim 0.056\,\mathrm{pc}$. For a source to lie within our calibrated sensitivity, its distance must satisfy both the proper-motion and photon-count requirements simultaneously. The motion calibration restricts the distance to

$$\frac{v_\perp}{\mu_{\max}} \leq d \leq \frac{v_\perp}{\mu_{\min}},$$

while the photon-count calibration defines a separate distance interval through the source flux. Only the overlap of these two intervals contributes to the effective survey volume, as formalized in Section 5.2.1.

Photon directions recorded by Fermi-LAT scatter about the source according to the point-spread function (PSF) [10, 11]. In local tangent-plane coordinates, we model the source trajectory as

$$\boldsymbol{\theta}_s(t) = \boldsymbol{\theta}_0 + \boldsymbol{\mu}(t - t_0). \tag{1}$$

Here, $\boldsymbol{\theta}_0$ is the source position at reference time t_0 and $\boldsymbol{\mu}$ is its constant angular-velocity vector. Over a time baseline T, this corresponds to a total angular displacement $\Delta\theta_T = |\boldsymbol{\mu}|T$. For a constant detected photon rate in an interval of duration Δt, the expected angular density is the PSF convolved with a uniform line segment:

$$\Sigma_{E,\Delta t}(\boldsymbol{\theta}) = \frac{N_\gamma}{\Delta t}\int_{-\Delta t/2}^{\Delta t/2} P_E(\boldsymbol{\theta} - \boldsymbol{\theta}_0 - \boldsymbol{\mu}\tau)\,d\tau. \tag{2}$$

If the PSF core is approximated locally by an isotropic Gaussian with variance σ_E^2, the covariance of the photon distribution becomes

$$\mathbf{C}_{E,\Delta t} = \sigma_E^2\mathbf{I} + \frac{\Delta t^2}{12}\boldsymbol{\mu}\boldsymbol{\mu}^\top. \tag{3}$$

Appendix B.1 derives Eq. 3 and the condition under which the linear angular-trajectory approximation is valid. The motion-induced term in Eq. 3 is aligned with $\boldsymbol{\mu}$, so it increases the variance only along the direction of travel, giving $\sigma_\parallel^2 = \sigma_E^2 + (|\boldsymbol{\mu}|\Delta t)^2/12$, while the perpendicular variance remains $\sigma_\perp^2 = \sigma_E^2$. Motion-induced elongation is most apparent at high energies, where the PSF is narrower. Because the motion-induced contribution to the covariance grows as Δt^2, dividing the observation period into shorter intervals reduces this longitudinal smearing. On the other hand, shorter bins would divide the sparse photon counts among more intervals, weakening cluster detection. Annual bins retain more photons per cluster while still providing ten temporal samples across the observation period. We therefore cluster in annual time bins with energy-dependent radii and link these clusters found in separate time bins across years instead of treating the decade as one stationary source.

3. SEARCH ALGORITHM FOR FINDING MOVING-SOURCE CANDIDATES

3.1. *Overview*

We divide the observation period into ten consecutive, approximately annual time bins. We then

1. Identify photon clusters by energy band. A compact moving source produces a localized, potentially slightly elongated concentration of photons within each of ten annual bins. We identify these concentrations by clustering nearby photons within a time bin, provided their local density statistically exceeds the surrounding background (Section 3.3).

2. Use dynamic programming to score and link clusters across our ten time bins according to their consistency with a moving-source trajectory. We keep tracks whose centroids form an elongated spatial distribution and progress roughly linearly, consistent with Eq. 1 within the centroid uncertainties (Section 3.4).

3. Remove duplicate candidates (Section 3.5).

The number of candidates produced by this procedure is reported in Section 5.1, and Figure 1 illustrates the above sequence.

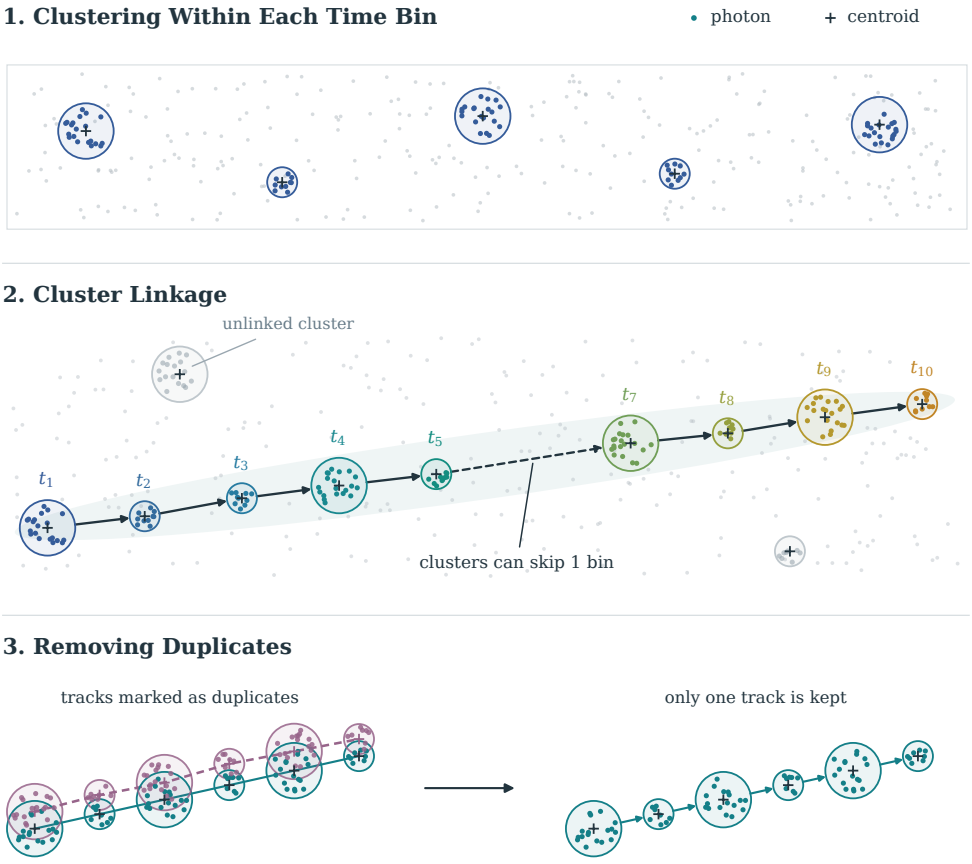


Figure 1. Visual representation of our moving-source search algorithm: we cluster photons within annual bins, then link clusters across bins using dynamic programming, and finally remove any duplicate moving-source tracks. Dots represent photons and crosses mark cluster centroids.

We select photons using the following criteria:

- **Time Range:** We require all photons to be within the time range 479174404–794707205 s of the Fermi-LAT mission elapsed time (MET) [19], which corresponds to 9 March 2016 at 00:00:00 UTC through 9 March 2026 at 00:00:00 UTC.
- **Energy Bound:** $1\,\mathrm{GeV} \leq E < 100\,\mathrm{GeV}$. The lower cut avoids the broad low-energy point-spread function (PSF) since the LAT's average 68% containment radius improves from about 5° at 100 MeV to 0.8° at 1 GeV [20]. The 100 GeV upper cut sets the calibrated upper boundary of the search as there are very few usable photons above 100 GeV.
- **Zenith Angle:** $\theta_z < 90°$. This is to suppress gamma rays from the Earth's limb [11, 20].
- **GTI:** `FPS_APPLY_GTI = true`. We only retain photons recorded during time intervals designated as "good time intervals" by Fermi-LAT [21].
- **Invalid Photons:** We also discard photons with nonfinite sky coordinates, energies, or arrival times, including values recorded as `inf` or `NaN`.

3.2. *Sky Division*

For computational purposes, we divide the sky into $N_{\mathrm{sky}} = 448$ overlapping spherical caps, each with angular radius 18°, and analyze each one separately. The centers are generated on a spherical Fibonacci lattice [22]. For $i = 0, \ldots, N_{\mathrm{sky}} - 1$, the Galactic latitude and longitude of center i are

$$(b_i, \ell_i) = \left(\arcsin\left(1 - \frac{2(i+1/2)}{N_{\mathrm{sky}}}\right), \left[i\pi(3-\sqrt{5})\right] \bmod 2\pi \right).$$

This construction is quasi-uniform because the spherical area element obeys $d\Omega = d\ell\, d(\sin b)$. The latitude prescription spaces $\sin b$ uniformly, and the longitude advances by the irrational golden angle $\pi(3-\sqrt{5})$, which avoids repeated meridians and distributes the centers quasi-uniformly in longitude.

We must guarantee that this construction covers the whole sky. To do this, we assign every sky position to its nearest center. The positions closer to one center than to any other form that center's spherical Voronoi cell, which serves as its unique ownership cell. We retain each track only in the cap whose Voronoi cell contains the track's spherical mean. Every point in that cell lies within 7.39° of its center. Across the ten annual bins, the cluster linking constraints in Section 3.4 ensure that every centroid lies within 5.13° of the track's spherical mean. The background annulus extends at most 2.94° from each centroid. Therefore, each track's source apertures and background annuli lie within $7.39 + 5.13 + 2.94 = 15.46°$ of the cap center, fully inside the 18° processing cap. We must also ensure that a track discovered in multiple regions is not double-counted, as described in Section 3.5.

3.3. *Cluster Criteria*

Within each annual time bin, we search for local photon concentrations using an energy-dependent angular radius that approximates the Fermi-LAT 68% PSF containment radius [20]:

$$r_{68}(E_{\mathrm{ref}}) = \max\left[0.8° \left(\frac{E_{\mathrm{ref}}}{1\,\mathrm{GeV}}\right)^{-0.8}, 0.1°\right],$$

where E_{ref} is a characteristic value in the center of the band spanning $[E_{\mathrm{min}}, E_{\mathrm{max}}]$.

Even over a single year, a source in our calibrated speed range moves 0.25–0.8°, compared with the 0.1° high-energy search radius. Such motion can broaden or split a source's photon concentration and reduce its recovery probability. The moving-source injections in Section 4 include this within-bin motion when measuring the search efficiency.

We split the photons into five seeding bands: $[1, 2.2]$, $[2.2, 10]$, $[10, 21.5]$, $[21.5, 46.4]$, and $[46.4, 100]$ GeV. This division is necessary because a single search radius would either miss a broader low-energy photon distribution or be influenced too heavily by the background around high-energy photons. Clusters are found using density-based spatial clustering of applications with noise (`DBSCAN`) [23, 24], with the haversine metric and radius $r_{\mathrm{DB}} = r_{68}(E_{\mathrm{ref}})$, since photons lie on the celestial sphere. `DBSCAN` is a particularly suitable choice for this problem because it is able to identify areas of high photon concentration without requiring a preset number of sources or a prescribed cluster shape.

We apply our clustering algorithm only when a band contains at least ten photons. For each photon, we estimate the expected local-background count μ_{local} within its r_{DB}-radius neighborhood from the surrounding $2r_{\mathrm{DB}}$–$5r_{\mathrm{DB}}$ annulus. The photon is designated a core photon when its neighborhood contains at least

$$m_{\mathrm{DB}} = \max[4, \lceil \mu_{\mathrm{local}} + 2.5\sqrt{\mu_{\mathrm{local}}} + 1 \rceil]$$

photons, including the tested photon itself. Core photons within r_{DB} are joined into the same cluster. For each cluster, we compute its centroid (α_c, δ_c) using a spherical mean to avoid the 0°/360° right-ascension wrap discontinuity.

We then mask known stationary sources using the `FL16Y` catalogue. A cluster is rejected if its centroid lies within $\max(0.5°, r_{\mathrm{DB}})$ of any catalogue source, which prevents the misidentification of stationary sources as moving-source candidates. We also exclude photons within $\max(0.5°, 2r_{\mathrm{DB}})$ of `FL16Y` sources from the background annulus and use the remaining unmasked solid angle to normalize the background.

For each surviving centroid, we count N_{on} photons in the r_{DB}-radius aperture and N_{off} photons in the unmasked $2r_{\mathrm{DB}}$–$5r_{\mathrm{DB}}$ annulus. Let $\kappa = A_{\mathrm{on}}/A_{\mathrm{off}}$ be the ratio of their usable solid angles. We reject an annulus retaining less than 25% of its cap-clipped area or containing fewer than ten off photons. For a positive excess, we require the Li–Ma significance (Eq. 17 of Ref. [25])

$$S_{\mathrm{LM}} = \left\{ 2 \left[N_{\mathrm{on}} \ln\left(\frac{1+\kappa}{\kappa} \frac{N_{\mathrm{on}}}{N_{\mathrm{on}} + N_{\mathrm{off}}} \right) + N_{\mathrm{off}} \ln\left((1+\kappa) \frac{N_{\mathrm{off}}}{N_{\mathrm{on}} + N_{\mathrm{off}}} \right) \right] \right\}^{1/2} \geq 2.5.$$

If $N_{\mathrm{on}} \leq \kappa N_{\mathrm{off}}$, we define $S_{\mathrm{LM}} = 0$. We use S_{LM} to select and weight cluster candidates, but defer a rigorous significance calculation for the full search to the scrambled campaigns described in Section 4.1.

3.4. *Cluster Linking*

We now treat the cluster candidates as vertices of a directed graph ordered by time. Let g_{ij} be the positive difference between the annual-bin indices of candidates i and j, and let Δt_{ij} be their elapsed time in years. Candidate j can follow candidate i only if $g_{ij} \in \{1, 2\}$ and their angular separation satisfies

$$d_{ij} \leq (1° \,\mathrm{yr}^{-1}) \Delta t_{ij}.$$

The $1° \,\mathrm{yr}^{-1}$ bound limits the separation of reconstructed centroids between annual bins. By contrast, the 0.25–$0.8° \,\mathrm{yr}^{-1}$ interval specifies the true source speeds over which we calibrate recovery efficiency.

Thus, an individual link can skip at most one intervening annual bin. We retain tracks containing at least five cluster detections, with at most two missing annual bins between the first and last linked detections.

We assign each partial track T a dynamic-programming linking score $Q(T)$. A one-detection track containing candidate j is initialized with $Q(\{j\}) = S_{\mathrm{LM},j}$. For each candidate j, we keep the highest-scoring track ending at j, maximizing Eq. 4 over allowed predecessors, and recover tracks with backtracking. If T ends at candidate i, appending an allowed successor j gives

$$Q(T \parallel j) = Q(T) + S_{\mathrm{LM},j} - \left[\frac{d_{ij}}{g_{ij}} + 0.5(g_{ij} - 1)\right], \tag{4}$$

where $T \parallel j$ denotes T extended by j, d_{ij} is the angular separation in degrees between the two centroids, and g_{ij} is their bin-index difference. Note that d_{ij}/g_{ij} is simply the angular step per bin. This numerical score rewards local photon concentration excesses while simultaneously penalizing missing detections and high angular speed, which reduces the likelihood of misidentifying noise as a moving source.

We apply two further geometric checks to reject tracks inconsistent with steady linear motion:

1. **Anisotropy safeguard.** A real moving source's per-bin centroids form a thin, elongated line on the sky, whereas a stationary source whose centroid jitters from photon-counting noise scatters in a roughly circular blob around its true position. We measure this by projecting the per-bin centroids onto a flat plane tangent to the sphere at the center of the track, giving each of the m_T detections a planar position (x_i, y_i). The two principal axes of the resulting cloud are obtained from the eigendecomposition of the 2×2 sample covariance matrix

$$C = \frac{1}{m_T - 1} \sum_{i=1}^{m_T} \begin{pmatrix} (x_i - \bar{x})^2 & (x_i - \bar{x})(y_i - \bar{y}) \\ (x_i - \bar{x})(y_i - \bar{y}) & (y_i - \bar{y})^2 \end{pmatrix},$$

   where m_T is the number of centroid detections in track T. Its eigenvectors are the long and short axes of the cloud, and its eigenvalues are the variances along those axes, $\lambda_\parallel \geq \lambda_\perp$. The anisotropy

$$A = 1 - \frac{\lambda_\perp}{\lambda_\parallel}$$

   approaches 1 for a perfectly thin tube ($\lambda_\perp = 0$) and 0 for an isotropic blob ($\lambda_\perp = \lambda_\parallel$). Candidates with $A < 0.7$ are rejected as insufficiently elongated.

2. **Constant-direction safeguard.** We project each bin's centroid onto the long axis identified above, giving a scalar position ξ_i for the ith bin, and compare it to the bin midpoint t_i in years. The Pearson correlation

$$r = \frac{\sum_i (t_i - \bar{t})(\xi_i - \bar{\xi})}{\sqrt{\sum_i (t_i - \bar{t})^2}\sqrt{\sum_i (\xi_i - \bar{\xi})^2}}$$

   quantifies the linear relationship between position along the principal axis and time. A perfect moving source gives $|r| \to 1$, with later bins reliably farther along the trajectory. Candidates with $|r| < 0.85$ are rejected as inconsistent with steady progression along the principal axis.

3.5. *Removing Duplicates*

Duplicates can arise in three ways:

- The same photon concentration is detected in multiple energy bands.
- The same cluster detections appear in multiple linked tracks.
- The same track is reconstructed in overlapping sky regions.

Cross-band duplicates. Within each time bin and processing region, we rank cluster candidates by decreasing Li–Ma significance S_{LM}. We retain the highest-ranked candidate and remove any remaining candidate whose centroid separation from it is no greater than the larger of their two clustering radii. We then repeat this process with the highest-ranked remaining candidate until all candidates have been considered. This step precedes temporal linking so that repeat detections in different bands do not yield redundant nodes.

Overlapping linked tracks. After applying the motion filters, we rank tracks within each region by decreasing number of linked detections, breaking ties by the dynamic-programming score. We discard tracks that are already contained within or substantially overlap with a higher-ranked track. Specifically, a track T is discarded if any previously retained track T' satisfies

$$\frac{|C(T) \cap C(T')|}{\min\{|C(T)|, |C(T')|\}} \geq 0.8,$$

where $C(T)$ denotes the set of cluster detections in T. Applying this rule after the motion filters ensures that an invalid extension cannot eliminate an otherwise valid shorter trajectory.

Cross-region duplicates. Finally, we address repetition across the overlapping sky regions. We first assign each track to the Voronoi cell containing the spherical mean of its cluster centroids and retain it only in that cell's processing region. However, minor reconstruction differences can nevertheless cause the same track to be retained by different regions. Two tracks are considered to be duplicates if they overlap in at least three time bins and all centroids are separated by no more than 0.6° in every shared bin. We represent tracks as vertices of a graph and matches as edges, and each connected component therefore represents a group of potentially redundant tracks. From each component, we retain the single track with the largest sum of per-bin Li–Ma significances.

4. CALIBRATION AND VALIDATION

Before interpreting our candidates, we need to determine how often random photons form a plausible moving track by chance and how often the search recovers a moving source when one is present. We estimate the background with time-scrambled skies and test the result on a separate set of scrambled skies. We then inject artificial sources to estimate how effective our algorithm is and where these measurements remain reliable.

4.1. *Estimating False-Positive Detections with Time Scrambling*

We measure the combined track-count distribution using 60 complete time-scrambled all-sky campaigns in order to estimate how often chance concentrations of photons produce moving candidates under our algorithm. Each campaign preserves each photon's direction and energy while drawing a replacement arrival time as described below. This procedure retains the spatial and spectral background along with the local time distribution in expectation while removing the original association between direction and time within each cell.

The time histogram uses one-day bins, with the final bin truncated at the observation boundary. Let $c(i)$ denote the nearest-center cell of photon i, let τ_k be the edges of the time histogram, and

let n_{ck} be the number of events from cell c in interval $[\tau_k, \tau_{k+1})$. The empirical probability that the replacement time for photon i in scrambled campaign b falls in this interval is

$$\Pr\left(t_i^{(b)} \in [\tau_k, \tau_{k+1})\right) = \frac{n_{c(i)k}}{\sum_j n_{c(i)j}}.$$

For photon i in campaign b, we select bin $K_i^{(b)}$ with the probabilities above, then draw a time uniformly within that bin:

$$t_i^{(b)} \sim \mathcal{U}\left(\tau_{K_i^{(b)}}, \tau_{K_i^{(b)}+1}\right).$$

If no local time distribution is available for a cell, we instead use the time distribution of all photons across the full sky.

We generate each pseudorandom draw deterministically from the campaign seed and the photon's run and event identifiers. As a result, a photon present in several overlapping caps receives the same replacement time in all of them. The counts from the 60 scrambled skies define our empirical background distribution.

4.2. *Statistical Validation of the Search Results*

Using 30 calibration null campaigns and 30 held-out validation campaigns, we compare three predictive models for the background candidate count. Let $\widehat{\mu}$ and $\widehat{v}$ denote the calibration-sample mean and variance and C_a denote calibration counts. Our three candidate models are

- **Poisson.** This model assumes the variance is equal to the mean:

$$p_{\mathrm{P}}(n) = e^{-\widehat{\mu}} \frac{\widehat{\mu}^n}{n!}.$$

- **Negative binomial.** This model allows overdispersion, $\widehat{v} > \widehat{\mu}$, using

$$p_{\mathrm{NB}}(n) = \frac{\Gamma(n+r)}{\Gamma(r)n!} \left(\frac{r}{r+\widehat{\mu}}\right)^r \left(\frac{\widehat{\mu}}{r+\widehat{\mu}}\right)^n,$$

  where

$$r = \frac{\widehat{\mu}^2}{\widehat{v} - \widehat{\mu}}.$$

- **Campaign-resampling model.** This empirical model makes no parametric assumption and uses the observed calibration frequencies directly:

$$p_{\mathrm{E}}(n) = \frac{1}{30} \sum_{a=1}^{30} \mathbf{1}(C_a = n).$$

A two-sample Kolmogorov–Smirnov test found the calibration and validation count samples consistent, with a p-value of 0.958. For the negative-binomial candidate, we use the expression above when $\widehat{v} > \widehat{\mu}$, and when the sample is not overdispersed, we replace that candidate with its Poisson limit.

We evaluate each model on a held-out count y using the continuous ranked probability score (CRPS) [26]. If $p(n)$ is the probability that the model assigns to count n, then

$$\mathrm{CRPS}(p, y) = \sum_{n=0}^{\infty} p(n)|n - y| - \frac{1}{2} \sum_{n=0}^{\infty} \sum_{k=0}^{\infty} p(n)p(k)|n - k|,$$

which evaluates both calibration and sharpness rather than just the fitted variance.

For each model M and held-out campaign h, let V_h be its observed track count, let F_M be the cumulative distribution corresponding to p_M, and let $F_M(V_h^-)$ denote the probability assigned to counts strictly below V_h. Discrete predictive calibration is checked with the randomized probability integral transform [27]

$$U_{M,h} = F_M(V_h^-) + W_h p_M(V_h), \qquad W_h \sim \mathcal{U}(0,1),$$

followed by a Kolmogorov–Smirnov test against $\mathcal{U}(0,1)$ and a check that the central 90% predictive interval contains at least 75% of the held-out counts. If $M_\star$ minimizes the mean holdout score, the final one-standard-error condition [28] is

$$\overline{\mathrm{CRPS}}_{\mathrm{E}} \le \overline{\mathrm{CRPS}}_{M_\star} + \mathrm{SE}\Big(\overline{\mathrm{CRPS}}_{M_\star}\Big).$$

Here the negative binomial has the smallest mean score, 15.04, while the campaign-resampling model gives 15.14, within one standard error of the minimum. It also gives a randomized-PIT p-value of 0.442 and held-out interval coverage 0.867, so we retain this model.

The combined null has a variance-to-mean ratio of 1.37, which indicates overdispersion relative to a Poisson model. Let N_b denote the number of tracks found in scrambled sky b. For $N_{\mathrm{null}} = 60$ combined null campaigns and observed count n_{obs}, the plus-one upper-tail probability [29] is

$$\widehat{p}_{\mathrm{count}} = \frac{1 + \sum_{b=1}^{N_{\mathrm{null}}} \mathbf{1}(N_b \ge n_{\mathrm{obs}})}{N_{\mathrm{null}} + 1},$$

with Gaussian equivalent $Z_{\mathrm{emp}} = \Phi^{-1}(1 - \widehat{p}_{\mathrm{count}})$, where Φ is the standard normal cumulative distribution.

The signal-count limit follows from the same empirical background sample $\mathcal{B} = \{N_b\}_{b=1}^{N_{\mathrm{null}}}$. Under the additive signal model, a null-campaign background count N_b is augmented by an independent signal count $K_s \sim \mathrm{Poisson}(s)$. Conditioning on N_b and averaging over campaigns with the plus-one correction gives the lower-tail function

$$G_{\mathcal{B}}(n; s) = \frac{1 + \sum_{b=1}^{N_{\mathrm{null}}} \Pr(N_b + K_s \le n \mid N_b)}{N_{\mathrm{null}} + 1} = \frac{1 + \sum_{b=1}^{N_{\mathrm{null}}} F_{\mathrm{Pois}}(n - N_b; s)}{N_{\mathrm{null}} + 1}, \tag{5}$$

where $F_{\mathrm{Pois}}(k; s) = 0$ for $k < 0$. Because $G_{\mathcal{B}}(n; s)$ is nonincreasing in s, the one-sided $(1-\zeta)$ upper endpoint is

$$\widehat{s}_{1-\zeta}(n) = \begin{cases} 0, & G_{\mathcal{B}}(n; 0) \le \zeta, \\ \text{the unique } s > 0 \text{ satisfying } G_{\mathcal{B}}(n; s) = \zeta, & G_{\mathcal{B}}(n; 0) > \zeta. \end{cases}$$

Finite-background uncertainty is assessed separately by resampling the 60 campaign counts with replacement 2,000 times [30] and recomputing the upper limit. We validate the signal upper-limit construction in 24 parameter combinations using 20,000 pseudoexperiments per combination, each with a background estimate from 60 simulated campaigns.

- Signal means: 0, 10, 25, 50, 87, and 150 tracks.
- Background variance-to-mean ratios: 1, 1.37, 1.5, and 2.

For positive signals, measured coverage ranges from 0.909 to 0.915. At zero signal, coverage is unity because every nonnegative upper limit covers the physical boundary at zero. Together, these results validate the construction across the tested count models and signal range. The cellwise 95% Wilson intervals have half-widths of approximately 0.004 and lie entirely between 0.905 and 0.919.

In this coverage calculation, a background with mean μ_{bg} and variance-to-mean ratio ϕ is drawn from

$$X_{\mathrm{bg}} \sim \begin{cases} \mathrm{Poisson}(\mu_{\mathrm{bg}}), & \phi = 1 \\ \mathrm{NB}\left(\dfrac{\mu_{\mathrm{bg}}}{\phi - 1}; \dfrac{1}{\phi}\right), & \phi > 1 \end{cases},$$

which has $\mathrm{Var}(X_{\mathrm{bg}}) = \phi\mu_{\mathrm{bg}}$. For replicate m, let $\mathcal{B}_m$ be its 60 simulated background-training counts. Its observation is $X_{\mathrm{bg},m}^{\mathrm{obs}} + K_s^{(m)}$. The reported Monte Carlo coverage for $s > 0$ is exactly

$$\widehat{C}_\phi(s) = \frac{1}{20000} \sum_{m=1}^{20000} \mathbf{1}\left[G_{\mathcal{B}_m}\left(X_{\mathrm{bg},m}^{\mathrm{obs}} + K_s^{(m)}; s\right) \geq 0.1\right], \tag{6}$$

which is equivalent, by monotonicity, to counting $s \leq \widehat{s}_{0.9}(X_{\mathrm{bg},m}^{\mathrm{obs}} + K_s^{(m)})$.

Population-level agreement between the observed and null samples is computed for five track properties: track score, angular speed, number of linked detections, total photon count, and number of contributing energy bands. We use two-sample Kolmogorov–Smirnov distances calibrated for each campaign. The score property is defined as $S(T) = \sum_{j \in T} S_{\mathrm{LM},j}$, excluding the linker penalties in Eq. 4.

For property j, let $\widehat{F}_{R,j}$ denote the empirical cumulative distribution function (ECDF) for the real campaign, $\widehat{F}_{0,j}$ the ECDF pooled over all null campaigns, and $\widehat{F}_{b,j}$ the ECDF for null campaign b. The observed statistic and its leave-one-campaign-out null counterparts are

$$\begin{aligned} D_j^{\mathrm{obs}} &= \sup_x \left|\widehat{F}_{R,j}(x) - \widehat{F}_{0,j}(x)\right|, \\ D_{j,b}^0 &= \sup_x \left|\widehat{F}_{b,j}(x) - \widehat{F}_{-b,j}(x)\right|, \\ p_j &= \frac{1 + \sum_{b=1}^{N_{\mathrm{null}}} \mathbf{1}(D_{j,b}^0 \geq D_j^{\mathrm{obs}})}{N_{\mathrm{null}} + 1}, \end{aligned}$$

where $\widehat{F}_{-b,j}$ pools the remaining $N_{\mathrm{null}} - 1$ null campaigns and $N_{\mathrm{null}} = 60$. For the ordered raw probabilities $p_{(1)} \leq p_{(2)} \leq p_{(3)} \leq p_{(4)} \leq p_{(5)}$, the step-down Holm-adjusted [31] values are

$$\widetilde{p}_{(k)} = \min\left\{1, \max_{1 \leq \ell \leq k} (6 - \ell) p_{(\ell)}\right\}.$$

Calibration against complete campaigns accommodates within-campaign dependence without treating candidate tracks as independent, while the Holm adjustment controls the familywise error rate across the five fixed property tests. Candidate-level p_{global} values are calibrated separately by applying the same plus-one campaign ranking to each campaign's maximum track score, as defined in Appendix A.

4.3. *Detection Efficiency of the Moving-Source Algorithm*

We measure how effectively our algorithm detects moving sources by injecting plausible sources into representative background fields and calculating how often the algorithm recovers them.

Detected photon energies are drawn over 1–100 GeV from the normalized distribution

$$p_{\mathrm{det}}(E) = \frac{E^{-2}}{\displaystyle\int_{1\,\mathrm{GeV}}^{100\,\mathrm{GeV}} E^{-2}\,dE}. \tag{7}$$

Each sampled energy then enters the analysis in three ways:

1. It replaces the cloned event's `ENERGY` value.
2. It selects the appropriate energy-dependent double-King PSF used to scatter that photon.
3. It determines which seeding band receives the photon during the search.

Thus, p_{det} controls the injected photons' energy distribution and hence their PSF and clustering efficiency. It does not alter the energies of real or scrambled photons.

For each instrument response function (IRF) cell q, a King component [33] has a radial CDF

$$F_{\mathrm{K}}(\psi_{\mathrm{rad}};\sigma,\gamma,E) = 1 - \left(1 + \frac{\psi_{\mathrm{rad}}^2}{2\gamma\sigma^2\mathcal{R}_{\mathrm{PSF}}(E)^2}\right)^{1-\gamma},$$

where ψ_{rad} is the angular separation in radians and the tabulated PSF scale is

$$\mathcal{R}_{\mathrm{PSF}}(E) = \left[\left\{c_0\left(\frac{E}{100\,\mathrm{MeV}}\right)^{\beta}\right\}^2 + c_1^2\right]^{1/2}.$$

The `P8R3_SOURCE_V3` tables give $\beta = -0.8$ for both conversion types.

For the core and tail widths $\sigma_{\mathrm{c},q}$ and $\sigma_{\mathrm{t},q}$, the cell CDF is

$$F_q(\psi_{\mathrm{rad}}) = f_{\mathrm{c},q}F_{\mathrm{K}}(\psi_{\mathrm{rad}};\sigma_{\mathrm{c},q},\gamma_{\mathrm{c},q},E) + (1-f_{\mathrm{c},q})F_{\mathrm{K}}(\psi_{\mathrm{rad}};\sigma_{\mathrm{t},q},\gamma_{\mathrm{t},q},E)\,,$$

with core fraction

$$f_{\mathrm{c},q} = \left[1 + N_{\mathrm{tail},q}\left(\frac{\sigma_{\mathrm{t},q}}{\sigma_{\mathrm{c},q}}\right)^2\right]^{-1}.$$

The sampler mixes up to four neighboring $(\log_{10}E, \cos\theta)$ cell distributions with bilinear weights w_q:

$$F_{\mathrm{PSF}}(\psi_{\mathrm{rad}} \mid E,\theta,\mathrm{type}) = \sum_{q\in\mathcal{N}(E,\theta)} w_q F_q(\psi_{\mathrm{rad}})$$

with the weights determined by the photon's position relative to the neighboring IRF cells.

After choosing a cell and component, inverse-CDF sampling with $U \sim \mathcal{U}(0,1)$ gives

$$\psi_{\mathrm{rad}} = \mathcal{R}_{\mathrm{PSF}}(E)\sigma\left\{2\gamma\left[(1-U)^{1/(1-\gamma)} - 1\right]\right\}^{1/2}.$$

We perform 50 trials for each of the 45 combinations of photon rate, angular speed, and Galactic-latitude band listed in Table 1. Within each latitude band, the trials are balanced across ten fixed background regions, with five sources injected into each region.

Injection condition		R (detected photons per annual bin)				
$\lvert b\rvert$	μ	8	12	18	25	35
0–10°	$0.25^\circ\,\mathrm{yr}^{-1}$	2%	10%	8%	24%	34%
0–10°	$0.50^\circ\,\mathrm{yr}^{-1}$	0%	0%	8%	16%	34%
0–10°	$0.80^\circ\,\mathrm{yr}^{-1}$	0%	2%	2%	4%	12%
10–30°	$0.25^\circ\,\mathrm{yr}^{-1}$	2%	16%	50%	70%	76%
10–30°	$0.50^\circ\,\mathrm{yr}^{-1}$	2%	6%	40%	66%	90%
10–30°	$0.80^\circ\,\mathrm{yr}^{-1}$	2%	0%	12%	44%	68%
30–90°	$0.25^\circ\,\mathrm{yr}^{-1}$	12%	24%	70%	88%	86%
30–90°	$0.50^\circ\,\mathrm{yr}^{-1}$	8%	10%	62%	90%	96%
30–90°	$0.80^\circ\,\mathrm{yr}^{-1}$	4%	4%	30%	64%	80%

Table 1. Recovery Frequency of Injected Moving Sources. We identify the rate at which our moving-source algorithm recovers artificially injected moving sources for each combination of Galactic-latitude band, angular speed, and photons per annual bin using 50 trials per cell. Recovery generally improves with photon count and is substantially lower near the Galactic plane.

At fixed speed and latitude, recoveries k_i in $n_i = 50$ trials at ordered count nodes R_i give the weighted isotonic estimate

$$(\widehat{\epsilon}_i)_{i=1}^5 = \underset{0\le e_1\le\cdots\le e_5\le 1}{\arg\min} \sum_{i=1}^5 n_i \left(e_i - \frac{k_i}{n_i}\right)^2. \tag{8}$$

When neighboring recovery frequencies violate monotonicity, we pool the adjacent violating points into a shared weighted average until the fitted efficiency sequence is nondecreasing [34].

For all nine latitude–speed curves, we use the binomial deviance to test how well the monotonic fit matches the observed recovery counts:

$$D = 2\sum_{i=1}^5 \left[k_i \ln \frac{k_i}{n_i\widehat{\epsilon}_i} + (n_i - k_i)\ln\frac{n_i - k_i}{n_i(1-\widehat{\epsilon}_i)}\right],$$

with zero-numerator terms set to zero. Bootstrap draws $k_i^{(m)} \sim \mathrm{Binomial}(n_i, \widehat{\epsilon}_i)$ are refitted to obtain D_m^* and the goodness-of-fit probability

$$\widehat{p}_{\mathrm{GOF}} = \frac{1 + \sum_{m=1}^{10000} \mathbf{1}(D_m^* \ge D)}{10001}$$

for 10,000 simulations. The smallest value obtained is $\widehat{p}_{\mathrm{GOF}} = 0.242$, so there is no evidence that our injections are inconsistent with the monotonic efficiency assumption.

To propagate finite-injection uncertainty, we draw each cell's recovery efficiency from the Jeffreys posterior [35],

$$\epsilon_{\mathrm{raw}} \mid k, n \sim \mathrm{Beta}\left(k + \tfrac{1}{2}, n - k + \tfrac{1}{2}\right),$$

with the weighted isotonic projection in Eq. 8 applied to each of 100,000 joint draws. For every projected draw, the full-sky response at a measured count and speed is the weighted average across the three Galactic-latitude bands:

$$\epsilon_{\mathrm{sky}}^{(d)}(R_i, \mu_j) = \frac{78\,\widehat{\epsilon}_{1ji}^{(d)} + 146\,\widehat{\epsilon}_{2ji}^{(d)} + 224\,\widehat{\epsilon}_{3ji}^{(d)}}{448}.$$

The weights 78/448, 146/448, and 224/448 are the fractions of the 448 survey regions in the three latitude bands.

Between specific grid values, we use bilinear interpolation in $\ln R$ and μ within the ranges of $8 \leq R \leq 35$ and $0.25^\circ\,\mathrm{yr}^{-1} \leq \mu \leq 0.8^\circ\,\mathrm{yr}^{-1}$. For interpolation weights w_R and w_μ for the photon count and angular speed, respectively,

$$w_R = \frac{\ln R - \ln R_i}{\ln R_{i+1} - \ln R_i}, w_\mu = \frac{\mu - \mu_j}{\mu_{j+1} - \mu_j}.$$

Then, within a latitude band (and identically for each posterior draw),

$$\widetilde{\epsilon}(R, \mu) = (1 - w_\mu)\left[(1 - w_R)\widehat{\epsilon}_{j,i} + w_R\widehat{\epsilon}_{j,i+1}\right] + w_\mu\left[(1 - w_R)\widehat{\epsilon}_{j+1,i} + w_R\widehat{\epsilon}_{j+1,i+1}\right]. \tag{9}$$

Figure 2 shows the measured recovery frequencies at the sampled photon counts and angular speeds. Recovery rises sharply with photon count; at 25–35 detected photons per annual bin, recovery spans 64–96% at $|b| = 30^\circ$–90° and 44–90% at $|b| = 10^\circ$–30°, but only 4–34% within 10° of the Galactic plane. The higher diffuse background near the plane (lower $|b|$) raises the local density threshold and lowers the significance of an injected cluster, so fewer clusters survive into the track-linking stage. The $0.8^\circ\,\mathrm{yr}^{-1}$ curves also generally lie below the slower-speed curves because faster motion broadens or splits the source within an annual bin, while the linking score penalizes the larger step between consecutive centroids.

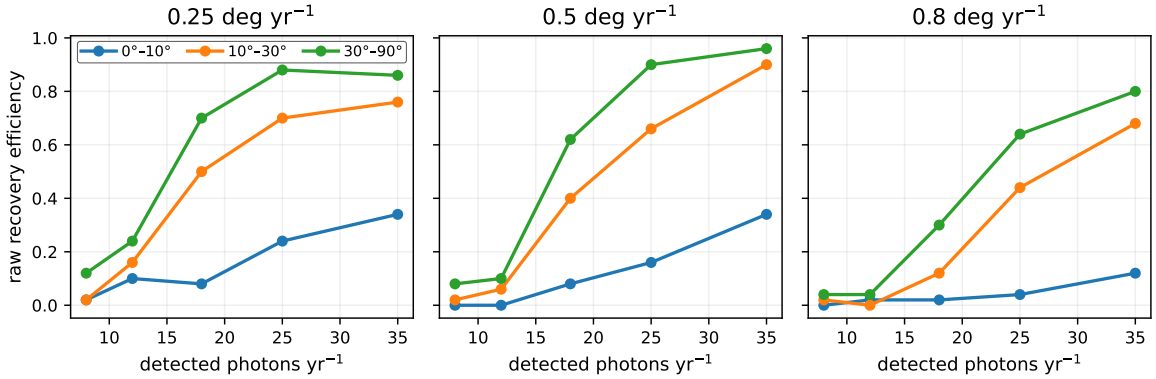


Figure 2. Measured injection-recovery efficiency as a function of detected photons per annual bin. Panels show angular speeds of 0.25, 0.5, and 0.8 deg yr^{-1}, while the curves show the three Galactic-latitude bands. Each point represents 50 injection trials.

4.4. *Detection Efficiency as a Function of Source Flux*

A steady flux F_γ produces fluctuating annual counts rather than the same number of photons per year. The position-dependent conversion from integrated photon flux to expected counts uses the LAT exposure averaged with the corresponding incident E^{-2} spectrum. For the resulting spectrum-weighted window exposure $\mathcal{E}_q$ in $\mathrm{cm}^2\,\mathrm{s}$ at region q, the mean per year is $\lambda_q = F_\gamma \mathcal{E}_q / T_{\mathrm{yr}}$, where $T_{\mathrm{yr}} = 10.01$ is the total observation duration in years. Approximating the observing bins by 10 annual bins, the recovery efficiency is Poisson-convolved to represent steady emission, giving

$$\epsilon_{\mathrm{P}}(\lambda, \mu, b) = \sum_{K=0}^{\infty} \frac{e^{-10\lambda}(10\lambda)^K}{K!}\widetilde{\epsilon}(K/10, \mu, b). \tag{10}$$

Poisson addition gives the correct total-count distribution, but at fixed K, the event distribution across years may still affect linkage. Appendix B.2 derives the resulting scalar-response approximation. We now validate this approximation with 120 injections across 12 cells that cover all latitude bands and speeds. Annual injection rates follow the host field's yearly photon counts, with an RMS

modulation of 6.16%. The Poisson response predicts 53.37 recoveries, compared with 51 observed, giving $z = -0.493$ within the conventional two-sided 95% criterion $|z| \leq 1.96$. Paired exact-count injections yield 56 recoveries, with no significant directional difference under a two-sided binomial test with a p-value of 0.458.

4.5. *Injection Validation and Calibration Boundaries*

We test two potential sources of response error: the accuracy of the double-King angular approximation relative to `gtobssim` [36] and the treatment of exact-count efficiencies outside the calibrated range during Poisson averaging.

Double-King approximation check. We wish to verify that our double-King injections are roughly equivalent to the official Fermi-supported `gtobssim` injection method. To do this, we compare 110 unique matched injections and ensure both versions preserve photon counts and trajectories.

At exact detected rate R (photons per annual bin), let $\epsilon^{\rm pair}_{\rm gto}(R)$ and $\epsilon^{\rm pair}_{\rm King}(R)$ be the fractions of matched injections recovered with `gtobssim` and the double-King approximation, respectively. Their difference is $\Delta\epsilon_{\rm pair}(R) = \epsilon^{\rm pair}_{\rm gto}(R) - \epsilon^{\rm pair}_{\rm King}(R)$. The resulting fractional shift in the survey-level source limit, based on the survey-weighted efficiency $\epsilon^{\rm survey}_{\rm King}(R)$ from the 50-trial-per-combination double-King calibration, is

$$\delta_{\rm lim}(R) = \left| \frac{\epsilon^{\rm survey}_{\rm King}(R)}{\mathrm{clip}_{[0,1]}[\epsilon^{\rm survey}_{\rm King}(R) + \Delta\epsilon_{\rm pair}(R)]} - 1 \right|.$$

We use McNemar's test [37] on the number of sources n_{01} recovered only with `gtobssim`, and n_{10}, the number of sources recovered only with the double-King approximation. At 18, 25, and 35 photons per annual bin, $(n_{01}, n_{10}) = (1, 2)$, $(8, 5)$, and $(5, 4)$, yielding two-sided p-values of 1.00, 0.581, and 1.00 and source-limit shifts of 7.90%, 8.15%, and 3.87%. None provides statistically significant evidence of a directional difference in recovery.

Poisson-response boundaries. We measure recovery for exact detected rates $8 \leq R \leq 35$ photons per annual bin. Poisson averaging nevertheless assigns nonzero probability to realized rates outside this range, even when the physical mean λ lies within it. At calibrated angular speeds, we therefore extend the conditional response $\widetilde{\epsilon}$ from Eq. 9 as

$$\widetilde{\epsilon}_{\rm ext}(R, \mu, b) = \begin{cases} 0, & R < 8, \\ \widetilde{\epsilon}(R, \mu, b), & 8 \leq R \leq 35, \\ \widetilde{\epsilon}(35, \mu, b), & R > 35. \end{cases} \tag{11}$$

Relative to the exact-count calculation, the primary Poisson convolution reduces the central effective volumes by 0.20–1.08%.

5. RESULTS

We compare the moving-track candidates found in the Fermi-LAT data to those found by chance in time-scrambled skies. This tells us how many additional moving tracks could still be consistent with the data. We then convert this result into bounds on the number of moving sources across the full sky and consider what these limits mean for two possible populations of primordial black holes.

5.1. *Upper Limit on Detected Tracks*

Our search returned 577 unique moving-track candidates; the 15 with the largest summed track scores are listed in Appendix A. Meanwhile, the 60 time-scrambled null campaigns find a mean of 532.8 and a standard deviation of 27.0 tracks. Since 3 null campaigns produced at least as many moving tracks as the real data, we find an empirical upper-tail probability of 0.066, which maps to a one-sided Gaussian significance of 1.51σ, and we cannot interpret this as evidence of a moving-source population. We therefore apply Eq. 5 with $n = 577$ and $N_{\mathrm{null}} = 60$. Solving $G_{\mathcal{B}}(577; s) = 0.10$ gives the one-sided 90% confidence upper limit $s_{90} = 80.8$ expected detected signal tracks. The 5th and 95th percentiles of the bootstrap estimates of s_{90} give a 90% confidence interval of $[73.2, 86.3]$. Figure 3 compares the observed FL16Y track count with the counts from all 60 matched null campaigns.

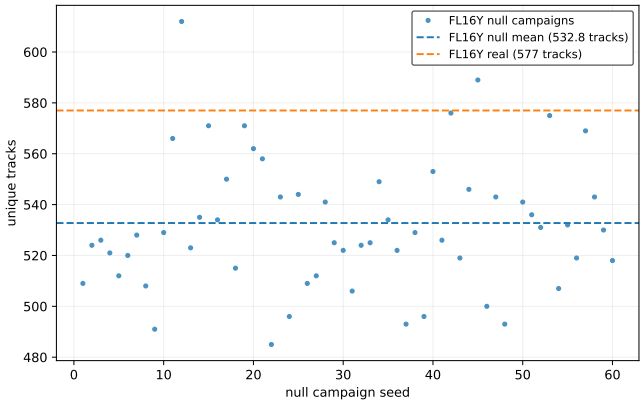


Figure 3. Unique track counts from the 60 time-scrambled FL16Y null campaigns. The dashed blue line marks their mean of 532.8 tracks, and the dashed orange line marks the observed FL16Y track count of 577. We find no significant moving-source excess with a p-value of 0.066 and set a 90% confidence upper limit of 80.8 expected detected signal tracks.

We also test whether the candidate population differs from the null campaigns in summed Li–Ma score, angular velocity, track length, photon count, or energy-band occupancy. Figure 4 summarizes these comparisons, and all five p-values exceed 0.05, providing no evidence of an unusual population in these properties.

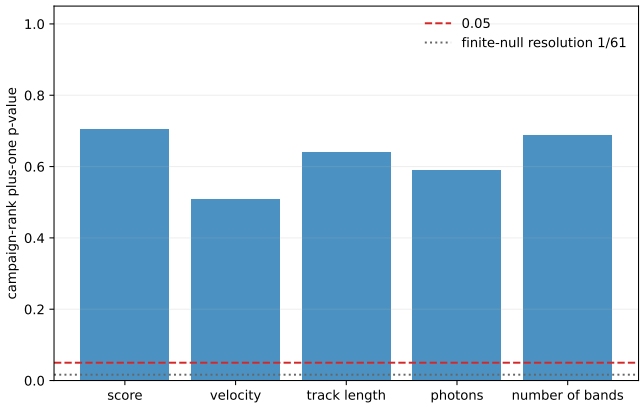


Figure 4. Population screening using campaign-rank plus-one p-values for score, angular velocity, track length, photon count, and energy-band count. The dashed line marks a p-value of 0.05 and the dotted line marks the final null resolution of 1/61. None of the five properties differs significantly between the observed candidates and the null populations.

5.2. *Limits on the Full-Sky Source Population*

Consider steady sources with a common photon flux F_γ and angular speed μ, distributed uniformly in sky direction. Averaging over the 448 survey regions gives the detection probability

$$\bar{\epsilon}(F_\gamma, \mu) = \frac{1}{448} \sum_{q=1}^{448} \epsilon_{\mathrm{P}} \left(\frac{F_\gamma \mathcal{E}_q}{T_{\mathrm{yr}}}, \mu, b_q \right),$$

where b_q denotes the Galactic-latitude band containing region q. Under the Poisson population model, a source-count normalization N_{src} produces mean detected count $s = N_{\mathrm{src}}\bar{\epsilon}$. The full-sky upper limit is therefore

$$N_{90}(F_\gamma, \mu) = \frac{s_{90}}{\bar{\epsilon}(F_\gamma, \mu)}.$$

The strongest central limit is $N_{90} = 122$ at $F_\gamma = 7.08 \times 10^{-10}\,\mathrm{ph\,cm^{-2}\,s^{-1}}$ and $\mu = 0.5^\circ\,\mathrm{yr}^{-1}$, with average efficiency 0.659 and joint 90% interval of $[110, 135]$. Conditioning instead on exactly 35 photons in every annual bin and region gives 97.0 sources, with interval $[87.8, 105.8]$. The physical limit is weaker because a common flux produces exposure-dependent Poisson means, and the realized photon counts fluctuate about those means. At high flux, an increasing number of regional means exceed the calibrated upper boundary, producing the upturn shown in Figure 5.

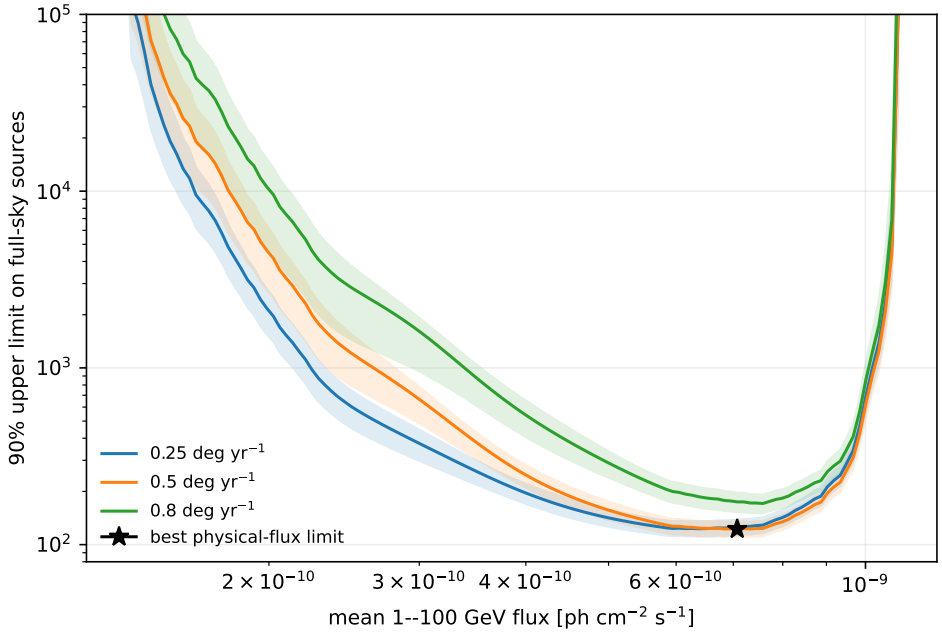


Figure 5. Full-sky 90% upper limits on the number of steady moving sources as a function of mean photon flux. Curves show the central limits for angular speeds of 0.25, 0.5, and 0.8° yr^{-1} under the Poisson photon-count response, while the shaded bands show the joint 90% uncertainty intervals. The black star marks the strongest central limit of $N_{90} = 122$.

FL16Y sensitivity. A sufficiently bright moving source could enter the `FL16Y` stationary catalogue. Because our algorithm masks catalogue sources, that source would then be excluded from the discovery search. We assess this possibility through the ratio of its stationary-equivalent flux to the local catalogue threshold:

$$q_i = \frac{R_i T_{\mathrm{yr}}\, \Pi(\mu_i)}{\mathcal{E}_i F_{\mathrm{th},i} c_{\mathrm{cat}}},$$

where R_i is the detected photon count per annual bin, $T_{\rm yr}$ is the window duration in years, $\mathcal{E}_i$ is the local window exposure, and $F_{{\rm th},i}$ is the converted catalogue flux threshold [38]. The factor $\Pi(\mu_i)$ is the fraction of a moving source's flux recovered by a stationary point-source fit. We estimate this fraction by projecting a uniform linear trail over the Fermi-LAT catalogue baseline onto the survey-averaged stationary PSF. We combine the projections over 1–100 GeV using weak-source likelihood weights, with photon indices of 2.0 for the source and 2.3 for the diffuse background. We scale the catalogue threshold by $c_{\rm cat} = 0.957$, the median flux-to-threshold ratio of near-threshold hard FL16Y sources. A moving source could be misidentified as a catalogue source when $q_i \geq 1$. The nominal calculation gives $q_i < 1$ for all 2,250 injections, with a maximum ratio of 0.904.

5.2.1. *Effective Survey Volume*

For an unresolved source of 1–100 GeV photon luminosity L_γ, the flux and angular speed depend jointly on distance. Integrating the detection probability over position and transverse speed gives

$$V_{\rm eff}(L_\gamma) = \int d\Omega \int_0^\infty r^2\, dr \int_0^\infty dv_\perp\, p(v_\perp \mid \hat{\boldsymbol{n}}) \times \epsilon_{\rm P}\left(\frac{L_\gamma \mathcal{E}(\hat{\boldsymbol{n}})}{4\pi r^2 T_{\rm yr}}, \frac{v_\perp}{r}, b\right).$$

The response is zero outside the adopted physical mean-count (λ) and angular-speed ranges. At fixed direction and transverse speed, their intersection sets the lower distance boundary

$$r_{\rm lo} = \max\left\{\sqrt{\frac{L_\gamma \mathcal{E}(\hat{\boldsymbol{n}})}{4\pi T_{\rm yr}\lambda_{\rm max}}}, \frac{v_\perp}{\mu_{\rm max}}\right\}. \tag{12}$$

The upper boundary is

$$r_{\rm hi} = \min\left\{\sqrt{\frac{L_\gamma \mathcal{E}(\hat{\boldsymbol{n}})}{4\pi T_{\rm yr}\lambda_{\rm min}}}, \frac{v_\perp}{\mu_{\rm min}}\right\}.$$

Here $\lambda_{\rm min} = 8$ and $\lambda_{\rm max} = 35$ photons per year, while $\mu_{\rm min} = 0.25^\circ\,{\rm yr}^{-1}$ and $\mu_{\rm max} = 0.8^\circ\,{\rm yr}^{-1}$ are converted to radians per matching time unit in the distance ratios. The contribution vanishes when $r_{\rm lo} \geq r_{\rm hi}$. Increasing luminosity shifts the count-selected interval outward as $L_\gamma^{1/2}$, whereas the motion-selected interval at fixed speed is unchanged.

The Galactic-rest velocity components are independent, untruncated Gaussians with dispersion $\sigma_v = 238/\sqrt{2} = 168.3\,{\rm km\,s^{-1}}$, following the standard halo model speed recommended in Ref. [18]. We subtract Solar motion, comprising a local-standard-of-rest speed of $232\,{\rm km\,s^{-1}}$, consistent with the measured $232.8\pm3.0\,{\rm km\,s^{-1}}$ [39], and peculiar velocity $(U, V, W) = (11.1, 12.24, 7.25)\,{\rm km\,s^{-1}}$ [40]. The resulting transverse-speed distribution is

$$p(v_\perp \mid \hat{\boldsymbol{n}}) = \frac{v_\perp}{\sigma_v^2} \exp\left[-\frac{v_\perp^2 + \nu^2}{2\sigma_v^2}\right] I_0\left(\frac{v_\perp \nu}{\sigma_v^2}\right),$$

where $\nu(\hat{\boldsymbol{n}})$ is the transverse Solar speed and I_0 is the modified Bessel function of the first kind. Appendix B.4 derives this Rice distribution [41] and its role in the volume integral.

Figure 6 shows the effective survey volume as a function of intrinsic photon luminosity. At luminosities 10^{25}, 10^{26}, and $10^{27}\,{\rm ph\,s^{-1}}$, the central effective volumes are 1.67×10^{-6}, 1.66×10^{-4}, and $3.82\times10^{-4}\,{\rm pc^3}$, respectively. The corresponding equal-tail 90% intervals are $[1.55\times10^{-6}, 1.99\times10^{-6}]$, $[1.58\times10^{-4}, 1.80\times10^{-4}]$, and $[2.60\times10^{-4}, 5.02\times10^{-4}]\,{\rm pc^3}$. The calculation spans nine luminosity nodes from 10^{24} to $10^{28}\,{\rm ph\,s^{-1}}$, with linear interpolation in $\log L_\gamma$–$\log V_{\rm eff}$ confined to that interval. The volume rises and then falls because, at sufficiently high luminosity, the count-selected distances require transverse speeds in the sparse tail of the adopted velocity distribution.

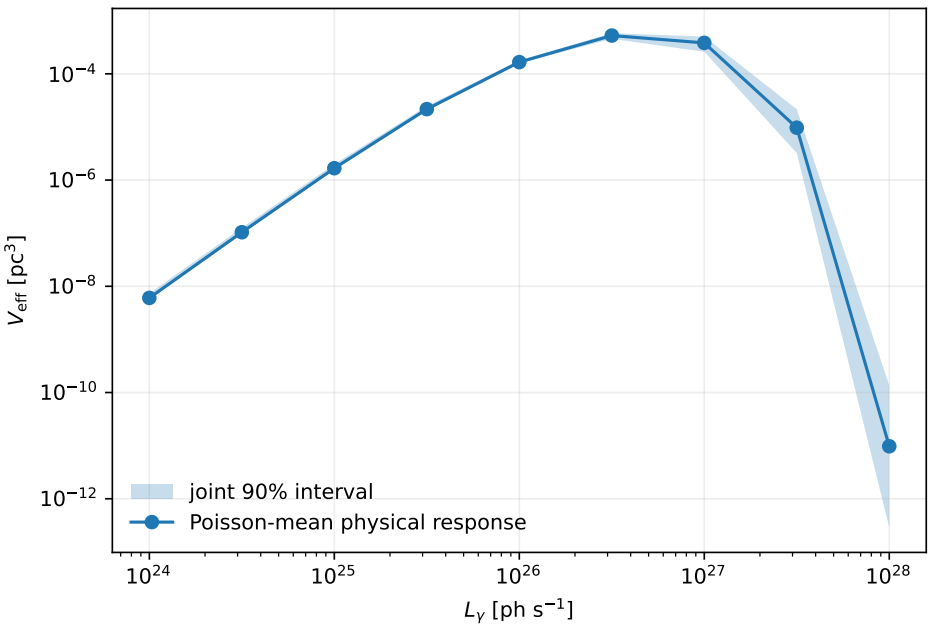


Figure 6. Effective survey volume versus intrinsic 1–100 GeV photon luminosity for the adopted heliocentric velocity distribution and Poisson-averaged response. The band gives the 90% volume interval.

Host-region sensitivity. We assess whether recovery is heterogeneous across different injection points within each latitude band. Although recovery varies among hosts within each latitude band, this variation has a negligible effect on the population limits as omitting any one host changes the strongest exact-count limit by at most 1.06%. At the flux giving the strongest physical limit, 100,000 bootstrap draws that resample whole hosts [30] change the joint 90% interval from 110.4–134.5 to 109.3–134.7 sources. Across the nine luminosity nodes, the largest change in an effective-volume interval endpoint is 7.52%.

5.2.2. *Local Number Density and Dark Matter Abundance*

For a population of approximately uniform local number density n, the expected detected count is $s = nV_{\rm eff}$. Thus, the corresponding limit is

$$n_{90}(L_\gamma) = \frac{s_{90}}{V_{\rm eff}(L_\gamma)}. \tag{13}$$

If each object has mass M, the population's local dark matter fraction is $f = Mn/\rho_{\rm DM}$, giving

$$f_{90}(M, L_\gamma) = \frac{M s_{90}}{\rho_{\rm DM} V_{\rm eff}(L_\gamma)}. \tag{14}$$

The central local mass density is $\rho_{\rm DM} = 0.4\,{\rm GeV\,cm^{-3}}$, the value near which estimates have converged [42] and which is consistent with $0.38 \pm 0.04\,{\rm GeV\,cm^{-3}}$ [39]. We assign a uniform uncertainty over 0.3–0.5 GeV cm^{-3}, the range preferred by global mass models [42]. At each luminosity, we propagate uncertainties in efficiency, exposure, and the velocity model through the Poisson-convolved response to obtain the distribution of $V_{\rm eff}$. The $\rho_{\rm DM}$ uncertainty enters only when calculating f_{90} through Eq. 14. The 5th and 95th percentiles of the $V_{\rm eff}$ distribution define the equal-tail 90% interval.

At luminosities 10^{25}, 10^{26}, and $10^{27}\,{\rm ph\,s^{-1}}$, the central number-density limits are 4.83×10^7, 4.86×10^5, and $2.11 \times 10^5\,{\rm pc^{-3}}$. The corresponding masses at which $f_{90} = 1$ are 2.17×10^{-10},

2.16×10^{-8}, and $4.97 \times 10^{-8}\, M_\odot$. At fixed luminosity, the fraction limit increases linearly with mass (Figure 7); values above unity do not exclude the assumed emitters from constituting all local dark matter.

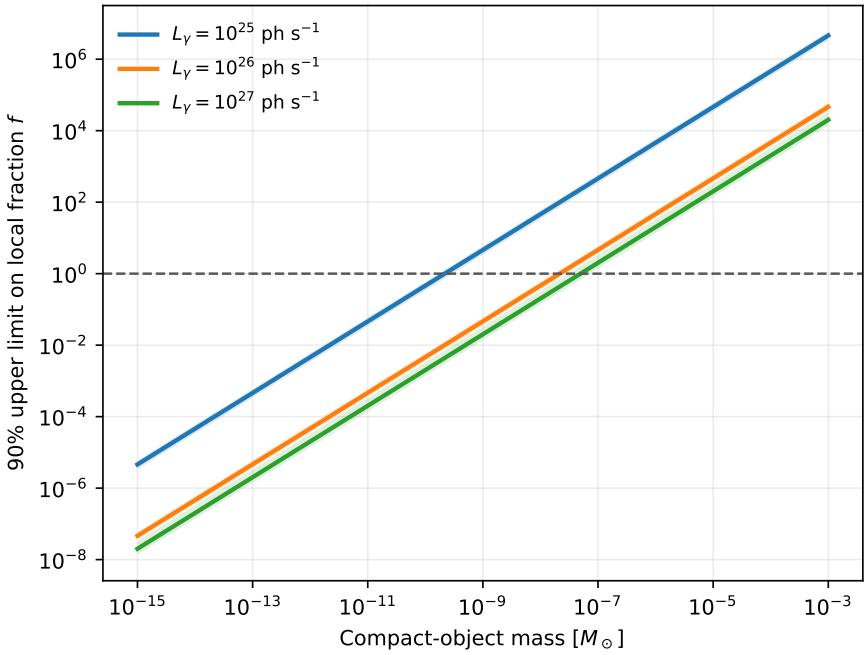


Figure 7. 90% upper limits on the local dark matter fraction in compact objects of mass M for varying luminosities. At a fixed luminosity, the fractional limit is proportional to mass.

5.3. *Primordial Black Hole Benchmarks*

The PBH results are translations of a common reference response, rather than spectrum-specific. In both cases, only the resulting band-integrated luminosity is passed through the reference response, not the full model spectrum. Using an E^{-2} law provides a broad, conventional benchmark that populates all five search bands and retains the principal energy dependence of the LAT exposure. A harder, softer, or strongly curved spectrum could redistribute photons among the bands and alter the limits at the $\mathcal{O}(1)$ level. We leave full spectrum-specific PBH exclusions to Section 7.

5.3.1. *Dressed Primordial Black Holes*

We first consider the monochromatic dressed-PBH model, which assumes that all PBHs have the same mass and are surrounded by dark matter halos formed before matter–radiation equality. The particle model is a self-conjugate $m_\chi = 1\,\mathrm{TeV}$ weakly interacting massive particle (WIMP) annihilating into $b\bar{b}$ with $\langle\sigma v\rangle = 3 \times 10^{-26}\,\mathrm{cm^3\,s^{-1}}$ [43]. Dark matter falling toward each PBH initially forms a steep density profile [9, 43],

$$\rho_{\rm sp}(r) = \frac{\rho_{\rm eq}}{2}\left(\frac{r_{\rm ta}}{r}\right)^{9/4},$$

with turnaround radius $r_{\rm ta} = (2GM_{\rm PBH}t_{\rm eq}^2)^{1/3}$, where $t_{\rm eq}$ and $\rho_{\rm eq}$ are the time and total density at equality.

Near the PBH, WIMP annihilations cap the density at $\rho_{\rm ann} = m_\chi/(\langle\sigma v\rangle t_{\rm age})$ [43], where $t_{\rm age} = 13.8\,\mathrm{Gyr}$ is the adopted halo age, equal to the age of the Universe [1]. The halo is truncated at the

outer radius $r_{\rm ta}$ [44] and integrated inward to the inner cutoff $r_{\rm min} = 4r_s$, where $r_s = 2GM_{\rm PBH}/c^2$ is the Schwarzschild radius [45]. Throughout the evaluated mass range, the plateau radius r_c defined in Appendix B.5 exceeds the kinetic-break radius $r_K \simeq 10\,(m_\chi/\mathrm{TeV})^{-9/2}(M_{\rm PBH}/M_\odot)^{-1}\,\mathrm{cm}$, inside which WIMP thermal velocities flatten the spike [44], so the $r^{-9/4}$ form holds wherever annihilation has not already capped the density.

For dressed PBHs, we calculate the physical 1–100 GeV luminosity by integrating the `PPPC4DMID` 1 TeV $b\bar{b}$ annihilation yield [46]. For the capped profile $\rho(r)$, the photon luminosity from self-conjugate annihilations is

$$L_\gamma = \frac{\langle\sigma v\rangle}{2m_\chi^2}\left(\int \rho^2(r)\,d^3r\right)\left(\int_{1\,\mathrm{GeV}}^{100\,\mathrm{GeV}} \frac{dN_\gamma}{dE}\,dE\right). \tag{15}$$

Appendix B.5 evaluates the spatial integral and shows that luminosity is approximately proportional to PBH mass. Because $V_{\rm eff}$ itself rises and falls with luminosity, the fraction limit is nonmonotonic: $f_{\rm PBH,90} \propto M_{\rm PBH}/V_{\rm eff}[L_\gamma(M_{\rm PBH})]$. Figure 8 shows the resulting abundance limit as a function of PBH mass. Its strongest central limit is $f_{\rm PBH,90} = 8.49 \times 10^{-3}$ at $5.79 \times 10^{-10}\,M_\odot$, rising to 1.70×10^{-2} at $10^{-9}\,M_\odot$. The conservative catalogue-veto check weakens the strongest dressed-PBH bound by 0.794%.

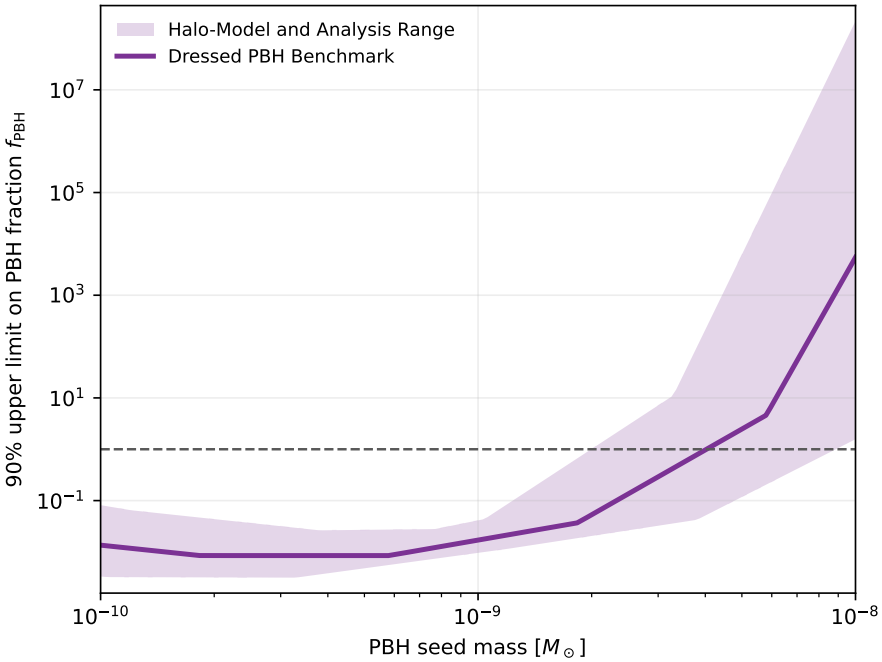


Figure 8. Dressed-PBH abundance limit for a surviving annihilation halo around each PBH. The central curve uses the $b\bar{b}$ luminosity model and the recovery response calibrated for an injected 1–100 GeV photon spectrum $dN_\gamma/dE \propto E^{-2}$. The curve and band are restricted to the evaluated benchmark mass range of $10^{-10}\,M_\odot$ to $10^{-8}\,M_\odot$. The strongest central 90% confidence limit is $f_{\rm PBH,90} = 8.49 \times 10^{-3}$ at $M_{\rm PBH} = 5.79 \times 10^{-10}\,M_\odot$.

At the mass where the central limit is minimized, $M = 5.79 \times 10^{-10}\,M_\odot$, the evaluated halo-model bracket [44] gives $f_{\rm PBH,90}$ limits from 5.56×10^{-3} to 2.72×10^{-2}.

5.3.2. *Quasi-Steady Evaporating Primordial Black Holes*

Ground-based searches with H.E.S.S. and HAWC have constrained terminal PBH evaporation bursts [47, 48]. Here we instead consider PBHs whose gamma-ray emission changes slowly enough

to be treated as quasi-steady over the ten-year Fermi-LAT baseline. The current PBH mass sets the Hawking temperature [5],

$$k_{\rm B}T_{\rm H} = \frac{\hbar c^3}{8\pi GM} \simeq 1.06\,{\rm GeV}\left(\frac{10^{13}\,{\rm g}}{M}\right).$$

Thus $T_{\rm H} \propto M^{-1}$: lower-mass PBHs are hotter and emit a harder Hawking spectrum.

For evaporating PBHs, we integrate the adopted Hawking/QCD spectrum [49, 50] over the same energy band and observing interval. Let $Y_\gamma(\tau)$ denote the resulting total 1–100 GeV photon yield emitted over a PBH's final τ years, using the analytic Standard Model parametrization adopted in Ref. [14]. For $\tau > \Delta t$, we compute the mean luminosity over the observing interval $\Delta t = 10\,{\rm yr}$ as $\overline{L}_\gamma(\tau) = [Y_\gamma(\tau) - Y_\gamma(\tau - \Delta t)]/\Delta t$. The analytic Fermi-LAT spectrum includes its published overall rescaling factor of 0.45. Separately, we adopt a 35% photon-energy fraction as the central model normalization and span 25%–45% [14], corresponding to multiplicative factors of 0.714–1.29 about the central spectrum.

To retain quasi-steady emission, we require a remaining lifetime of at least 100 years and less than 25% intrinsic luminosity evolution across the decade. Within the calibrated luminosity support, the accepted present masses are 1.99×10^{12}–$4.03 \times 10^{12}\,{\rm g}$. The low-mass endpoint is set by the 100-year lifetime requirement. The high-mass endpoint is the largest mass for which the lowest-normalization spectrum still lies within the calibrated luminosity range. The low- and high-mass endpoints correspond, respectively, to Hawking temperatures of 5.32 and 2.63 GeV, remaining lifetimes of 100 and 827 years, and mean luminosities of 6.12×10^{24} and $1.40 \times 10^{24}\,{\rm ph\,s^{-1}}$. The luminosity increases by at most 6.81% across the observing decade, well within the 25% quasi-steady threshold. The strongest central number-density limit occurs at the 100-year boundary, where $M = 1.99 \times 10^{12}\,{\rm g}$, $\overline{L}_\gamma = 6.12 \times 10^{24}\,{\rm ph\,s^{-1}}$, and $V_{\rm eff} = 5.10 \times 10^{-7}\,{\rm pc^3}$. Substituting $s_{90} = 80.8$ into Eq. 13 gives $n_{90} = 1.58 \times 10^{8}\,{\rm pc^{-3}}$ (Figure 9). To express this number-density result as a progenitor-equivalent dark matter fraction, we assign each accepted PBH the reference progenitor mass $M_* = 5 \times 10^{14}\,{\rm g}$, for which the lifetime is close to the age of the Universe [4, 14]. Together with the adopted local dark matter density, Eq. 14 then gives the strongest progenitor-equivalent fraction limit, $f_{\rm PBH,90} = 3.79 \times 10^{-9}$. This fraction refers to progenitor mass assigned to PBHs in the accepted evaporation phase, rather than to their present mass.

For a steady local evaporation-rate density $\dot{n}$, the number density in a remaining-lifetime interval of width $d\tau > 0$ is $dn = \dot{n}\,d\tau$: each PBH spends $d\tau$ in that interval because $d\tau/dt = -1$. The expected count is $s = \dot{n}\mathcal{V}_\tau$, where the volume weighted by residence time is

$$\mathcal{V}_\tau = \int_{\tau_{\rm min}}^{\tau_{\rm max}} V_{\rm eff}[\overline{L}_\gamma(\tau)]\,d\tau.$$

The adopted evaporation law [5, 14] has $\tau = C_\tau \vartheta^{-3}$ for $\vartheta = k_{\rm B}T_{\rm H}$ expressed in GeV. Thus the temperature quadrature used in the calculation is

$$\mathcal{V}_\tau = \int_{\vartheta_{\rm min}}^{\vartheta_{\rm max}} V_{\rm eff}[\overline{L}_\gamma(\vartheta)]\frac{3C_\tau}{\vartheta^4}\,d\vartheta,$$

where C_τ is the lifetime normalization of the adopted spectrum and the integration limits correspond to the accepted phase. The resulting evaporation-rate bound is

$$\dot{n}_{90} = \frac{s_{90}}{\mathcal{V}_\tau}.$$

The central volume $\mathcal{V}_\tau = 5.92 \times 10^{-5}\,{\rm pc^3\,yr}$ gives $\dot{n}_{90} = 1.37 \times 10^{6}\,{\rm pc^{-3}\,yr^{-1}}$, with combined model and analysis envelope 5.60×10^{5}–$3.59 \times 10^{6}\,{\rm pc^{-3}\,yr^{-1}}$, which uses the E^{-2} response.

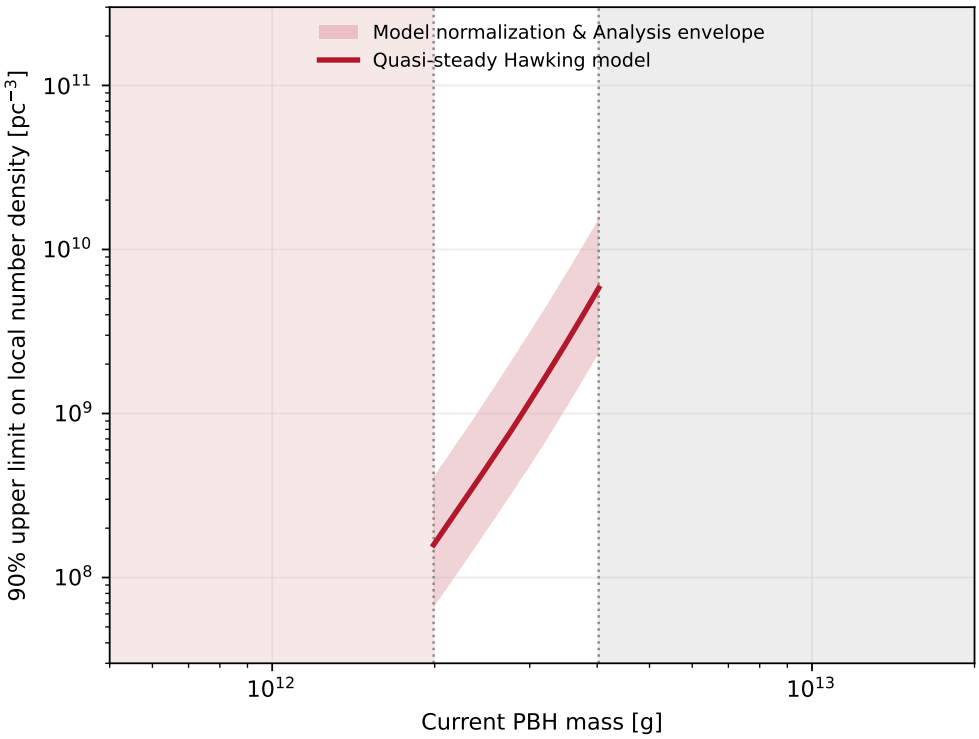


Figure 9. Quasi-steady evaporating-PBH number-density limit versus present mass. The evaluated phase has remaining lifetimes of 100–827 years. Shaded sidebands mark the rapidly evolving final phase and the region outside the calibrated luminosity support. The strongest central 90% confidence limit is $n_{90} = 1.58 \times 10^{8}\,\mathrm{pc}^{-3}$ at a present PBH mass of $1.99 \times 10^{12}\,\mathrm{g}$.

6. DISCUSSION

We compare our background-corrected source limit with a more conservative bound that treats every candidate as a possible source. We then discuss how the calibrated photon-count and angular-speed ranges set the distance range over which a source of a given luminosity and transverse speed can be detected.

6.1. *Background-Corrected Source Limits*

The observed count and candidate properties are consistent with accidental tracks in the matched null skies and do not provide evidence for a moving-source population. Within the validated additive signal model, the observed track count relative to this calibrated null distribution constrains an additional moving-source population to $s_{90} = 80.8$ expected detected tracks. Treating all observed tracks as possible signal gives the Poisson bound

$$\sum_{j=0}^{577} \frac{e^{-s_{90}^{\mathrm{all}}} (s_{90}^{\mathrm{all}})^{j}}{j!} = 0.10. \tag{16}$$

Solving Eq. 16 gives $s_{90}^{\mathrm{all}} = 609$, the background-free upper limit obtained by treating all observed tracks as possible signal.

6.2. *Effects of Source Luminosity and Motion*

The population limits use the recovery response only within our calibrated photon range and our calibrated angular-speed range. At fixed luminosity and transverse speed, these ranges each

restrict the contributing source distances. As luminosity increases, this joint distance window moves outward, so accepted sources require velocities increasingly deep in the rare high-speed tail. This tradeoff causes the effective volume to rise and then fall, producing corresponding minima in the abundance-limit curves. Solar motion makes the transverse-speed distribution dependent on direction, which means the selection response also varies across the sky.

The logarithmic slope of a mass-dependent population limit follows directly from Eq. 14. Where a model luminosity $L_\gamma(M)$ and the interpolated volume are differentiable,

$$\frac{d\ln f_{90}}{d\ln M} = 1 - \frac{d\ln V_{\rm eff}}{d\ln L_\gamma}\frac{d\ln L_\gamma}{d\ln M}.$$

For the approximately linear dressed-PBH luminosity relation, the fraction limit decreases with mass where volume grows faster than luminosity and increases where it grows more slowly. A minimum can therefore occur even when the luminosity is monotonic. Keeping the mass and luminosity axes independent separates this measured selection effect from the emission model.

7. FUTURE WORK

We propose expanding the number of null campaigns used to calibrate the search and investigating two further algorithms for identifying moving-source candidates: a trajectory search using the Hough transform and a U-Net convolutional neural network. We then discuss three potential applications of moving-source searches to other physical scenarios: asteroid-mass PBHs, UCMHs, and nearby sources with annual parallax.

Additional Background Simulations. Increasing the number of complete time-scrambled skies beyond the current 60 would improve calibration of the background tail and simultaneously reduce uncertainty on the signal-count upper limit. The minimum attainable empirical p-value is $1/(N_{\rm null}+1)$, so a larger set of null skies would allow us to assess candidates at lower global false-positive probabilities. For independent scrambles, the standard error of the estimated mean background count scales as $N_{\rm null}^{-1/2}$. To establish any gain in faint-source recovery, we would compare search thresholds using matched injections and perform validation on held-out null campaigns.

Hough Detection Algorithm. Trajectory-based detection could recover sources that fall below the annual-cluster threshold for our DP Linker algorithm. A natural extension is a search that uses the Hough transform to detect lines [51]. For tangent-plane photon positions $\boldsymbol{\theta}_i$, the count in a reference-position cell C is

$$H_C(\boldsymbol{\mu}) = \sum_i \mathbf{1}_C[\boldsymbol{\theta}_i - \boldsymbol{\mu}(t_i - t_0)]\,.$$

Here t_0 is the reference epoch. At the source velocity, photons from a linear trajectory accumulate near one reference position, combining evidence across the full observing period without requiring an accepted cluster in each year.

U-Net CNN. Each region could be represented as an RGB image, with photon arrival time mapped to color and brightness proportional to photon density. A U-Net segmentation model [52] could then identify elongated, tubelike structures, after which the associated photons would be fitted in position and time to form trajectory proposals. Training would use injected moving-source trajectories as target masks and stationary-source or background-only fields as negative examples. Comparison with our current algorithm and Hough searches would use the same injections and null campaigns to measure recovery at the same false-positive rate.

Asteroid-Mass Primordial Black Holes. The approximate mass interval 10^{-17}–$10^{-10}\,M_\odot$ lies between strong evaporation constraints and optical microlensing constraints, and the PBH abundance in this interval remains poorly constrained observationally [53]. In particular, Subaru/HSC

microlensing constraints weaken sharply near $10^{-10}\,M_\odot$, where wave-optics and finite-source-size effects suppress the lensing signal [54]. We could extend our dressed-PBH calculation below $10^{-10}\,M_\odot$ to test whether gamma-ray emission from annihilating halos can constrain PBHs at masses beyond the reach of methods such as optical microlensing.

UCMHs. If the constituent dark matter particles of an ultracompact minihalo (UCMH) annihilate, the UCMH's dense interior could produce gamma rays [8, 55]. Its luminosity follows the density-squared integral in Eq. 15, with the UCMH profile and inner core replacing the dressed-PBH halo model. Nearby surviving UCMHs could then appear as moving sources, as proposed for microhalos [12]. Such a calculation could connect moving-source limits to the abundance of early dark matter structures and the primordial power spectrum on small scales.

Nearby-Source Searches. Extending the search to higher angular speeds would enable the discovery of closer moving sources. However, at smaller distances, the Fermi-LAT telescope's shifting position produces larger apparent source displacements relative to the background. For sufficiently nearby sources, Earth's orbit around the Sun also introduces annual parallax, which can cause deviations from our approximately linear trajectories that are assumed in the current search. We could model parallax together with the source's motion and recalibrate recovery and false-positive rates.

An interesting application is the PBH interpretation of Planet Nine, a hypothetical outer Solar System body proposed to explain the orbital clustering of trans-Neptunian objects [56]. A captured PBH surrounded by an annihilating dark matter halo could provide a gamma-ray target [57]. Searching for this target would require modeling its orbit around the Sun together with parallax from the observer's motion, then calibrating recovery and false-positive rates for these trajectories.

8. CONCLUSION

We construct an all-sky search for nearby moving gamma-ray sources using ten years of 1–100 GeV Fermi-LAT data. Our algorithm identifies photon clusters within annual time bins and links them into candidate trajectories using dynamic programming. We then use matched time-scrambled skies to characterize the background and synthetic moving-source injections to calibrate the search's recovery efficiency. The observed candidate population shows no significant moving-source excess. We derive an upper limit of 80.8 expected detected signal tracks at 90% confidence. After accounting for the calibrated recovery probability, this corresponds to a central upper limit of 122 steady sources over the full sky at $F_\gamma = 7.08 \times 10^{-10}\,\mathrm{ph\,cm^{-2}\,s^{-1}}$ and $\mu = 0.5^\circ\,\mathrm{yr^{-1}}$ at 90% confidence.

Researchers can apply our response by combining source luminosity and velocity distributions from their own models to derive constraints on moving gamma-ray populations. Since our search selects sources by flux and proper motion, its effective survey volume does not vary monotonically with source luminosity. The resulting population limits can therefore tighten and then weaken as luminosity increases.

Although this initial search finds no significant excess, it presents an entirely new search strategy that currently probes only a restricted region of signal space. Future searches spanning additional energies and motions could therefore uncover moving gamma-ray sources beyond the reach of the present analysis, while larger null ensembles would improve the calibration of their statistical significance.

APPENDIX

A. HIGHEST-RANKED CANDIDATES

Table 2 lists the 15 candidates with the largest summed track scores $S(T) = \sum_{j\in T} S_{\mathrm{LM},j}$, in decreasing order. Their global empirical probabilities compare each score S with the maximum score in every complete scrambled sky:

$$p_{\mathrm{global}}(S) = \frac{1 + \sum_{k=1}^{60} \mathbf{1}\left(S^{\mathrm{max}}_{\mathrm{null},k} \geq S\right)}{61}. \tag{17}$$

Comparing each candidate's score with the maximum score in every null sky accounts for the search across the full sky. The smallest attainable empirical probability is 1/61.

ID	$(\alpha,\delta)_{\mathrm{start}}$	$(\alpha,\delta)_{\mathrm{end}}$	Linked year bins	$\langle v\rangle$ ($^\circ\,\mathrm{yr}^{-1}$)	p_{global}
ROI301-C001	$(84.96, -7.43)$	$(83.71, -5.49)$	$1, 2, 3, \mathbf{4}, 6, 7, 8, 10$	0.255	0.180
ROI301-C003	$(84.96, -7.43)$	$(86.98, -8.94)$	$1, 2, 3, 4, 6, 7, 9$	0.313	0.492
ROI280-C001	$(86.32, -1.07)$	$(86.80, 0.329)$	$1, 2, 3, 4, 5, 6, 7, 8, 9, 10$	0.164	0.557
ROI301-C002	$(84.96, -7.43)$	$(83.71, -5.49)$	$1, 2, 3, \mathbf{5}, 6, 7, 8, 10$	0.255	0.672
ROI211-C004	$(308.79, 39.59)$	$(308.94, 38.66)$	$5, 6, 8, 9, 10$	0.188	0.967
ROI266-C001	$(313.84, 31.58)$	$(314.30, 31.27)$	$1, 2, 4, 5, 6, 7, 8, 9, 10$	0.056	0.984
ROI232-C003	$(312.20, 37.26)$	$(308.94, 38.66)$	$4, 6, 8, 9, 10$	0.487	0.984
ROI280-C002	$(86.40, -1.82)$	$(88.04, 1.35)$	$1, 2, 3, 4, 5, 6, 7, 8, 10$	0.395	0.984
ROI298-C001	$(53.30, 31.12)$	$(54.50, 31.15)$	$1, 3, 4, 6, 7, 8, 9, 10$	0.114	0.984
ROI280-C003	$(86.32, -1.07)$	$(85.59, -2.10)$	$1, 2, 3, 4, 5, 6, 7, 9$	0.159	0.984
ROI238-C001	$(91.48, 16.20)$	$(91.30, 14.72)$	$1, 3, 4, 6, 7, 8, 9, 10$	0.165	0.984
ROI293-C001	$(80.29, 8.03)$	$(81.96, 6.47)$	$1, 2, 3, 4, 5, 7, 8, 9, 10$	0.253	0.984
ROI267-C001	$(91.92, -7.12)$	$(91.47, -5.72)$	$1, 2, 3, 4, 5, 6, 7, 9$	0.184	1.00
ROI285-C001	$(63.64, 28.31)$	$(66.44, 26.88)$	$1, 2, 3, 4, 5, 6, 7, 9, 10$	0.319	1.00
ROI289-C001	$(286.38, -37.22)$	$(285.70, -37.28)$	$1, 3, 4, 5, 6, 8, 9$	0.068	1.00

Table 2. The 15 highest-ranked moving-track candidates. IDs identify the region of interest and the candidate within that region. Global probabilities are defined by Eq. 17, with the smallest p-value being 0.180, meaning no candidate is individually significant.

B. DERIVATIONS OF SOURCE MOTION, DETECTION EFFICIENCY, AND POPULATION LIMITS

B.1. *Effects of Motion on Source Shape and Trajectory*

For the constant detected rate and time-independent PSF assumed in Eq. 2, we can write the reconstructed photon direction as $\boldsymbol{\Theta} = \boldsymbol{\theta}_0 + \boldsymbol{\mu}\tau + \boldsymbol{\eta}$. The arrival time is then uniform over $\tau \in [-\Delta t/2, \Delta t/2]$, and the centered PSF error $\boldsymbol{\eta}$ is independent of τ. The law of total covariance gives

$$\begin{aligned}\mathrm{Cov}(\boldsymbol{\Theta}) &= \mathbb{E}[\mathrm{Cov}(\boldsymbol{\Theta} \mid \tau)] + \mathrm{Cov}[\mathbb{E}(\boldsymbol{\Theta} \mid \tau)] \\ &= \mathbf{C}_{\mathrm{PSF}} + \frac{\Delta t^2}{12}\boldsymbol{\mu}\boldsymbol{\mu}^\top.\end{aligned}$$

Setting $\mathbf{C}_{\mathrm{PSF}} = \sigma_E^2\mathbf{I}$ yields Eq. 3.

For rectilinear space motion, let $\boldsymbol{x} = r\hat{\boldsymbol{n}}$ have constant velocity $\boldsymbol{v}$. Define the radial component $v_r = \boldsymbol{v}\cdot\hat{\boldsymbol{n}}$ and the tangent-plane projector $\mathbf{P} = \mathbf{I} - \hat{\boldsymbol{n}}\hat{\boldsymbol{n}}^\top$. Since $\dot{r} = v_r$, differentiating the direction

gives $\dot{\hat{\boldsymbol{n}}} = \frac{\mathbf{P}\boldsymbol{v}}{r}$. Its magnitude is $\mu = v_\perp/r$. Using $\dot{v}_r = v_\perp^2/r$, a second differentiation yields

$$\ddot{\hat{\boldsymbol{n}}} = -\mu^2 \hat{\boldsymbol{n}} - \frac{2v_r}{r}\dot{\hat{\boldsymbol{n}}}.$$

The normal term preserves unit direction, while the tangent component is the perspective acceleration,

$$\mathbf{P}\ddot{\hat{\boldsymbol{n}}} = -\frac{2v_r}{r}\dot{\hat{\boldsymbol{n}}}.$$

Over a short interval T, the resulting displacement from a linear angular trajectory is approximately $|v_r|\mu T^2/r$, which is small relative to μT when $|v_r|T/r \ll 1$ [58].

B.2. *Poisson Count Fluctuations and Detection Efficiency*

Let the annual counts be independent $N_t \sim \text{Poisson}(\lambda_t)$ for $t = 1, \ldots, 10$, with total mean $\Lambda = \sum_t \lambda_t = 10\lambda$. Their probability-generating functions multiply to give

$$\mathbb{E}[z^{\sum_t N_t}] = \prod_{t=1}^{10} \exp[\lambda_t(z-1)] = \exp[\Lambda(z-1)].$$

The total $K = \sum_t N_t$ is therefore Poisson with mean Λ, even when the annual means differ. For annual fractions $a_t = \lambda_t/\Lambda$ and counts satisfying $\sum_t n_t = K$, the conditional allocation follows by dividing the joint Poisson mass by the total-count mass:

$$\Pr(\boldsymbol{N} = \boldsymbol{n} \mid K) = \frac{K!}{\prod_t n_t!} \prod_{t=1}^{10} a_t^{n_t}.$$

Let $h(\boldsymbol{n})$ be the recovery probability for annual count vector $\boldsymbol{n}$ at fixed source and host conditions, averaged over the other calibrated injection variables. The exact response is

$$\epsilon_{\text{exact}} = \sum_{K=0}^{\infty} \frac{e^{-\Lambda}\Lambda^K}{K!} \underbrace{\mathbb{E}_{\boldsymbol{N}|K}[h(\boldsymbol{N})]}_{H_K(\boldsymbol{a})}.$$

The scalar response of Eq. 10 replaces the allocation-averaged recovery $H_K(\boldsymbol{a})$ with the response inferred from equal-count injections, continued at the boundaries as in Eq. 11. Its approximation error is

$$\epsilon_{\text{exact}} - \epsilon_{\text{P}} = \sum_{K=0}^{\infty} \frac{e^{-\Lambda}\Lambda^K}{K!} \left[H_K(\boldsymbol{a}) - \widetilde{\epsilon}_{\text{ext}}(K/10, \mu, b)\right].$$

This difference isolates the effects of annual allocation and the response-boundary prescription. The 120 Poissonized injections in Section 4.4 test the aggregate error over sampled hosts. For independent recoveries $Y_i \in \{0, 1\}$ with predicted probabilities p_i, the Poisson-binomial residual is

$$z = \frac{\sum_i (Y_i - p_i)}{\sqrt{\sum_i p_i(1-p_i)}}.$$

B.3. *Existence and Uniqueness of the Signal Upper Limit*

For integer $k \geq 0$ and signal mean $s > 0$, differentiating the Poisson cumulative probability and cancelling adjacent terms gives

$$\frac{\partial}{\partial s} F_{\text{Pois}}(k; s) = \frac{\partial}{\partial s}\left(e^{-s}\sum_{j=0}^{k}\frac{s^j}{j!}\right) = -e^{-s}\frac{s^k}{k!}.$$

For a background sample $\mathcal{B} = \{N_b\}_{b=1}^{N_{\rm null}}$, terms with $n - N_b < 0$ vanish, so the empirical construction satisfies

$$\frac{\partial G_{\mathcal{B}}(n;s)}{\partial s} = -\frac{1}{N_{\rm null}+1} \sum_{b:N_b \le n} e^{-s} \frac{s^{n-N_b}}{(n-N_b)!}.$$

The function $G_{\mathcal{B}}(n;s)$ is continuous and strictly decreasing for $s > 0$ whenever at least one null count satisfies $N_b \le n$. The plus-one term sets its asymptotic floor,

$$\lim_{s\to\infty} G_{\mathcal{B}}(n;s) = \frac{1}{N_{\rm null}+1}.$$

Consequently, a unique positive solution to $G_{\mathcal{B}}(n;s) = \zeta$ exists if and only if

$$\frac{1}{N_{\rm null}+1} < \zeta < G_{\mathcal{B}}(n;0).$$

The observed count satisfies this condition for $N_{\rm null} = 60$ and $\zeta = 0.10$. When $G_{\mathcal{B}}(n;0) \le \zeta$, the prescribed endpoint is zero. The plus-one floor also prevents finite inversion at arbitrarily small ζ. For $s > 0$ at the adopted confidence level, monotonicity gives $s \le \widehat{s}_{1-\zeta}(n)$ if and only if $G_{\mathcal{B}}(n;s) \ge \zeta$, establishing the equivalence used in Eq. 6.

B.4. *Source Velocities and Effective Survey Volume*

Projecting the heliocentric velocity $\boldsymbol{u} - \boldsymbol{v}_\odot$, where $\boldsymbol{u} \sim \mathcal{N}(\mathbf{0}, \sigma_v^2 \mathbf{I})$, onto an orthonormal tangent basis preserves component variance σ_v^2. The transverse mean has magnitude ν, with

$$\nu^2(\hat{\boldsymbol{n}}) = |\boldsymbol{v}_\odot|^2 - (\boldsymbol{v}_\odot \cdot \hat{\boldsymbol{n}})^2.$$

Choose the first tangent axis along the projected mean. Integration over polar angle, including the Jacobian $v_\perp$, yields

$$\begin{aligned} p(v_\perp \mid \hat{\boldsymbol{n}}) &= \frac{v_\perp}{2\pi\sigma_v^2} e^{-(v_\perp^2+\nu^2)/(2\sigma_v^2)} \int_0^{2\pi} e^{v_\perp \nu \cos\varphi/\sigma_v^2}\, d\varphi \\ &= \frac{v_\perp}{\sigma_v^2} e^{-(v_\perp^2+\nu^2)/(2\sigma_v^2)} I_0\!\left(\frac{v_\perp \nu}{\sigma_v^2}\right). \end{aligned}$$

Here $I_0(x) = (2\pi)^{-1} \int_0^{2\pi} e^{x\cos\varphi}\, d\varphi$. The density is normalized as a radial Gaussian marginal and reduces to the Rayleigh distribution when $\nu = 0$.

The numerical volume integral is evaluated over angular speed. At fixed distance, $v_\perp = r\mu$ and $dv_\perp = r\, d\mu$, so the change of variable gives

$$\begin{aligned} V_{\rm eff}(L_\gamma) = \int d\Omega \int_0^\infty r^3\, dr \int_{\mu_{\rm min}}^{\mu_{\rm max}} d\mu\, p(r\mu \mid \hat{\boldsymbol{n}}) \\ \times \epsilon_{\rm P}\!\left(\frac{L_\gamma \mathcal{E}(\hat{\boldsymbol{n}})}{4\pi r^2 T_{\rm yr}}, \mu, b\right). \end{aligned}$$

Here μ is in radians per matching time unit; the factor r is the velocity Jacobian, and the response enforces the count support.

B.5. *Annihilation Integral for the Capped Density Spike*

The central dressed-PBH profile reaches the annihilation plateau at r_c, defined by $\rho_{\rm sp}(r_c) = \rho_{\rm ann}$. For the $r^{-9/4}$ spike,

$$r_c = r_{\rm ta}\left(\frac{\rho_{\rm eq}}{2\rho_{\rm ann}}\right)^{4/9}. \tag{18}$$

The surviving profile is

$$\rho(r) = \begin{cases} \rho_{\rm ann}, & r_{\rm min} \le r \le r_c, \\ \rho_{\rm ann}(r_c/r)^{9/4}, & r_c < r \le r_{\rm ta}, \\ 0, & \text{otherwise.} \end{cases}$$

Throughout the evaluated benchmark, $r_{\rm min} = 4r_s < r_c < r_{\rm ta}$. Splitting the annihilation integral $J = \int \rho^2\, d^3r$ at the core radius gives

$$\begin{aligned} J &= 4\pi\rho_{\rm ann}^2 \left[\int_{r_{\rm min}}^{r_c} r^2\, dr + r_c^{9/2}\int_{r_c}^{r_{\rm ta}} r^{-5/2}\, dr\right] \\ &= 4\pi\rho_{\rm ann}^2 r_c^3 \left\{\frac{1-(r_{\rm min}/r_c)^3}{3} + \frac{2}{3}\left[1-(r_c/r_{\rm ta})^{3/2}\right]\right\}. \end{aligned}$$

At fixed equality, particle, and age parameters, $r_c/r_{\rm ta}$ is constant and $r_c^3 \propto r_{\rm ta}^3 \propto M_{\rm PBH}$. Hence $J \propto M_{\rm PBH}$, up to the inner-cutoff correction. Substitution into Eq. 15 yields

$$L_\gamma \propto M_{\rm PBH}\left[1 + O\Big((r_{\rm min}/r_c)^3\Big)\right].$$

REFERENCES

[1] N. Aghanim *et al.* (Planck Collaboration), "Planck 2018 results. VI. Cosmological parameters," *Astron. Astrophys.* **641**, A6 (2020) [arXiv:1807.06209].

[2] G. Bertone, D. Hooper, and J. Silk, "Particle dark matter: Evidence, candidates and constraints," *Phys. Rep.* **405**, 279 (2005) [hep-ph/0404175].

[3] L. Baudis and S. Profumo, "Dark Matter," in S. Navas *et al.* (Particle Data Group), *Review of Particle Physics*, *Phys. Rev. D* **110**, 030001 (2024), and 2025 update, https://pdg.lbl.gov/2025/reviews/rpp2025-rev-dark-matter.pdf.

[4] A. M. Green and B. J. Kavanagh, "Primordial black holes as a dark matter candidate," *J. Phys. G* **48**, 043001 (2021) [arXiv:2007.10722].

[5] B. J. Carr, K. Kohri, Y. Sendouda, and J. Yokoyama, "New cosmological constraints on primordial black holes," *Phys. Rev. D* **81**, 104019 (2010) [arXiv:0912.5297].

[6] A. M. Green, S. Hofmann, and D. J. Schwarz, "The power spectrum of SUSY-CDM on subgalactic scales," *Mon. Not. R. Astron. Soc.* **353**, L23 (2004) [astro-ph/0309621].

[7] T. Bringmann, "Particle models and the small-scale structure of dark matter," *New J. Phys.* **11**, 105027 (2009) [arXiv:0903.0189].

[8] T. Bringmann, P. Scott, and Y. Akrami, "Improved constraints on the primordial power spectrum at small scales from ultracompact minihalos," *Phys. Rev. D* **85**, 125027 (2012) [arXiv:1110.2484].

[9] Y. N. Eroshenko, "Dark matter density spikes around primordial black holes," *Astron. Lett.* **42**, 347 (2016) [arXiv:1607.00612].

[10] W. B. Atwood *et al.* (Fermi-LAT Collaboration), "The Large Area Telescope on the Fermi Gamma-ray Space Telescope mission," *Astrophys. J.* **697**, 1071 (2009) [arXiv:0902.1089].

[11] M. Ackermann *et al.* (Fermi-LAT Collaboration), "The Fermi Large Area Telescope on orbit: event classification, instrument response functions, and calibration," *Astrophys. J. Suppl.* **203**, 4 (2012) [arXiv:1206.1896].

[12] S. M. Koushiappas, "Proper motion of gamma-rays from microhalo sources," *Phys. Rev. Lett.* **97**, 191301 (2006) [astro-ph/0606208].

[13] A. Geringer-Sameth and S. M. Koushiappas, "Detecting unresolved moving sources in a diffuse background," *Mon. Not. R. Astron. Soc.* **425**, 862 (2012) [arXiv:1012.1873].

[14] M. Ackermann *et al.* (Fermi-LAT Collaboration), "Search for Gamma-Ray Emission from Local Primordial Black Holes with the Fermi Large Area Telescope," *Astrophys. J.* **857**, 49 (2018) [arXiv:1802.00100].

[15] F. Acero *et al.* (Fermi-LAT Collaboration), "Fermi Large Area Telescope Third Source Catalog," *Astrophys. J. Suppl.* **218**, 23 (2015) [arXiv:1501.02003].

[16] S. Abdollahi *et al.* (Fermi-LAT Collaboration), "Incremental Fermi Large Area Telescope Fourth Source Catalog (4FGL-DR3)," *Astrophys. J. Suppl.* **260**, 53 (2022) [arXiv:2201.11184].

[17] J. Ballet, P. Bruel, T. H. Burnett, and B. Lott (Fermi-LAT Collaboration), "Fermi-LAT 16-year Source List" (2026) [arXiv:2602.22148].

[18] D. Baxter *et al.*, "Recommended conventions for reporting results from direct dark matter searches," *Eur. Phys. J. C* **81**, 907 (2021) [arXiv:2105.00599].

[19] Fermi Science Support Center, "Time in Fermi data analysis", https://fermi.gsfc.nasa.gov/ssc/data/p7v6/analysis/documentation/Cicerone/Cicerone_Data/Time_in_ScienceTools.html.

[20] S. Abdollahi *et al.* (Fermi-LAT Collaboration), "Fermi Large Area Telescope Fourth Source Catalog," *Astrophys. J. Suppl.* **247**, 33 (2020), https://doi.org/10.3847/1538-4365/ab6bcb.

[21] Fermi Science Support Center, "LAT data selection recommendations", https://fermi.gsfc.nasa.gov/ssc/data/analysis/scitools/lat_data_selection.html.

[22] Á. González, "Measurement of areas on a sphere using Fibonacci and latitude–longitude lattices," *Math. Geosci.* **42**, 49 (2010), https://doi.org/10.1007/s11004-009-9257-x.

[23] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Proceedings of KDD-96*, pp. 226–231 (1996), https://cdn.aaai.org/KDD/1996/KDD96-037.pdf.

[24] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.* **12**, 2825 (2011), https://jmlr.org/papers/v12/pedregosa11a.html.

[25] T.-P. Li and Y.-Q. Ma, "Analysis methods for results in gamma-ray astronomy," *Astrophys. J.* **272**, 317 (1983), https://doi.org/10.1086/161295.

[26] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *J. Am. Stat. Assoc.* **102**, 359 (2007), https://doi.org/10.1198/016214506000001437.

[27] C. Czado, T. Gneiting, and L. Held, "Predictive model assessment for count data," *Biometrics* **65**, 1254 (2009), https://doi.org/10.1111/j.1541-0420.2009.01191.x.

[28] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. (Springer, New York, 2009), Sec. 7.10, https://doi.org/10.1007/978-0-387-84858-7.

[29] B. Phipson and G. K. Smyth, "Permutation p-values should never be zero: calculating exact p-values when permutations are randomly drawn," *Stat. Appl. Genet. Mol. Biol.* **9**, 39 (2010), https://doi.org/10.2202/1544-6115.1585.

[30] B. Efron, "Bootstrap methods: Another look at the jackknife," *Ann. Stat.* **7**, 1–26 (1979), https://doi.org/10.1214/aos/1176344552.

[31] S. Holm, "A simple sequentially rejective multiple test procedure," *Scand. J. Stat.* **6**, 65 (1979), https://www.jstor.org/stable/4615733.

[32] P. Bruel, T. H. Burnett, S. W. Digel, G. Jóhannesson, N. Omodei, and M. Wood (Fermi-LAT Collaboration), "Fermi-LAT improved Pass 8 event selection" (2018) [arXiv:1810.11394].

[33] I. King, "The structure of star clusters. I. An empirical density law," *Astron. J.* **67**, 471 (1962), https://doi.org/10.1086/108756.

[34] M. Ayer, H. D. Brunk, G. M. Ewing, W. T. Reid, and E. Silverman, "An empirical distribution function for sampling with incomplete information," *Ann. Math. Stat.* **26**, 641 (1955), https://doi.org/10.1214/aoms/1177728423.

[35] L. D. Brown, T. T. Cai, and A. DasGupta, "Interval estimation for a binomial proportion," *Stat. Sci.* **16**, 101 (2001), https://doi.org/10.1214/ss/1009213286.

[36] Fermi Science Support Development Team, "Fermitools: Fermi Science Tools," Astrophysics Source Code Library, ascl:1905.011 (2019), https://ascl.net/1905.011.

[37] Q. McNemar, "Note on the sampling error of the difference between correlated proportions or percentages," *Psychometrika* **12**, 153 (1947), https://doi.org/10.1007/BF02295996.

[38] Fermi-LAT Collaboration, "Map of the 16-year LAT detection threshold" (2026), https://fermi.gsfc.nasa.gov/ssc/data/access/lat/fl16y/detthresh_P8R3_16years_PL22.fits.

[39] P. J. McMillan, "The mass distribution and gravitational potential of the Milky Way," *Mon. Not. R. Astron. Soc.* **465**, 76 (2017) [arXiv:1608.00971].

[40] R. Schönrich, J. Binney, and W. Dehnen, "Local kinematics and the local standard of rest," *Mon. Not. R. Astron. Soc.* **403**, 1829 (2010) [arXiv:0912.3693].

[41] S. O. Rice, "Mathematical analysis of random noise," *Bell Syst. Tech. J.* **23**, 282 (1944), https://doi.org/10.1002/j.1538-7305.1944.tb00874.x.

[42] P. F. de Salas and A. Widmark, "Dark matter local density determination: recent observations and future prospects," *Rep. Prog. Phys.* **84**, 104901 (2021) [arXiv:2012.11477].

[43] J. Adamek, C. T. Byrnes, M. Gosenca, and S. Hotchkiss, "WIMPs and stellar-mass primordial black holes are incompatible," *Phys. Rev. D* **100**, 023506 (2019) [arXiv:1901.08528].

[44] B. Carr, F. Kühnel, and L. Visinelli, "Black holes and WIMPs: All or nothing or something else," *Mon. Not. R. Astron. Soc.* **506**, 3648 (2021) [arXiv:2011.01930].

[45] P. Gondolo and J. Silk, "Dark matter annihilation at the galactic center," *Phys. Rev. Lett.* **83**, 1719 (1999) [astro-ph/9906391].

[46] M. Cirelli *et al.*, "PPPC 4 DM ID: A poor particle physicist cookbook for dark matter indirect detection," *JCAP* **03**, 051 (2011); **10**, E01 (2012) [arXiv:1012.4515].

[47] F. Aharonian *et al.* (H.E.S.S. Collaboration), "Search for the evaporation of primordial black holes with H.E.S.S.," *JCAP* **04**, 040 (2023) [arXiv:2303.12855].

[48] A. Albert *et al.* (HAWC Collaboration), "Constraining the local burst rate density of primordial black holes with HAWC," *JCAP* **04**, 026 (2020) [arXiv:1911.04356].

[49] E. Bugaev, P. Klimai, and V. Petkov, "Photon spectra from final stages of a primordial black hole evaporation in different theoretical models," in *Proceedings of the 30th International Cosmic Ray Conference* **3**, 1123 (2008) [arXiv:0706.3778].

[50] V. B. Petkov *et al.*, "Searching for very-high-energy gamma-ray bursts from evaporating primordial black holes," *Astron. Lett.* **34**, 509 (2008) [arXiv:0808.3093].

[51] R. O. Duda and P. E. Hart, "Use of the Hough transformation to detect lines and curves in pictures," *Commun. ACM* **15**, 11 (1972), https://doi.org/10.1145/361237.361242.

[52] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Lect. Notes Comput. Sci. **9351**, 234 (2015) [arXiv:1505.04597].

[53] B. Carr, A. J. Iovino, G. Perna, V. Vaskonen, and H. Veermäe, "Primordial black holes: constraints, potential evidence and prospects," *Riv. Nuovo Cim.* (2026) [arXiv:2601.06024].

[54] N. Smyth, S. Profumo, S. English, T. Jeltema, K. McKinnon, and P. Guhathakurta, "Updated constraints on asteroid-mass primordial black holes as dark matter," *Phys. Rev. D* **101**, 063005 (2020) [arXiv:1910.01285].

[55] P. Scott and S. Sivertsson, "Gamma-rays from ultracompact primordial dark matter minihalos," *Phys. Rev. Lett.* **103**, 211301 (2009); **105**, 119902(E) (2010) [arXiv:0908.4082].

[56] K. Batygin, F. C. Adams, M. E. Brown, and J. C. Becker, "The Planet Nine Hypothesis," *Phys. Rep.* **805**, 1 (2019) [arXiv:1902.10103].

[57] J. Scholtz and J. Unwin, "What if Planet 9 is a primordial black hole?" *Phys. Rev. Lett.* **125**, 051103 (2020) [arXiv:1909.11090].

[58] L. Lindegren and D. Dravins, "Astrometric radial velocities for nearby stars," *Astron. Astrophys.* **652**, A45 (2021), https://doi.org/10.1051/0004-6361/202141344.