

First observations of solar disk gamma rays over a full solar cycle

Tim Linden,^{1,2,*} John F. Beacom^{2,3,4,†} Annika H. G. Peter,^{2,3,4,‡} Benjamin J. Buckman,^{2,3,§}
Bei Zhou^{2,3,5,||} and Guanying Zhu^{2,3,¶}

¹*Stockholm University and The Oskar Klein Centre for Cosmoparticle Physics,
AlbaNova, 10691 Stockholm, Sweden*

²*Center for Cosmology and AstroParticle Physics (CCAPP), Ohio State University,
Columbus, Ohio 43210, USA*

³*Department of Physics, Ohio State University, Columbus, Ohio 43210, USA*

⁴*Department of Astronomy, Ohio State University, Columbus, Ohio 43210, USA*

⁵*Department of Physics and Astronomy, Johns Hopkins University, Baltimore, Maryland 21218, USA*



(Received 31 December 2020; accepted 18 February 2022; published 21 March 2022)

The solar disk is among the brightest γ -ray sources in the sky. It is also among the most mysterious. No existing model fully explains the luminosity, spectrum, time variability, and morphology of its emission. We perform the first analysis of solar disk γ -rays over a full 11 year solar cycle, utilizing a powerful new method to differentiate solar signals from astrophysical backgrounds. We produce: (i) a robustly measured spectrum from 100 MeV to 100 GeV, reaching a precision of several percent in the 1–10 GeV range, (ii) new results on the anticorrelation between solar activity and γ -ray emission, (iii) strong constraints on short-timescale variability, ranging from hours to years, and (iv) new detections of the equatorial and polar morphologies of high-energy γ -rays. We find no significant energy dependence in the time variability of solar disk emission, indicating that strong magnetic-field effects close to the solar surface, rather than modulation throughout the heliosphere, must primarily control the flux and morphology of solar disk emission.

DOI: [10.1103/PhysRevD.105.063013](https://doi.org/10.1103/PhysRevD.105.063013)

I. INTRODUCTION

The Sun is a special astrophysical source. Its close proximity allows detailed studies critical to understanding other stars. The ability to spatially resolve solar emission is especially important for probing high-energy nonthermal processes, which can be highly local. These processes reveal charged-particle acceleration and interactions in the Sun's complex dynamic magnetic fields. In addition, the “space weather” induced by these processes affects Earth's atmosphere and our technological infrastructure, giving these studies practical as well as scientific importance [1].

The highest-energy processes are revealed by γ -ray observations up to ~ 200 GeV, which correspond to charged parent particles at $\sim 10\times$ higher energies.

Three processes produce GeV-range solar γ -rays. Solar flares—rapid and well-localized ejections of plasma from the solar surface—can accelerate charged particles, producing γ -rays up to a few GeV [2–7]. These are easy to separate from other sources. Higher-energy γ -rays are produced through passive bombardment by cosmic rays. Cosmic-ray electrons (and positrons) undergo inverse-Compton scattering (ICS) with solar photons, producing a γ -ray halo around the Sun [8–10]. Cosmic-ray protons (and nuclei) undergo hadronic interactions with matter in (and under) the solar photosphere, producing a bright disk [10–13]. For the halo, hadronic cosmic rays are irrelevant due to their small γ -ray production cross sections; for the disk, leptonic cosmic rays are irrelevant due to their small flux. The available angular resolution lets us resolve the disk and halo components. In this paper, we focus on the disk emission.

The Sun's γ -ray emission is dramatically affected by its magnetic fields. Without magnetic fields, the disk emission would have two components. At energies above ~ 1 GeV, the γ -ray direction increasingly follows that of the parent cosmic ray. Accordingly, only cosmic rays that graze the solar surface can interact and have the γ -rays escape [14]. The corresponding emission from the solar limb is too faint to be observed by the *Fermi* Gamma-Ray Space Telescope (*Fermi*). Near 1 GeV, there is also a “backsplash”

*linden@fysik.su.se; 0000-0001-9888-0971.

†beacom.7@osu.edu; 0000-0002-0005-2631.

‡peter.33@osu.edu; 0000-0002-8040-6785.

§buckman.12@buckeyemail.osu.edu; 0000-0002-0458-2402.

||beizhou@jhu.edu; 0000-0003-1600-8835.

¶zhu.1475@osu.edu; 0000-0003-0031-634X.

component from the whole disk, as kinematics allow low-energy γ -rays to be emitted at a large angle relative to the parent cosmic ray [15].

Of course, the Sun does have magnetic fields. Seckel, Stanev, and Gaisser (SSG) [16] hypothesized that surface fields allow emission from the full disk even at high energies. This requires the fields to deflect cosmic rays from incoming to outgoing before they interact and produce γ -rays. This requires a “Goldilocks” tuning between the solar magnetic field and gas density profiles. If cosmic rays are deflected too high in the solar atmosphere, then they will not encounter enough matter to produce γ -rays. If they are deflected too low, then they will produce only γ -rays that are pointed into, and are subsequently absorbed by, the Sun. To efficiently produce a broad γ -ray spectrum, SSG assume that cosmic rays impinging on the solar surface are funneled into magnetic flux tubes that mirror cosmic rays at the right depths. The γ -ray flux predicted by SSG [16] greatly exceeds that predicted for the solar limb [14]. Detailed simulations of cosmic-ray interactions within the solar atmosphere (including, e.g., solar composition, secondary production, coronal magnetic fields, etc.) have recently been completed by several groups [17,18].

The steady-state emission from the solar disk and halo was not detected until 2008, in a reanalysis of EGRET data [19]. In 2011, an analysis with the *Fermi* Large Aperture Telescope (*Fermi*-LAT) measured solar γ -ray emission between 100 MeV and 10 GeV, separating the disk and halo components [10]. The disk itself was not resolved, but the halo was detected out to $\sim 20^\circ$ from the Sun. Interestingly, the solar disk flux exceeded the SSG prediction by a factor of ~ 5 .

Since 2011, new studies of *Fermi*-LAT data by our group have greatly improved these observations, identifying several new features. Ng *et al.* [20] showed that the solar disk γ -ray spectrum extends to ~ 100 GeV and is significantly harder than the SSG predictions. Additionally, this paper showed surprising evidence for time variability, finding that the γ -ray flux¹ is strongly anticorrelated with solar activity. Linden *et al.* [21] took advantage of the high angular resolution of the *Fermi*-LAT above 10 GeV to perform the first resolved study of the disk, finding that the bright hard-spectrum emission observed during solar minimum is produced predominantly in the Sun’s equatorial plane, while the softer emission from polar regions is constant over the solar cycle. Tang *et al.* [22] showed that the solar-minimum spectrum scales roughly as $E^{-2.2}$ from 100 MeV to beyond 100 GeV, except for a deep dip between 30 and 50 GeV, which is unexplained.

¹Throughout this paper, we consider the γ -ray flux to be the total photon count per unit area per time on the *Fermi*-LAT instrument, e.g., photons $\text{cm}^{-2}\text{s}^{-1}$, which is sometimes also referred to (in solar physics literature) as a flux density.

In this paper, we utilize a powerful new technique to separate solar signals from astrophysical backgrounds, which allows us to significantly increase statistics and minimize systematic uncertainties. Using this model, we analyze solar disk γ -rays over a full solar cycle (2008–2020), including the first measurements below 1 GeV since 2011. This large dataset allows us to measure the energy dependence of the anticorrelation between the disk flux and solar activity. The long baseline also provides the statistics needed to search for short-period variations. Because our observations happen to start at solar minimum, we are able to compare the emission between two different minima, over which the polarity of the Sun’s dipole field has reversed. This allows us to compare the morphology and spectrum of the emission across the solar disk during periods with opposite polarities.

The outline of this paper is as follows. In Sec. II, we describe our models of the solar disk, solar ICS halo, and astrophysical backgrounds. In Sec. III, we measure the solar disk flux, spectrum, time variability, and morphology. In Sec. IV we interpret our results. In the Appendixes, we further examine our ICS halo and solar flare models. Overall, our results significantly extend the measurement of solar disk γ -ray emission at low energies, complementing our work at TeV energies with the High-Altitude Water Cherenkov Observatory (HAWC) Collaboration [23,24].

II. DATA ANALYSIS METHODOLOGY

We develop a novel sophisticated methodology to measure the solar disk γ -ray flux, improving on our previous analyses [20–22] in several ways. Most importantly, we use a data-driven model for the astrophysical γ -ray background that optimally isolates solar contributions and automatically accounts for systematic uncertainties in *Fermi* point-source determinations and diffuses background models. We utilize and compare several methods to remove the contributions of solar flares to the disk flux. Finally, we have developed and implemented detailed models for γ -ray production via the ICS solar halo, as well as the γ -ray production from the Moon. These technical improvements allow us to accurately determine the solar disk γ -ray flux down to 100 MeV.

A. *Fermi*-LAT data selection

We utilize 11.4 yr of P8_V3 *Fermi*-LAT data, from August 2008 to February 2020 (MET 239557417–603960406), including SOURCE class events from 100 MeV to 100 GeV with zenith angles $< 90^\circ$. P8_V3 events represent the latest event reconstruction algorithm (the conversion from electronic readout to γ -ray event information), while SOURCE class events represent the standard cuts utilized to optimize the removal of cosmic-ray contamination, while maintaining an acceptable γ -ray acceptance for analyses of individual γ -ray sources. [25].

MET (Mission Elapsed Time) is the number of seconds since January 1, 2001, which is used to count *Fermi*-LAT mission time. These cuts are placed on the data using the standard *Fermi*-LAT tools.²

We divide the data into eight equally spaced logarithmic energy bins per decade. We calculate the exposure in 10 equal bins of the instrumental coordinate ϕ , which describes the azimuthal angle surrounding the *Fermi*-LAT boresight ($\text{phibins}=10$). This choice is not typically necessary in *Fermi*-LAT studies, but is important here as the orientation of *Fermi*'s solar panels biases solar observations to regions near $\phi = 0$.

We analyze the data in helioprojective (Sun-centered, with the North pole defined by the Sun's spin axis) coordinates by separately transforming γ -ray events and the *Fermi*-LAT exposure. For events, we use the `sunpy` code [26] to convert every γ -ray into solar coordinates, utilizing the exact time of each event. For the exposure, we use the *Fermi*-LAT tools `gtltcube` and `gtexpcube2` to calculate the exposure in J2000 coordinates over small (8640-second) time bins. In each bin, the Sun moves only $\sim 0.1^\circ$ in J2000 coordinates. We convert this exposure map to helioprojective coordinates using the average time in each bin. This time binning is sufficient for our analysis, as the exposure is smooth on 0.1° scales.

B. *Fermi*-LAT solar exposure since March 2018

Solar analyses are complicated by a solar-panel malfunction that affected *Fermi* on March 16, 2018. While this issue does not significantly impact the vast majority of *Fermi* science goals, the necessity of locking the fixed solar panel onto the Sun (to power the instrument) decreases the solar exposure, as the *Fermi*-LAT instrument is inclined relative to the solar panels. For data beginning in March 2018, the monthly exposure is approximately half that recorded in previous periods. This is taken into account in our analyses of solar exposures and instrumental response functions. However, it worsens the statistical uncertainties in the most recent data.

C. Removal of solar flares

To isolate the effects from transient solar flares from the steady-state disk flux, we prepare three subsets of the data. In the first, we apply no cuts, producing a model that fits the total γ -ray flux from the solar disk, including both active emission from solar flares and passive emission from hadronic cosmic-ray interactions.

In the second, we apply a cut using the LAT Significant Flares List,³ removing events within 8640 s (2.4 hr) of a flare, the same size binning we use for exposure. We conservatively remove all events/exposure for which any

part of the bin overlaps with the full duration of the flare. This cut removes only $\sim 0.36\%$ of our total solar exposure but significantly lowers the solar flux by removing several bright flares.

In the third, we additionally apply a cut based on keV–MeV data from the *Fermi* Gamma-ray Burst Monitor (GBM), using the GBM Solar Flare List.⁴ To be conservative, we cut all γ -ray events that occur within 86400 s (1 day) of a flare from either the GBM or LAT. This removes a significant fraction (41.3%) of the solar exposure and allows us to test whether subthreshold flares impact the solar γ -ray spectrum.

D. Astrophysical background modeling

We produce a data-driven model of the astrophysical background produced by nonsolar events. The approach is simple to describe but computationally demanding. The key concept is the following. While the motion of the Sun relative to distant astrophysical sources adds significant technical complexity, it also allows the astrophysical background flux to be directly measured during periods when the Sun is not present.

To characterize the astrophysical backgrounds, we begin with the helioprojective coordinates of each γ -ray recorded by the *Fermi*-LAT, as well as the finely binned (2.4 hr) exposure files. We remove all events and exposure within 45° of the current solar position. To synchronize the event selection with our exposure calculation, we keep or remove *Fermi*-LAT events based on average right ascension and declination values of the Sun in each 2.4 hr bin. Last, we convert the counts and exposure maps back to J2000 coordinates, producing a flux map of the entire sky during periods when the Sun was far away.

This methodology significantly outperforms previous approaches. Because our model is data driven, it is insensitive to uncertainties in diffuse background modeling and point-source estimation. Additionally, our model automatically includes effects stemming from the *Fermi*-LAT angular and energy resolution. Compared to previous data-driven analyses (for example, the fake-Sun methods employed in [20–22]), this technique produces a significantly larger statistical sample, in effect using a continuum of fake Suns instead of a handful. Even more, we produce a Sun-centered skymap that contains nearly every *Fermi*-LAT γ -ray. Our eventual analysis includes 49600 photons that are assigned to the solar disk template, including 8400 above 1 GeV and 513 above 10 GeV, providing a significant statistical handle on the solar disk flux, time variability, and morphology.

Figure 1 shows the results of this method in two example energy bins, 100–316 MeV and 3.16–10 GeV. The γ -ray counts map in helioprojective coordinates is shown before

²<https://fermi.gsfc.nasa.gov/ssc/data/analysis/software/>.

³https://hesperia.gsfc.nasa.gov/fermi/lat/qlook/lat_events.txt.

⁴https://hesperia.gsfc.nasa.gov/fermi/gbm/qlook/fermi_gbm_flare_list.txt.




FIG. 1. Left: γ -ray count maps in helioprojective coordinates over 11 years of *Fermi*-LAT data (Mollweide projection). Periods with significant *Fermi*-LAT flares are removed. Results are shown for 100–316 MeV (top) and 3.16–10 GeV (bottom). The outer ring is an edge effect due to pixels that are partially in the region of interest (ROI). Right: Residual γ -ray emission maps after the empirically determined background is subtracted. The color bars apply to both left and right panels. Our background-removal methodology efficiently removes both diffuse emission (the haze in the left panels) and point sources (the streaks in the left panels), leaving only the solar emission (disk and halo components) and statistical noise.

and after background subtraction. Bright diffuse emission (with a flux nearly half that of the solar disk at low energies), as well as bright point sources (observed as tracks that streak past the Sun), are entirely removed, leaving only the solar (disk and halo) components, plus statistical noise. We note that each track corresponds to a single bright γ -ray source, with the direction of the track depending on the sources' exact motion through helioprojective coordinates.

In Appendix E we provide an additional test of our astrophysical background modeling by reproducing our full solar emission model for a source centered on the

“fake Sun” location, which trails the true location of the Sun by six months. We show that, in this case, the astrophysical background modeling accurately reproduces the observed γ -ray data, producing a measured intensity for the fake solar disk that is consistent with 0 in every energy bin.

E. Additional model considerations

There are three additional complications that affect our background methodology (and those of previous papers).

First, while our method produces a robust model for all sources that do *not* move in the J2000 coordinate system, it

produces a skewed model of any source that moves with respect to background stars. The two important sources in this class are (1) the Moon and (2) the diffuse solar ICS halo. While the Moon moves compared to the Sun, the solar ICS halo is stationary in helioprojective coordinates. Thus, the two components are treated using different techniques.

To account for the Moon's emission, we slightly amend the above analysis to additionally remove any γ -rays and instrumental exposure within 20° of the Moon in each 2.4 hr time bin. While the Sun moves only 0.1° during this time period, the Moon moves approximately 1.2° . However, this is still small compared to both the angular dependence of the *Fermi*-LAT exposure, as well as the size of the point-spread function (PSF) at the ~ 100 MeV energies where the Moon is bright. Our amended model thus produces an analysis of the astrophysical background in regions where neither the Moon nor Sun are present. We note that the removal of ROIs consistent with the solar and lunar position do not affect the expected astrophysical background flux in these regions, because both the photon counts and instrumental exposure from these sky regions is removed self-consistently. This means that the observed photon flux from the astrophysical coordinate is based on equivalent sets of counts and exposures taken while the Sun and Moon are not nearby.

However, simply removing lunar photons will not provide an accurate model for solar γ -rays during regions when the Moon is close to the Sun. In order to avoid throwing out events during such periods, we additionally produce a model for the Moon's γ -ray emission by utilizing `gtmodel`, which is a *Fermi*-LAT analysis tool designed to produce a simulated observation of a source, to produce a source with radius 0.26° at the current position of the Moon, and then convert this template into helioprojective coordinates for each of the 8640 s exposures in our model. This produces a template for lunar emission that traces the total lunar exposure at each point in helioprojective coordinates over the full observation time of our analysis. We fit the flux of the Moon component in our likelihood analysis. This enables us to avoid throwing out approximately 25% of our exposure, which is taken during periods when the Moon lies within 45° of the Sun.

Solar ICS emission is concentrated within 45° of the Sun, a region that is already removed from our analysis of the astrophysical backgrounds. We build a solar ICS emission model utilizing the `StellarICs` code [27,28], which we slightly modify to include updated local interstellar spectra for cosmic-ray electrons and positrons [29]. The modulated spectra are calculated using the standard force-field approximation [30], which assumes separability of the heliospheric diffusion coefficient into radial- and energy-dependent functions. For the radially-dependent force-field potential, we use Eq. (7) of Ref. [8], which is derived from the radial dependence of the cosmic-ray mean free path for solar cycles 20/22 [31]. The largest uncertainty in our

study, particularly at low energies, comes from our modeling of the ICS emission. We provide more details in Sec. II F.

Because the astrophysical background template is based on real data, it already includes the effects of the PSF. However, the solar disk, solar ICS halo, and Moon templates must be smeared by the *Fermi*-LAT angular resolution. The Moon's γ -ray flux is smeared using `gtmodel`. For the disk and ICS flux, we note that the PSF near the Sun may systematically differ from the average *Fermi*-LAT PSF due, e.g., to the extreme ϕ dependence of solar observations. Thus, we calculate the PSF in each 2.4 hr exposure using `gtpsf`. The solar PSF is then weighted by the solar exposure in each bin. Because the exact structure of the PSF farther from the Sun is less critical for our study, we utilize a PSF calculated at the solar position throughout the ROI. The average energy resolution of *Fermi*-LAT events ($\sim 10\%$) is smaller than the energy binning of our analysis, and the solar spectrum is relatively smooth, so we can safely neglect simulating the energy dispersion.

The second complication stems from variable γ -ray sources, which may have different luminosities when they are close to the Sun, producing residuals in the γ -ray analysis. However, this does not significantly affect our results for two reasons. First, the 11 yr lever arm of our analysis means that each source has made multiple passes near the Sun, decreasing the effects of single source flares by a factor of ~ 10 . Second, these source residuals produce bright streaks across the solar position, which are not degenerate with the spherically symmetric solar emission, as seen in Fig. 1. The combination of these factors implies that variable γ -ray sources have a very small effect on our determination of the solar disk flux.

The third complication is that the unique nature of the Sun makes it difficult to entirely eliminate systematic effects specific to the solar region. For example, because *Fermi*'s solar panels are locked onto the Sun, the instrumental phase space of solar events is unique. This potentially affects the effective area, angular resolution, and energy resolution of solar events. Moreover, radiation or cosmic-ray effects on the instrument itself cannot be entirely ruled out. Notably, our previous study [21] uncovered a systematic error in events near $\phi = 0$, which has since been corrected in the latest *Fermi*-LAT data release. Our thorough studies of systematics in Ref. [21] suggest that any remaining issues must be very subtle.

To conclude, we are confident that our methodology goes far beyond previous analyses. This allows us to utilize nearly every solar γ -ray, regardless of the solar position compared to background γ -ray sources, while simultaneously removing all background contributions with high fidelity.

F. Fitting the ICS component

Modeling the background from the solar ICS halo is especially challenging, both because it is stationary in

heliospheric coordinates and because its emission extends very close to the disk. The degeneracy between disk and halo emission is particularly bad below 1 GeV, where the *Fermi*-LAT angular resolution worsens significantly.

The uncertainty is magnified by our ignorance of the ICS morphology, which prevents us from reliably extrapolating the ICS flux from positions far from the Sun, where the disk flux is irrelevant. Instead, we test multiple models of the ICS morphology and marginalize our results over them. We begin with three “seed” models for the ICS morphology that use modulation potentials defined at 1 AU of $\Phi_0 \in \{0, 400, 1000\}$ MV. These choices liberally span the range of modulation potentials that are compatible with local cosmic-ray observations. We note that we do not intend to measure the modulation potential of the Sun in this analysis—instead we develop a model where the measured solar γ -ray flux is independent of the ICS modulation model. In a forthcoming publication [32], we will measure and analyze the ICS component in detail.

In our default analysis, we divide the ICS templates for each modulation model into seven angular components, spanning $\{0-1, 1-2, 2-4, 4-6, 6-8, 8-10, \text{and } 10-45\}^\circ$ from the solar position. These angular regions are defined *before* the emission is smoothed by the PSF. Thus, these templates represent physical emission on each angular scale. Utilizing this binning, our three distinct seed models provide different ICS angular profiles *within* each radial bin, but the floating normalizations of each independent bin allow each model to have similar global properties on angular scales much larger than a bin. Our goal, in this analysis, is to show that all three seed models produce the same disk flux, which will indicate that small scale features in the ICS morphology do not affect our results. We note that these templates are allowed to float to different best-fit values within each energy bin.

This approach has two benefits. First, it accounts for uncertainties in the ICS halo profile near the Sun due to uncertainties in solar modulation. Second, and more subtly, it disentangles uncertainties in the relative normalization of the ICS and astrophysical background far from the Sun with the ICS and solar disk components near the Sun. While the ICS surface brightness is highest within $\sim 2^\circ$ of the Sun, the large extent of the halo means that the average ICS γ -ray is produced $\sim 10-15^\circ$ from the Sun. Thus, in a model with only one ICS component, the ICS normalization would be set by the low-surface brightness emission observed far from the Sun, which is then extrapolated to calculate the ICS flux at the solar position. The effect of this extrapolation is shown in Appendix A.

We utilize several methods to control variation in the normalization of each ICS angular bin. These choices moderately affect the measured solar emission, but more significantly affect its uncertainty. In our default analysis, we conservatively allow the normalization of each ICS angular bin to float freely, demanding only that it be non-negative.

This produces robust error bars that encompass the systematic uncertainties of more fine-tuned normalization models. However, this potentially produces unphysical “saw-tooth” features in the ICS angular profile. In Appendix A, we show the results from physically motivated models that force the ICS normalization coefficients to either (1) be monotonic in angle, or (2) incur a renormalization penalty for adjustments that stray from unity.

We stress two important results. First, the modeling uncertainties affect the solar disk spectrum only below ~ 1 GeV. At higher energies, we robustly separate solar disk γ -ray emission from ICS and astrophysical emission regardless of our ICS modeling choices. Second, our choice to set extremely conservative constraints on the variation of the ICS modeling template implies that the uncertainty in our solar disk fits encompass both statistical errors, as well as systematic errors, from ICS halo mismodeling.

G. The solar disk component

Finally, we include a template for the solar disk. In our analysis, we assume a disk with a constant surface brightness out to a radius of R_\odot , which corresponds to $\sim 0.264^\circ$. We note that this is smaller than the *Fermi*-LAT angular resolution at all but the highest γ -ray energies. In Sec. III C, we investigate the solar disk γ -ray morphology, in detail, at energies above 10 GeV, where the solar disk emission can be resolved.

H. Likelihood analysis

To measure the solar disk γ -ray energy flux, we employ a template-based likelihood minimization analysis using the public code *iminuit* [33]. We utilize an ROI of radius 45° centered on the Sun, binned with a HEALPix grid of $n_{\text{side}} = 512$ (pixel size $\sim 0.11^\circ$). We carry out our analysis in 24 independent bins from 100 MeV to 100 GeV, equally spaced in logarithmic energy, so that we can measure the spectrum shape.

We fit the data with the four-component model described above: (1) a constant surface-brightness solar disk of radius 0.26° , the size of the physical disk, (2) the multicomponent ICS template, following Sec. II F, (3) the Moon template from Sec. II E, and (4) the background model described in Sec. II D. Formally, we then compute the likelihood independently in each energy bin as follows:

$$\ln(\mathcal{L}_e) = \sum_i \ln \left(\frac{\exp(-M_i) M_i^{X_i}}{X_i!} \right), \quad (1)$$

where X_i are the observed counts in each spatial pixel, and M_i represents the number of model counts given by

$$M_i = c_0 D_i + c_1 S_i + c_2 L_i + c_3 I_{0:1,i} + c_4 I_{1:2,i} + c_5 I_{2:4,i} \\ + c_6 I_{4:6,i} + c_7 I_{6:8,i} + c_8 I_{8:10,i} + c_3 I_{10:45,i}, \quad (2)$$

where c_0 through c_9 are arbitrary fitting parameters subject to limits discussed earlier in the text, D_i is the number of expected diffuse counts in each pixel, S_i is the number of solar disk counts expected in each pixel, L_i is the number of lunar events expected in each pixel, and the $I_{X:Y}$ models include the number of expected ICS counts expected in each pixel for ICS maps with physical emission between X° and Y° from the solar position.

While we allow the normalization of the astrophysical background to float, our method of constructing this template should force the normalization to unity. Indeed, we find that the normalizations of the templates vary by $<1\%$ from unity at most energies when they are allowed to float freely in each energy bin. We also find that our analysis is robust to changes in the Moon and astrophysical background models but can be affected by ICS modeling choices below 1 GeV (see Appendix A).

III. MEASUREMENTS OF THE SOLAR DISK EMISSION

Utilizing the methodology outlined above, we measure the solar disk flux and spectrum (Sec. III A), time variability (Sec. III B), and morphology (Sec. III C).

A. Flux and spectrum

Figure 2 shows the γ -ray spectrum for three different choices of solar-flare cuts. In each case, we utilize our multibinned ICS halo model, using a modulation potential of $\Phi_0 = 400$ MV at 1 AU. We show in Appendix A that this choice does not affect our results. We note several results common to each choice of solar-flare cut: (1) we produce a robust measurement of the solar disk γ -ray

spectrum down to 100 MeV, with uncertainties of 5–10% in 0.1–1 GeV, (2) we obtain excellent precision for the flux, reaching uncertainties of 2–5% in the 1–10 GeV range, and (3) we reproduce the unexpected “spectral dip” feature between 30 and 50 GeV, first identified in Ref. [22]. We note that these flux uncertainties do not include a $\sim 10\%$ uncertainty in the *Fermi*-LAT effective area, which affects all flux measurements produced from *Fermi* analyses [34].

Importantly, we stress that these results include significant freedom in the global profile of the ICS emission component (which we find is the largest systematic uncertainty in the model)—thus the error bars shown here are conservative and include both statistical and systematic contributions. Below ~ 1 GeV, we find that the systematic uncertainties from the global ICS morphology are dominant. In particular, we find that the total significance of the solar disk emission is 63σ , but this low significance is primarily due to the degeneracy with the surrounding ICS template. In models with more constrained (and theoretically motivated ICS) templates, the solar disk significance can exceed 112σ (see Appendix A).

Notably, our analysis finds that the solar spectrum is remarkably flat, following a power law very close to $E^{-2.0}$ from 100 MeV to nearly 30 GeV. This result is noteworthy and theoretically unexpected because the parent cosmic-ray proton population is observed to have an $E^{-2.7}$ spectrum across the relevant energy range. Because the outgoing γ -ray spectrum from isotropic cosmic-ray interactions is equivalent to parent proton spectrum, our results imply that the probability of a proton interacting with the solar atmosphere and producing an outgoing γ -ray increases strongly ($\sim E^{+0.7}$) with energy. The mechanism that would produce such a feature is unknown, though Ref. [17] notes




FIG. 2. Average solar disk γ -ray spectrum over the full analysis period (Aug. 2008–Feb. 2020). Results are shown for three choices of solar-flare cuts: no cuts (purple), removing high-significance *Fermi*-LAT detected flares (bold orange), and additionally removing high-significance *Fermi*-GBM flares (blue). Small offsets on the x axis are made to improve visibility. Removing *Fermi*-LAT flares is important below ~ 1 GeV. Removing *Fermi*-GBM flares is not important. The GBM-cut energy fluxes are slightly higher because the steady-state solar disk energy flux is dimmer during solar maximum, when more flares are present. The y axis is explicitly calculated using the number of photons in each bin, dividing by the width of the energy bin, and multiplying the result by the square of the average energy in the bin.

that the γ -ray spectrum can be significantly hardened if solar magnetic fields are significantly stronger than indicated by current observations and models.

Below 1 GeV, the results strongly depend on the solar-flare cuts. In the “full dataset,” which includes all solar flares, we observe a pronounced spike in the flux between 100 MeV and 1 GeV. Nearly 30% of the total solar emission below 1 GeV is produced by a handful of intense flares (see Sec. II C). More importantly, however, we obtain the same γ -ray spectrum regardless of whether we choose to remove only the brightest “LAT-detected” flares, or to more conservatively remove all “LAT- or GBM-detected” flares. This implies that subthreshold flares contribute negligibly to the total γ -ray flux. Thus, removing the LAT-detected flares (encompassing only $\sim 0.3\%$ of the total exposure) sufficiently isolates the steady-state solar disk emission, and we use that in what follows.

Focusing on the emission above 1 GeV, we note that the solar disk flux is larger during when LAT- or GBM-detected flares are removed compared to when only LAT-detected flares are removed. This is an artifact of our selection criteria. The LAT- or GBM-detected flare cut removes 41.3% of the solar exposure, predominantly during solar maximum, when the solar flare rate is highest and the steady-state disk emission is faintest [20–22]. Thus, the “GBM-detected” cut biases the 11 yr disk flux towards periods when the disk is brighter. In Appendix B, we analyze the γ -ray flux above 1 GeV in yearly bins, finding that removing GBM-detected solar flares does not affect the solar γ -ray flux during any particular phase of the solar cycle, but instead reweights the emission from different time periods. We cannot, however, rule out the interesting possibility that an underlying variable affects both the solar flare cycle and the steady state flux of the solar disk.

Focusing on our default model (with the *Fermi*-LAT flares removed), we find that the spectrum is well fit by a power law with a spectral index of $E^{-2.03 \pm 0.014}$ in the energy range between 100 MeV and 10 GeV. This simple power law provides a good fit to the γ -ray data in this energy range, with $\chi^2 = 16.6$ given 14 degrees of freedom (d.o.f.). We find evidence for a break or cutoff in the spectral index near 30 GeV, a result which has been studied in significantly more detail in Ref. [21].

Against this power law, we do note three possible spectral features. The most significant appears as a slight spectral dip at a central energy of 650 MeV. We evaluate the significance of this dip by comparing its energy flux (and energy-flux uncertainty) with the average energy flux and energy-flux uncertainty of surrounding bins. This technique is accurate because the best fit spectrum is a simple power law with an index that is nearly E^{-2} . We obtain a significance of 1.64σ , which (given the trial factors resulting from many spectral bins) is not significant. However, we note that the majority of the uncertainty in our model stems from degeneracies with the ICS component, which is

likely to be correlated between bins. Replacing our measured uncertainty with the Poisson uncertainty based on photon count, we obtain a significance of 7.14σ , which represents the maximum significance of this feature. Further study is necessary to determine the effect of correlated and uncorrelated systematic uncertainties on this possible feature. We note that two other apparent spectral features (at 1.54 GeV and 11.5 GeV, respectively) do not exceed 2σ local significance even when only Poisson uncertainties are considered.

B. Time variability

Our detailed background modeling and choice to remove only significant LAT-detected flares allows us to retain nearly 99.7% of solar disk exposure. Using this vast dataset, we can probe both long- and short-timescale variability.

1. Long-term variability over a solar cycle

Figure 3 shows the γ -ray flux in year-long bins. In each bin, we have sufficient data to refit our four-component model (astrophysical background, Moon, solar ICS halo, and solar disk), allowing us to account for variations in background sources or ICS emission. While our models fit emission over the full energy range, here we focus on two different energy scales. The first covers the range 1–10 GeV, where our measurements are the most precise; the second covers the range 0.1–1 GeV, where the effects of time-dependent solar modulation are expected to be greatest. At higher energies the time variation appears consistent with the lower-energy data, but the statistical uncertainties become significantly worse; we provide a more detailed analysis in Appendix D.




FIG. 3. Solar disk γ -ray flux for 1–10 GeV (red) and 0.1–1 GeV (blue) in year-long analysis windows beginning approximately on August 4 of each year (the last bin ends in February 2020). The points are slightly offset for clarity. We find strong evidence that the 11 yr variability, previously shown to anticorrelate with the sunspot number [20], is energy independent, which is surprising.

We find a highly significant anticorrelation between the solar disk flux and solar activity, echoing previous results [20,22]. Specifically, we find a bright peak in the γ -ray emission in 2008 (near solar minimum), which falls to a valley during solar maximum (and with no appreciable changes during the 2013 heliospheric polarity flip), before rising again in 2020 (near solar minimum). Importantly, we find that the solar disk fluxes during the prior and present solar minima are similar, despite their opposite polarities for the solar magnetic field (see Appendix D for a dedicated discussion at high energies). The larger error bars after March 2018 correspond to a decrease in *Fermi*-LAT exposure.

Figure 4 shows the low-energy solar disk γ -ray spectrum in four different temporal periods. The first roughly




FIG. 4. Top: Low-energy solar disk γ -ray spectrum in four periods, each of duration approximately three years, as marked, measured utilizing our standard analysis routine and compared to the average flux over the full observation period. For clarity, we display results using only four bins per decade. Bottom: Same as above, but the γ -ray flux has been renormalized in each time period to match the average flux in the four bins closest to 1 GeV. Our results do not indicate strong spectral differences between the solar minimum periods (represented by the first and last time periods) and the solar maximum time period (represented by the middle two temporal periods). Small offsets in the x axis on the bottom plot are to improve readability and do not correlate with shifts in the actual energy values.

corresponds to solar minimum with polarity $A < 0$, the second to solar maximum with $A < 0$, the third to solar maximum with $A > 0$, and the final to solar minimum with $A > 0$. While we have analyzed the data in each temporal period utilizing our standard energy binning of eight logarithmic bins per decade, for visual clarity we show the resulting spectra after using a weighted average to decrease this to four bins per decade.

Intriguingly, our analysis does not find any substantial spectral variations, once the overall change in the flux normalization is accounted for. We do, however, find mild changes in the 1–10 GeV energy spectrum, which are of moderate statistical significance due to the extreme statistical precision of our *Fermi*-LAT measurements. In the first time period we obtain a best-fit power-law spectrum of $\alpha = -2.048 \pm 0.018$, in the second $\alpha = -1.922 \pm 0.038$, in the third $\alpha = -2.016 \pm 0.022$, and in the final bin we obtain $\alpha = -2.076 \pm 0.024$. When the errors from each analysis are added in quadrature, we obtain a peak difference between the second and fourth temporal bins of 3.4σ . The second and fourth bins are both well fit by power laws with $\chi^2/\text{d.o.f.}$ much smaller than 1, while the fits for the first and third temporal bins have $\chi^2/\text{d.o.f.}$ that are slightly higher than one, owing to their statistical downturns in the highest energy bin.

We note that a time-variable spectrum, especially in the lowest energy bins, may have been expected—because the degree of cosmic-ray solar modulation near the Earth is a strongly energy-dependent phenomenon [35]. While we do see some evidence of spectral hardening in the second and third temporal bins (consistent with solar maximum), the amplitude of this effect is only at the $\sim 10\%$ level in the lowest (100 MeV) energy bin. The lack of a strong feature, especially at the low-energy range, implies that magnetic fields in the solar atmosphere, rather than the solar wind or cosmic-ray modulation far from the Sun, dominate the time variability of the solar disk emission.

2. Monthly variability

Figure 5 shows the time variability of the 1–10 GeV flux over even smaller time ranges of approximately 3×10^6 seconds (34.7 days) per bin. In these small time bins, the quantity of data is insufficient to precisely remove the degeneracy between solar disk emission and the ICS emission in the inner ($0-1^\circ$) angular bin. Thus, in this subsection, we fix the morphology of the ICS emission template (the relative flux in each ICS bin) to its best-fit value from our 11.4 yr analysis. We do, however, allow the normalization of this ICS component, along with the normalizations of the solar disk, Moon, and astrophysical background, to float independently in each energy and temporal bin. In our main analysis, we find that the ICS morphology is only slightly degenerate with the best-fit disk flux, indicating that this choice should only marginally affect the results here.




FIG. 5. Total solar disk γ -ray flux between 1 and 10 GeV (where the flux is best measured) in ~ 35 day (3 Ms) analysis windows. The 11 year periodicity of the solar disk signal is clearly visible, but no significant periodicity is visually obvious on smaller timescales. The yearly best-fit fluxes from Fig. 3 are also shown for comparison.

We typically obtain $>10\sigma$ detections of solar γ -ray emission in each time bin. We recover the 11 yr cycle found in Fig. 3. However, there are no visually obvious variations on the shorter timescales probed here. This implies that the short duration magnetic outbursts (e.g., solar flares, coronal mass ejections) are unlikely to determine the solar disk γ -ray flux.

3. High-frequency variability

Finally, we utilize the precise timing information of each solar γ -ray event to search for high-frequency periodic variations in the flux. Such signals might provide important evidence into the regions of the Sun that produce bright solar emissions. Additionally, this analysis provides an additional handle to test the contribution of solar flares to the steady-state solar disk γ -ray data.

One question of particular interest is whether certain regions of the Sun are more important than others for γ -ray production. Due to the solar rotation, this may manifest as a periodicity with a period of approximately 27 days (1 Bartels rotation, which characterizes the Sun's apparent rotation as seen from the moving Earth). The large *Fermi*-LAT dataset would, in principle, make it possible to fold the *Fermi*-LAT counts and exposure into the solar-phase space. However, because the solar atmosphere is not solid (polar regions rotate more slowly than equatorial regions), it is difficult to fold multiple solar rotations without smearing out any transient signal.

To maximize the time resolution of our analysis, in this subsection we employ a different analysis technique. We begin with the γ -ray data above 1 GeV, binned into 8640 s periods, which corresponds to our choice of exposure binning. We select all events observed within 2° of the solar position, noting that the average 68% confidence

interval of the PSF for γ -rays above 1 GeV is always smaller than 1° . This data sample then contains the sum of all emission within 2° of the Sun coming from the solar disk, solar ICS halo, Moon, and astrophysical background.

To isolate the variability of the solar disk, we note the following: (1) we can safely ignore the flux contribution from the Moon, because the 1 GeV minimum energy cutoff removes the vast majority of its γ -ray emission, (2) we assume that the ICS emission is relatively time invariant, moreover, above 1 GeV; the emission from the solar disk is significantly larger than the ICS flux within $\sim 2^\circ$ of the solar position [10].

We note, however, that the astrophysical background can be extremely variable in helioprojective coordinates, as the Sun regularly (and periodically) moves past regions of high and low diffuse and point-source emission. To account for this, in each 8640 s exposure bin, we calculate the number of γ -ray counts expected from the astrophysical background. We set this as the expectation value for the background flux during that temporal bin. Utilizing this energy cut and ROI, we observe 20661 counts compared to a background expectation of 9074 counts, implying that our signal is roughly 60% produced by the solar disk. Because our model of the data is very accurate (see, e.g., Fig. 1), we expect the difference between the background data and the background model to be mostly time independent. The solar emission, on the other hand, may have significant variability on a variety of timescales.

To measure the variability of the solar disk flux, and compute its autocorrelation function, we first calculate the median flux along with non-Gaussian uncertainties ($\pm 1\sigma$ values) using a Bayesian technique. We assume that the solar and background fluxes are each Poisson-distributed and that the priors on each component are uniform in the logarithm.




FIG. 6. Tests of time variability of the solar disk γ -ray flux on a wide ranges of scales. Top: Autocorrelation function for offsets less than 400 days. The vertical lines denote the potential periodicities shown in the bottom panel. Bottom: Power spectrum for offsets less than 3000 days. Peaks above the 99% false alarm level (blue line) are denoted by the period of the peak value. The red vertical line at 27 days corresponds to 1 Bartels rotation and indicates a region where we may expect a periodic signal due to solar rotation.

To calculate the autocorrelation function at a time offset t , we calculate the Spearman correlation coefficient for the collection of all flux pairs separated by time t . We use the Spearman correlation coefficient because it is robust to outliers. We account for flux uncertainties through bootstrapping. Specifically, we calculate the autocorrelation function 10000 times, each time using a different light curve consistent with the non-Gaussian flux uncertainties. We then extract the mean and standard deviation of the autocorrelation function at each time offset. We calculate the autocorrelation function out to 3000-day offsets, slightly below the length of our total observation period. We then calculate the power spectrum of the solar disk light curve by calculating the Lomb-Scargle periodogram (LSP) of the autocorrelation function using the public code *astropy* [36]. The LSP takes into account the uncertainties of the autocorrelation function.

Figure 6 shows the results of this analysis, including both the autocorrelation function and its power spectrum. This analysis confirms significant power in the autocorrelation function on the longest timescales, finding significant and constant power at all periods exceeding 1000 days. This confirms the significant long-term periodicity (consistent with a solar cycle) is found in our primary dataset, utilizing

a separate analysis technique without ICS or Moon modeling.

We find several additional periodicities that are statistically significant. However, most of these are explained due to a combination of astrophysical and instrumental factors. We find a significant peak at periods ~ 181.5 days, which is consistent with the two passages of the Sun across the galactic plane per year. The analysis of this subsection does a poorer job subtracting the astrophysical background. The increased variability in our background signal as the Sun passes the bright galactic plane is likely to induce a periodic signal. We also find significant periodicities at 60.8 and 90.4 days that are well explained as higher harmonics of the physical 181.5 day periodicity.

Additionally, we find a significant line at 46.2 days, which is connected with the 53 day orbital precession period of the *Fermi*-LAT.⁵ This 53 day orbital precession period occurs with respect to J2000 coordinates, and the Sun's motion across 14% of the sky over that period decreases the orbital precession period (with respect to the Sun) to ~ 45 days.

⁵<https://fermi.gsfc.nasa.gov/ssc/observations/types/exposure/>.

Finally, we find reasonably significant evidence for a one-day periodicity of unknown origin. We note several key features of this periodicity: (1) it is approximately a 4σ effect, (2) our flux bins are 0.1 days, implying that binning effects may be present, though our analysis does not necessarily fail for periodicities near our temporal binning, (3) the periodicity is exactly 1 day (within the few minutes resolution of our scan), (4) the periodicity is found in the instrumental exposure (at the level of a $\sim 1\%$ fluctuation) and does not appear to be present in the counts data, (5) a one-day periodicity has been found in other *Fermi*-LAT observations, though these studies have not definitively determined its origin.⁶ Combining these factors, we believe this effect to be of unknown instrumental origin.

Most importantly, we find no evidence of significant power in the variability spectrum on periods between 24 and 34 days, which correspond to the rotation periods of the solar atmosphere around the equatorial and polar regions of the Sun, respectively. This strongly constrains models where certain “hot spot” regions of the solar atmosphere could persist and produce significant γ -ray fluxes over multiple solar rotations.

C. Morphology

Finally, we utilize the high *Fermi*-LAT angular resolution above 10 GeV to study the γ -ray emission morphology across the solar disk. This analysis mostly repeats Ref. [21] but includes two more years of data from the recent solar minimum. Compared to the rest of the paper, this subsection includes six extra months of data (up until August 29, 2020) because here it is easy to do and because maximizing the total exposure is important, due to the scarcity of high-energy γ -rays.

Unlike other analyses in this paper, in this section we do not perform any of the background subtraction methods described in Sec. II. Instead, we copy the techniques of Ref. [21], to which we refer the reader for details. The reasons are twofold. First, every γ -ray above 10 GeV is very well localized to within $\sim 0.3^\circ$ of its true location. The solar disk intensity in this ROI is more than $10\times$ brighter (at the highest energies, $100\times$ brighter) than both ICS and astrophysical backgrounds, rendering their modeling irrelevant. Second, with such a high angular resolution, the Sun itself further inhibits the background. γ -rays from astrophysical objects will be blocked by the Sun, producing a hole in astrophysical γ -ray emission that fills roughly 25% of the 0.5° ROI we employ.

In our previous analysis, we found statistically significant, and completely unexpected, evidence for two morphological components of solar disk γ -ray emission. The first emanates predominantly from the Sun’s polar regions and has a flux that is nearly independent of the solar cycle.

The second emanates predominantly from the Sun’s equatorial plane and has a flux that peaks strongly during solar minimum and disappears during the remainder of the solar cycle. The equatorial component produces a harder γ -ray spectrum that continues to γ -ray energies above 100 GeV.

Here we focus on comparing *Fermi*-LAT data from the 2008–2009 solar minimum with data from the 2017–2020 solar minimum. In advance, there are two possible hypotheses:

- (i) The equatorial and polar emission components depend on the polarity of the heliospheric magnetic field. In this case, we might expect a significant shift in the morphology of solar disk γ -ray emission between the cycle 23/24 and cycle 24/25 solar minima.
- (ii) The equatorial and polar emission components relate to other magnetic phenomena closer to the solar photosphere. In this case, we may expect the γ -ray emission morphology of these two minima to be similar.

We divide the data into four categories, based on the minima to which the event corresponds and the event energy. For the cycle 23/24 minimum, we analyze γ -ray data between August 4, 2008 (the beginning of the *Fermi* mission) and January 1, 2010. For the cycle 24/25 minimum, we analyze γ -ray data between November 1, 2017 and August 29, 2020. We cut our events into two energy bands, the first spanning from 10–50 GeV and the second for events above 50 GeV. For this analysis, we do not make any cuts on solar flares, as none are observed above our energy threshold. Additionally, we produce no model for the ICS, Moon, or astrophysical backgrounds, noting (as in Ref. [21]) that these backgrounds combine to produce only $\sim 1\%$ of the total flux in our ROI. We do however, make one cut on the data, eliminating all events observed while the Sun lies within 5° of the galactic plane, a choice that significantly decreases the level of astrophysical contamination but only moderately affects the total exposure.

Despite the fact that observations during the current solar minimum span a duration twice as long as the previous minimum (1032 vs 515 days), observations of the current minimum include nearly the same (99%) exposure, due to the decreased solar exposure of the *Fermi*-LAT after March 2018. The present results for solar cycle 23/24 are almost identical to those of Ref. [21], with small fluctuations due to the adoption of the more-recent P8_V3 IRFs (instrument response functions), which have slightly changed the angular uncertainty of several events.

Figure 7 shows the morphology of γ -ray emission across the solar disk. During both solar minima, the emission at energies of 10–50 GeV is nearly uniform across the solar surface, while the emission above 50 GeV shows a prominent equatorial component and little polar emission.

To further probe the polar and equatorial components of the solar disk flux, we follow the procedure of Ref. [21] and

⁶https://fermi.gsfc.nasa.gov/ssc/data/analysis/LAT_caveats_temporal.html.

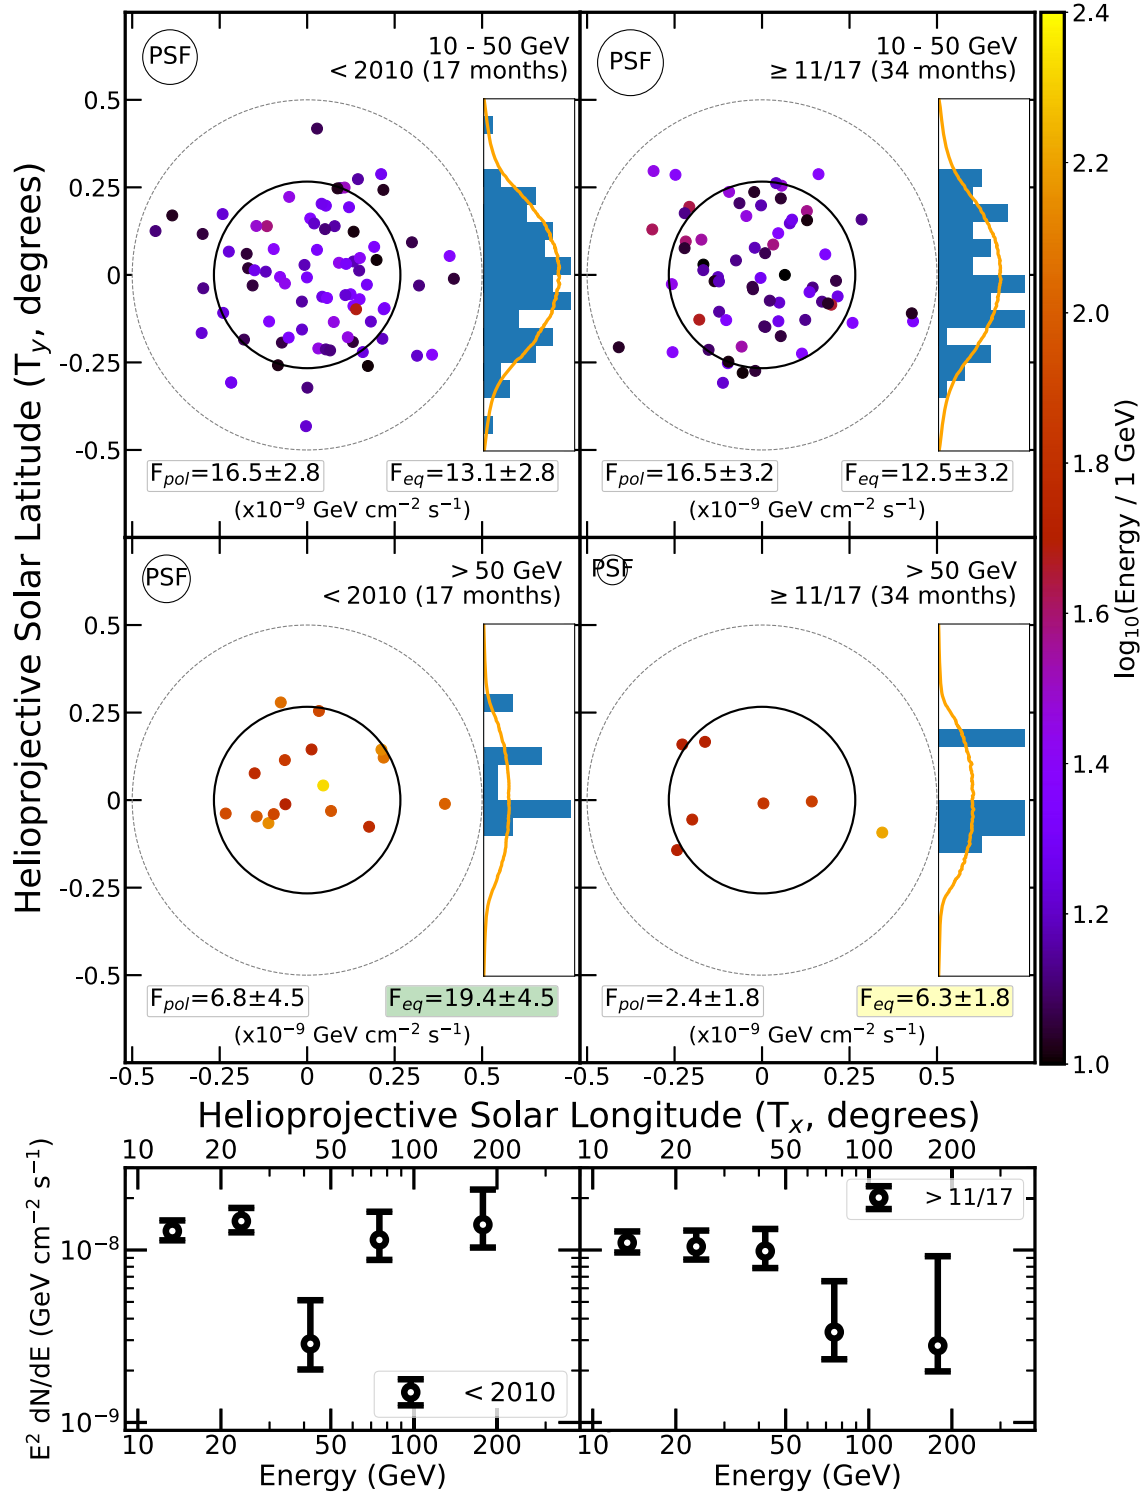


FIG. 7. Top: Locations and energies of high-energy solar disk γ -rays from the cycle 23/24 and cycle 24/25 minima in helioprojective coordinates. Data are separated into two temporal bins and two energy bins. The physical solar disk is represented by the solid circle and the 0.5° ROI by the dashed circle. The 68% containment region of the PSF for the average γ -ray is shown in the top left of each panel. The T_y positions of γ -rays are shown in the histogram, which is compared to the profile expected from isotropic emission smeared by the PSF (orange line). In each panel, we report the flux from the modeled polar and equatorial components, as described in the text. Bottom: The γ -ray spectra during the cycle 23/24 and cycle 24/25 minima.

fit the γ -ray emission utilizing a two-component model that includes homogeneous γ -ray emission across either: (a) the polar half of the solar disk, or (b) the equatorial half of the solar disk. These portions are divided at a helioprojective latitude T_y of $\pm 0.108^\circ$. We note that this helioprojective division equally divides the Sun into a 2D circle on the sky. If we were to instead ask about the total surface area of the 3D sun with respect to Earth, we would pick a cut at a slightly higher value of $T_{y,3D}$ of $\pm 0.125^\circ$. The “correctness” of either method depends on the average angle between the outgoing γ -ray and the normal to the solar surface, which can be quite complicated [37]. However, we note that the changes produced by this choice have no significant effect on our results. The emission from each component is smeared by the best-fit PSFs for each individual γ -ray, utilizing the King function model that describes the angular uncertainty of photons in the P8_V3 source PSF.

Utilizing this model, we find the best-fit fluxes of the polar and equatorial components, as well as their relative uncertainties. The polar and equatorial flux uncertainties depict the fractional contributions of each component to the recorded number of photons and *do not* include the statistical uncertainty of the total γ -ray flux (i.e., those coming from Poisson fluctuations in the number of photons that were observed).

Our analysis shows that the γ -ray morphology during the current and previous solar minima are similar. In both cases, the polar and equatorial γ -ray fluxes are similar between 10 and 50 GeV, while equatorial emission dominates at higher energies. For comparison, Ref. [21] found that the polar emission was dominant during the remaining solar cycle (2010–2017), where $75.6 \pm 4.1\%$ of the emission between 10–50 GeV and $93.1 \pm 6.8\%$ of the emission above 50 GeV, came from polar regions of the solar disk. It is worth noting that the relatively large statistical preference for equatorial emission in the high-energy data during the current solar minimum may be surprising, given the very low number of observed photons (7). However, we note that the stronger statistics stem from our photon-by-photon analysis of the PSF, combined with the fact that the three events belonging to the PSF3 class (signaling the best-recorded events) are the three events which lie closest to the equatorial plane. Still, due to the fact that the F_{pol} and F_{eq} fluxes are perfectly anticorrelated in this analysis, we note that the statistical preference for equatorial emission in this analysis is only about 1.1σ —though the flux ratios are also in close agreement with the equatorial preference observed during the 2010 solar minimum.

Figure 7 also shows the total γ -ray flux and uncertainty in four logarithmic energy bins over 10–100 GeV, along with one additional bin that includes the data between 100 and 316 GeV (encompassing all high-energy events in this analysis). For each bin, we include uncertainties based on the Poisson fluctuations in the γ -ray count in each bin (producing the asymmetric error bars that are observed).

The γ -ray spectrum during the current solar minimum is somewhat softer than the previous minimum. We find only seven γ -rays with energies above 50 GeV during the current minimum, compared to 16 events during the previous minimum, despite almost identical exposures. Moreover, events observed during the current minimum have lower energies (73.4 ± 37.3 GeV), compared to the previous minimum (97.4 ± 39.8 GeV). This combination leads to a factor of ~ 3 reduction in the γ -ray energy flux from the last minimum to the present one. Indeed, the γ -ray flux during the cycle 24/25 minimum barely exceeds the solar maximum flux, which had an energy flux of 6.9×10^{-6} MeV cm $^{-2}$ s $^{-1}$, compared to the solar minimum value of 8.7×10^{-6} MeV cm $^{-2}$ s $^{-1}$. Upcoming high-energy observations by the *HAWC* telescope, with a much larger effective area in the high-energy regime, will be important for understanding the high-energy disk emission during solar minimum. Observations during the solar maximum found no evidence for significant γ -ray emission above an energy threshold of ~ 1 TeV [23].

Finally, we note that the ~ 30 –50 GeV spectral dip (first identified in Ref. [22] and clearly detected during the first solar minimum here) is not observed during the most recent minimum. In fact, the 30–50 GeV flux from the current minimum significantly exceeds the flux from 2008–2010. The reason is unknown. One possibility is that the spectral dip depends on polarity and disappears during positive polarities. Another possibility is that the energy-range of the spectral dip shifts. Notably, if the dip moved to higher energies (e.g., 50–80 GeV), then it would also explain the different high-energy fluxes. At present, there is no model that favors either scenario.

IV. DISCUSSION AND CONCLUSIONS

The solar disk is one of the brightest γ -ray sources observed by the *Fermi*-LAT. Although the basic cause of the γ -ray emission is known to be bombardment by hadronic cosmic rays, many basic facts are unexplained. Compared to model predictions, the observed γ -ray emission is surprising: its luminosity is higher, its spectrum is harder (and features a mysterious dip), its time variation is stronger, and its morphology of the surface emission is less uniform [20–22]. A key way forward is to seek new observational clues. In other work, we have sought to extend the spectrum to higher energies through searches with the *HAWC* telescope [23,24]. Here we include a push to lower energies in the observational analysis. Understanding the Sun is important in its own right and is also an input to other physics, for example, tests of dark matter [37–46].

In this paper, we have produced detailed observations of the solar γ -ray flux over a full solar cycle, spanning an energy range between 100 MeV and 100 GeV. Utilizing a novel methodology to measure and subtract astrophysical backgrounds, we have made new high-precision measurements of the solar disk γ -ray spectrum, temporal variation,

and morphology. The solar disk γ -ray spectrum is significantly harder ($dN/dE \sim E^{-2.0}$ at low energies) compared to the cosmic-ray spectrum ($dN/dE \sim E^{-2.7}$).

Our observations indicate that the dominant temporal trend is a strong anticorrelation of solar γ -ray emission with solar activity (measured, e.g., via sunspot number). The γ -ray flux increases by approximately a factor of 2 during solar minimum and falls smoothly as the Sun approaches solar maximum. Notably the fall and rise of the solar disk γ -ray flux is symmetric around the solar maximum, with no evidence of a deviation based on the polarity of the heliospheric magnetic field. This trend matches previous observational data [20] but runs contrary to modeling efforts aimed at understanding solar γ -ray emission [18].

Further, we observe that the 11 yr variation has no dependence on energy. And, in the low-energy spectra (100 MeV to 10 GeV), we see no significant differences other than the overall variation. These are clues that the controlling reason for the variability is magnetic fields in the solar atmosphere as opposed to cosmic-ray modulation (as posited in Refs. [14,17,21,22]). This is supported by hints of features in the energy spectra and time profiles shown above, as the complexity of those fields would introduce a variety of energy and time scales.

Interestingly, our analysis finds no additional evidence for time variability beyond the 11 yr solar cycle (excluding contributions from specific solar flares). Three separate analyses have been performed to test this phenomenon: (1) yearly analyses where the full morphology of ICS emission was allowed to float, (2) near-monthly analyses where the morphology of ICS emission was fixed to the global value, and (3) a novel autocorrelation function approach that searched for general time variability.

In addition to our focus on low energies, we also conduct a new study of the highest-energy *Fermi*-LAT data, comparing the γ -ray emission observed during the 2007–2010 solar minimum and the 2017–2020 solar minimum. Because these solar minima both have low-solar activity, but have opposite magnetic polarities, comparing the γ -ray signal during these periods can allow us to constrain physical models for high-energy solar γ -ray emission. We obtain several important results: (1) The flux and spectrum below ~ 30 GeV are compatible between the two minima, indicating that the Sun’s polarity does not significantly affect the γ -ray signal through most of the *Fermi*-LAT energy range, (2) restricting ourselves to the best localized γ -rays above 10 GeV, we find a similar preference for equatorial, rather than polar, emission during each minimum, indicating that the magnetic polarity also does not affect the location at which cosmic rays preferentially interact with the solar surface, (3) we find no evidence for the 30–50 GeV “dip” during the current solar minimum, a feature that was most pronounced during the 2008 solar minimum [22], and (4) we find less evidence for emission above 100 GeV during the current solar minimum

(an observation of one γ -ray, compared to eight γ -rays in the previous minimum, with similar exposures). However, we note that the statistics are not sufficient to rule out the possibility that these results are consistent.

Finally, we note that the large time-dependent dataset produced by our novel methodology allows us to correlate solar disk γ -ray emission against nearly any time-dependent solar observable. For example, studies correlating solar γ -ray emission with the detection of coronal mass ejections, helmet streamers, coronal holes, and other solar phenomena are ongoing. Furthermore, we can reverse the process described in this paper to remove disk emission and measure the spectrum, time variability, and morphology of the ICS halo surrounding the solar position, for which results will be released in a forthcoming publication [32].

ACKNOWLEDGMENTS

We thank Gulli Jóhannesson, Igor Moskalenko, Mehr Un Nisa, Elena Orlando, and Andy Strong for helpful comments. T. L. is partially supported by the Swedish Research Council under Contract No. 2019-05135, the Swedish National Space Agency under Contract No. 117/19, and the European Research Council under Grant No. 742104. J. F. B., B. Z., and G. Z. were supported by NSF Grants No. PHY-1714479 and No. PHY-2012955. A. G. H. P. was supported by NASA Grants No. 80NSSC18K1728, No. 80NSSC19K1519, and No. 80NSSC20K1354. B. B. was supported by NASA Grant No. 80NSSC19K1709. T. L., B. Z., and G. Z. were additionally supported by NASA Grants No. 80NSSC18K1728 and No. 80NSSC19K1519.

APPENDIX A: EFFECT OF INVERSE-COMPTON MODELS

In addition to emission from the disk, the Sun produces a bright γ -ray flux via the interaction of cosmic-ray electrons with solar photons. The total flux of this component is a factor of a few brighter than the disk. However, the surface brightness of the ICS halo is much lower, with significant emission detected out to more than 20° from the Sun. Moreover, kinematic constraints (the similar directions of outgoing solar photons and cosmic-ray electrons) significantly inhibit ICS very close to the Sun. In this Appendix, we utilize these facts to differentiate between disk and halo emission.

If the Sun had no magnetic fields, the ICS halo morphology and energy spectrum would be trivial to calculate. Cosmic-ray electrons are produced by interstellar processes and are homogeneous in the solar system. The energy density of solar photons falls as $1/r^2$. However, solar wind and magnetic fields affect the propagation of electrons throughout the heliosphere. The strength of this interaction depends on the microphysics of heliospheric plasma [47,48]. However, the bulk effect of solar modulation can be modeled based on the phase of the solar cycle,

the tilt angle of the heliospheric current sheet, and the polarity of the heliospheric magnetic field [35,49]. Many models simplify the physics further, modeling solar modulation using a charge-, time-, and rigidity-dependent potential that cosmic rays must climb before encountering their target [50]. The value of this potential is fit to experiments.

Solar modulation also affects the spectrum and flux of the solar disk by modulating the cosmic-ray protons that impinge on the Sun. We discuss this in detail in [15]. However, the morphological impact of this modulation is small. At the level of the *Fermi*-LAT PSF, we can model the Sun as a homogeneous disk. The morphology of ICS emission, however, depends sensitively on the choice of modulation parameters.

Uncertainties in the ICS make it difficult to model the solar disk. ICS emission is brighter than the disk, making its accurate modeling important. Because the majority of ICS γ -rays are located several degrees from the Sun, the likelihood model fits the ICS intensity (flux per solid angle) far from the Sun and then extrapolates that result close to the Sun based on the assumed ICS morphology. Thus, ICS models with different modulation potentials will produce different solar disk fluxes. The effect is most pronounced at low energies, where the poor PSF of the *Fermi*-LAT smears out the small-scale differences between ICS and disk morphologies. In this paper, we concern ourselves only with the impact of ICS mismodeling on the calculated solar disk flux and spectrum. We leave our analysis of the ICS template itself to future work [32].

Figure 8 shows the solar disk spectrum for various ICS emission models. First (left), we plot the solar disk emission assuming ICS morphologies produced using a single potential over its full extent (0–45°). We note two

important results. First, above an energy of ~ 1 GeV, the degeneracy between our ICS model and the solar disk flux disappears entirely. In this range, the *Fermi*-LAT PSF is sufficiently accurate to differentiate between the disk and ICS halo. Second, we note a significant degeneracy at lower energies, producing a factor of $\sim 10 - 20\%$ uncertainty in the solar disk flux in the lowest energy bands.

The direction of this degeneracy is intuitive. The average ICS γ -ray is located 10–15° away from the Sun—and the critical region for fitting the ICS template (in terms of the log-likelihood) is 5–10° from the Sun. The ICS model is driven by fits within these regions. Because models with small modulation potentials have a large fraction of their flux close to the disk, the fit will produce a bright ICS flux near the Sun. Models with large modulation potentials have a small flux close to the disk and produce a dim ICS flux near the Sun. The solar disk template is the only parameter capable of responding to this change, and its flux floats accordingly. As the energy increases, the degeneracy disappears because all ICS models fall off very close to the Sun, due to kinematic considerations (when the angle between the momentum of a cosmic-ray electron and a solar photon is 0, the ICS cross section is 0).

There are many ways to account for this degeneracy. In Fig. 8 (right), we show our default choice. We divide the ICS template into seven angular bins, spanning $\{0 - 1, 1 - 2, 2 - 4, 4 - 6, 6 - 8, 8 - 10, 10 - 45\}^\circ$. The ICS template normalizations are allowed to float independently in each bin (but are fixed to be non-negative), allowing the ICS flux near the Sun to float independently of more distant regions. We note three results: (1) All ICS models produce a consistent measurement of the disk flux, implying that the ICS morphology within each bin does not




FIG. 8. Left: Solar disk γ -ray flux calculated using three different values for the solar modulation potential of the ICS template. Because ICS emission surrounds the Sun, is brighter than the disk (but with a lower surface brightness), and has a mean γ -ray distance several degrees from the Sun, different modulation models affect the observed disk flux. The effect is most significant at low energies, where the *Fermi*-LAT angular resolution is poor. Right: Same, but dividing the ICS model into seven angular bins spanning $\{0 - 1, 1 - 2, 2 - 4, 4 - 6, 6 - 8, 8 - 10, 10 - 45\}^\circ$, with the normalization of each bin allowed to float independently. Effects from our choice of solar modulation model are removed, and much greater latitude is given to the ICS morphology, producing conservative uncertainty bars that encompass both statistical and systematic errors.

affect the result. (2) The flux and uncertainty in the disk emission above ~ 1 GeV is unchanged, indicating that our modeling does not affect higher energy emission. (3) The error bars on the low-energy disk flux have increased, indicating that there is a significant degeneracy between the disk flux and the normalization of the inner-most ICS ring.

In Fig. 9, we test two alternative methods. First (left), we utilize the same binning but force the normalization factor of each inner ICS bin to be smaller than that of the larger ICS bin surrounding it. This is theoretically motivated because solar modulation primarily repels electrons—with stronger effects closer to the Sun. In this case, we find much smaller statistical uncertainties for each choice of modulation potential. This indicates that the largest uncertainty in the disk flux comes from scenarios where the ICS model adopts a saw-tooth configuration in the angular bins. However, in this method, the low-energy fits for different modulation models are not consistent, indicating that there is a residual systematic error in our model of the disk flux. For this reason, our default model allows the normalization of each ICS bin to float independently—as we believe that the statistical errors produced in that fit more accurately represent the true uncertainty.

In Fig. 9 (right), we again repeat our default analysis (allowing the flux in each bin to float to any non-negative value). However, we combine the two inner bins to form a single bin spanning $0-2^\circ$. This decreases the degeneracy of the ICS template at moderate energies (between 0.3 and 1 GeV), decreasing the uncertainty in the disk flux at those energies. The cost of this choice is the appearance (at low-significance) of a bump at ~ 0.25 GeV, which is produced when the angular resolution of the template becomes too poor, and the disk emission becomes degenerate with the larger ICS bin.

Overall, we find that these methods all provide rigorous fits to the solar disk flux. Models with a single ICS template provide results with excellent statistical precision, but include systematic errors that depend on the model choice. However, our default binning scheme produces results that: (i) do not depend on the choice of the modulation potential and (ii) have systematic uncertainties that are small compared to their conservative (and quantified) statistical uncertainties. Thus, we believe that our default analysis accurately represents our state of knowledge regarding the low-energy solar disk flux, though we note that the error bars may be conservative.

We note that this method differs from traditional *Fermi*-LAT analyses, which focuses on testing multiple physical models for the ICS emission and then minimizing the log-likelihood fit for each model. In this analysis, however, such a technique is suboptimal because the log-likelihood can be highly influenced from the fit of the ICS template far from the solar disk. We reserve a full likelihood analysis of the ICS morphology to a later paper [32].

However, we note that our analysis provides significant evidence for a modulation potential with a nontrivial rigidity dependence. In Table I, we show the $\Delta\text{LG}(\mathcal{L})$ in 24 energy bins between 0.1 and 100 GeV for models with a single ICS component. At low γ -ray energies ($\lesssim 0.5$ GeV), our fits prefer very high modulation potentials, which remove any low-energy electrons from regions near the Sun. However, at higher energies, this situation reverses and our models prefer modulation potentials near 0. The rapid shift (occurring over a single energy bin), may signal either an abrupt change in the effect of the heliospheric magnetic field on cosmic rays, or may alternatively indicate a systematic error in the differentiating the solar disk and ICS templates at low energies. Even if the abruptness of the




FIG. 9. Left: Same as Fig. 8 (right), except that the intensity normalization of each ICS angular bin is forced to be smaller than the larger angular bin surrounding it. This mimics the physical intuition that solar modulation can only prevent cosmic rays from reaching the Sun. The best-fit disk flux in this model is consistent with our default model, but the error bars are much smaller, and are not statistically consistent for different values of ϕ . Right: Same as Fig. 8 (right), except the two inner bins are replaced by a single bin that spans $0-2^\circ$. This removes some degeneracy between the disk and ICS flux between 0.3–1 GeV, at the cost of inducing a slight spectral feature at the lowest energies.

TABLE I. Change in log-likelihood fit for three choices of the solar modulation potential ($\phi_0 = 0$ MV, $\phi_{400} = 400$ MV, and $\phi_{1000} = 1000$ MV). The $\Delta\text{LG}(\mathcal{L})$ is compared to the best fit model, thus one column is always 0 (shaded green) by construction. We find that strong solar modulation is preferred for γ -ray energies below ~ 500 MeV, but the model quickly shifts to prefer very small modulation potentials at higher energies.

Energy (GeV)	$\Delta\text{LG}(\mathcal{L})_{\phi_0}$	$\Delta\text{LG}(\mathcal{L})_{\phi_{400}}$	$\Delta\text{LG}(\mathcal{L})_{\phi_{1000}}$
0.10–0.13	1.35	0.94	0.0
0.13–0.18	0.0	0.74	1.46
0.18–0.24	5.84	3.00	0.0
0.24–0.32	5.96	3.55	0.0
0.32–0.42	7.55	5.01	0.0
0.42–0.56	0.61	2.23	0.0
0.56–0.75	0.0	5.70	7.71
0.75–1.00	0.0	3.79	4.17
1.00–1.33	0.0	4.67	6.68
1.33–1.78	0.0	3.67	5.35
1.78–2.37	0.0	3.67	6.43
2.37–3.16	0.0	1.31	1.96
3.16–4.22	0.0	0.90	1.66
4.22–5.62	0.0	0.21	0.39
5.62–7.50	0.0	0.06	0.12
7.50–10.0	0.17	0.10	0.0
10.0–13.3	0.0	0.08	0.19
13.3–17.8	0.0	0.16	0.38
17.8–23.7	0.0	0.04	0.09
23.7–31.6	0.0	0.12	0.29
31.6–42.2	0.0	0.10	0.24
42.2–56.2	0.0	0.08	0.21
56.2–75.0	0.0	0.07	0.18
75.0–100.	0.0	0.14	0.34

change is unexpected, our model provides strong evidence for a important solar modulation effect at low energies and a very small effect in the higher energies. We will provide more details concerning this scenario in Ref. [32].

APPENDIX B: ACCOUNTING FOR BIAS IN *FERMI*-LAT/GBM CUT

In Sec. III A, we utilized two different cuts on solar flare activity intended to isolate the steady state γ -ray emission from the solar disk. The first cut removed data from time slices corresponding to events in the *Fermi*-LAT significant flares list. This removed only 0.3% of the total solar exposure, but removed the vast majority of flare activity. We additionally attempted a more liberal cut that removed not only the *Fermi*-LAT significant flares, but also all low-energy flares detected by the *Fermi*-GBM. This cut removed nearly 40% of the total data and strongly limited the potential contribution from subthreshold solar flares to the solar γ -ray flux.

In Fig. 2, we showed these methodologies produced similar γ -ray spectra, indicating that removing just a handful of the brightest solar flares was sufficient to clean




FIG. 10. Same as Fig. 3, but showing results for both our standard cut on *Fermi*-LAT detected flares and our conservative cut on both LAT- and GBM-selected flares. In each year, the γ -ray fluxes from both cuts are similar, verifying that subthreshold flares are not correlated (or anticorrelated) with γ -ray emission above 1 GeV.

the dataset of transient flare events. This conclusion was further strengthened by the lack of small-scale time-variability features, which would be indicative of transient flare activity.

One confusing result from Fig. 2 is the relative enhancement of the γ -ray flux above 1 GeV during periods when both LAT/GBM flares are removed. In the main text we noted that this is a correlation, but not a causation—GBM flares predominantly occur during periods of high solar activity, and periods with high solar activity have small steady-state γ -ray fluxes. That is, our choice to remove GBM cuts biased our solar exposure towards time periods where the solar disk happened to be bright.

To verify this, in Fig. 10 we show the γ -ray flux above 1 GeV using both our LAT and LAT/GBM time cuts. We further divide the data into yearly components, following the methodology from the main text. We find that in each period, the LAT and LAT-GBM cuts produce similar γ -ray fluxes. This provides further evidence supporting our assertion that the removal of GBM flares does not affect the steady state γ -ray spectrum.

APPENDIX C: COMPARISON WITH *FERMI*-LAT COLLABORATION [10]

The most recent previous analysis of the solar disk γ -ray flux below 1 GeV was performed by the *Fermi*-LAT collaboration in 2011 [10] (hereafter, *Fermi*2011). We note that the *Fermi*-LAT analysis was completed using the Pass 6 reconstruction algorithm, while our analysis utilizes the significantly updated Pass 8 event reconstruction. Here, we compare our results and discuss our interpretation of their differences.

In Fig. 11, we show our spectrum compared to *Fermi*2011. Two results immediately stand out. The first




FIG. 11. Comparison of the solar disk γ -ray flux determined in our analysis (orange), with the most recent results on the solar disk γ -ray flux below 1 GeV (blue, *Fermi*2011). Also shown are the best-fit γ -ray fluxes utilizing our analysis technique over the solar minimum time period, which closely matches the analysis period observed in *Fermi*2011 (brown). Datasets from this paper are binned at eight logarithmic energy bins per decade, while results from *Fermi*2011 are binned at seven logarithmic energy bins per decade.

is the significantly smaller error bars in our flux determinations, a result of the nearly 100-fold increase in exposure enabled by a combination of a longer observation time, as well as the elimination of cuts based on nearby bright sources, the galactic latitude of the Sun, and the position of the Moon.

The second difference is the factor of ~ 2 flux offset between both observations. This is primarily due to the time-variability of the disk γ -ray flux (first observed in [20] and described in detail in Sec. III B here). In Fig. 11, we also show a “Solar Min” flux computed over the period from August 4, 2008 to January 1, 2010 (similar to the August 4, 2008 to February 4, 2010 cut used in *Fermi*2011). This corrects the majority of the offset, in particular, above 1 GeV, where our results closely agree with those of *Fermi*2011. We note the improved precision from our analysis, even in data taken over the same time period.

At low energies, we note a small offset between *Fermi*2011 and our result, particularly near ~ 160 MeV, where a small bump appears in *Fermi*2011, which does not appear in our data. We also note that our solar-minimum analysis occasionally has error bars that are similar to, or slightly exceed that of *Fermi*2011. This is due to the significant freedom given to the normalization of the inner-ICS profile in our analysis—which implies that our uncertainties include both statistical and systematic uncertainties. If we were to use an analysis that was restricted to uniform or monotonic ICS profiles (based on a simple modulation potential), then we would obtain results with significantly smaller uncertainties (see, e.g., Fig. 9).

APPENDIX D: HIGH-ENERGY TIME VARIABILITY

In Sec. III B, we examined the yearly variability in the disk γ -ray spectrum, focusing on the energy range of

0.1–10 GeV, where the statistical power is highest. Here we examine evidence for variability at higher energies.

We note three critical results from previous studies. First, Ref. [20] searched for variability in the energy bin 10–100 GeV, but found no statistically significant evidence for a time-dependent component (though their results were also consistent with the time-dependence observed in the 1–10 GeV data). Secondly, Ref. [22] found statistically significant evidence for a large-amplitude spectral dip between 30 and 50 GeV, which is most significant during the 2008–2010 solar minimum (data after this period shows some evidence for a dip, but the amplitude is smaller and the data is not statistically significant on its own). Third, the dedicated analysis of the morphology of photons above 10 GeV [21] found nine γ -rays with energies above 100 GeV during solar minima, and no such events during the remainder of the solar cycle.

The combination of these three facts complicates the search for 10–100 GeV variability. We first note that the low statistics in this energy range (~ 500 events over 11 years of data) make it difficult to analyze the data (accounting for background modeling and the ICS component) at eight energy bins per decade. The pronounced spectral dip makes it difficult (and inaccurate) to combine energy bins above and below the dip, as the average energy of photons in each time period (and thus the normalized value of $E^2 dN/dE$) is not similar. The lack of any events recorded above 100 GeV during the solar maximum (years 3–8) makes it difficult to assign uncertainties (which would be very large in terms of $E^2 dN/dE$) in these energy bins.

Here, we make the following choices, noting that different choices could produce slightly different results. We rebin our solar disk and background models to four energy bins per decade, and then combine energy bins spanning from 10–31.6 GeV (below the dip) and 31.6–100 GeV (in/above the dip) in postprocessing, assuming an average

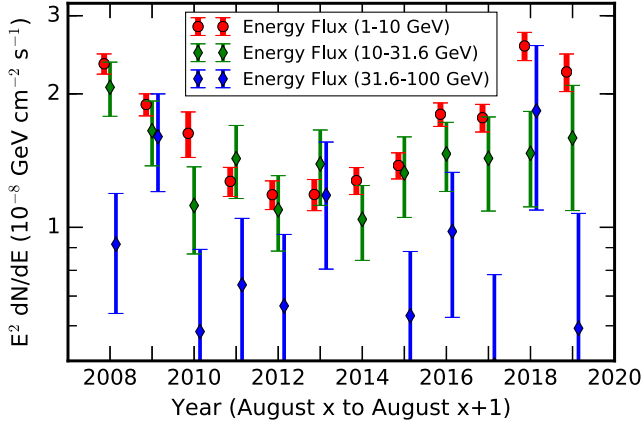


FIG. 12. Same as Fig. 3, except utilizing energy bins of 1–10 GeV (red), 10–31.6 GeV (green), and 31.6–100 GeV (blue). The points are slightly offset for clarity. We find strong evidence that the 11 yr variability, previously shown to anti-correlate with the sunspot number [20], continues up to an energy of ~ 30 GeV. Note that the highest energy data point in 2014 has a best-fit flux of 0 and falls off of the y axis of the plot. We find little evidence for periodic variability at higher energies, but note that the interpretation of this is extremely difficult, as it depends sensitively on our understanding of the spectral dip (see Ref. [22] and Fig. 2).

spectrum in the two ranges of E^{-2} . We do not include events above 100 GeV, but note that there is very strong evidence for temporal variability in this energy bin [21].

In Fig. 12, we show the resulting time variability in these energy bins, comparing our results to data from the 1–10 GeV energy bin shown in the main text. Our analysis shows some evidence for variability in the 10–31.6 GeV dataset, in particular, a statistically significant fall in the solar γ -ray emission between 2008 and 2012. The evidence for an increasing γ -ray flux during the more recent solar minimum is shakier, partially due to the effect of the diminished solar exposure since 2018. Between energies of 31.6–100 GeV, we find no strong evidence for solar disk variability, but note that in Ref. [22] this is best interpreted as offsetting effects between a brighter disk flux during solar minimum and a more pronounced solar dip feature during the same time period. While not shown here, we find that our analysis verifies the unique time variability of γ -ray emission above 100 GeV (first identified in Ref. [21]), with all γ -ray events observed at this energy stemming from the solar minimum period.

APPENDIX E: FAKE SOLAR POSITION TEST

As a final test of our astrophysical modeling, we note that the methodology described in Sec. II can be employed to calculate the flux of any γ -ray source that is moving with respect to the J2000 coordinate system. The only true constraint on this statement is that the source must move quickly enough so that every position (right ascension,

declination) is outside of the source ROI for at least some portion of the *Fermi*-LAT observation period.

We thus choose to verify our analysis by repeating the analysis chain for a fake Sun that trails the true position of the Sun by six months. This has the advantage that it moves through a similar background environment on average—meaning that many unique issues for our solar analysis (e.g., a source with multiple γ -ray flares) are likely to also affect the fake Sun analysis. While single flares will be constructed differently in this fake sky analysis (compared to the true solar analysis)—this is an irreducible issue for any data analysis test unless the exact flux and duration of each flare is known. Moreover, because the background astrophysical flux does not care about the solar cycle, the probability of such a one-off flare occurring is identical in each analysis.

We first calculate the fake helioprojective coordinates of every recorded γ -ray and then build a background model by removing all exposures and counts which lie within 45° of the fake solar position. We use this map of the astrophysical background to produce the astrophysical background map. We additionally add a Moon template (which tracks the true position of the Moon with respect to the fake solar position) and both a fake ICS and solar disk template (which surround the fake solar position). We then utilize our likelihood analysis (operating independently in each of 24 energy bins) in order to reconstruct the best-fitting solar disk flux.

In Fig. 13 we show the result of this analysis, which indicates that the “fake” solar disk flux lies orders of magnitude below the position of the real solar disk flux, and is identical with zero in every energy bin. The maximum local significance in any energy bin is 1.06σ in the energy




FIG. 13. The fake solar disk γ -ray flux obtained by repeating our entire analysis chain to study the γ -ray flux from a moving “false solar position” that trails the true solar position by six months. The fake solar disk flux is compatible with 0 in every energy bin (the maximum local significance in any bin is 1.06σ). The flux lies more than a factor of 100 below the true solar disk flux.

bin centered on 365 MeV. For comparison, we also replot the standard solar disk flux (obtained in our default analysis). Error bars are also shown for this fit but are invisible due to the large dynamic range between the real and fake solar disk fluxes. This analysis indicates that our astrophysical background subtraction technique is more than adequate to remove the astrophysical background and provide a robust fit for the solar disk component.

APPENDIX F: MODEL RESIDUALS

In Fig. 14 we show the residual map, calculated as the (data model)/model at energies between 0.1–0.3 GeV and

3–10 GeV. We show the calculated residual in two scenarios: (1) a model where only diffuse background is subtracted, but contributions from the Moon, ICS, and solar disk are ignored, and (2) the complete model where all solar and lunar components of the emission are included.

This analysis visually shows a large highly significant excess at the center of the γ -ray, which is consistent with the solar position; the remaining residuals from our full model have a much smaller amplitude and are randomly distributed across the sky like Poissonian fluctuations. We note that the average amplitude of these degree-scale variations is only at the level of $\sim 2\text{--}5\%$ at energies of 0.1–0.3 GeV, but rises to approximately 10–15% between

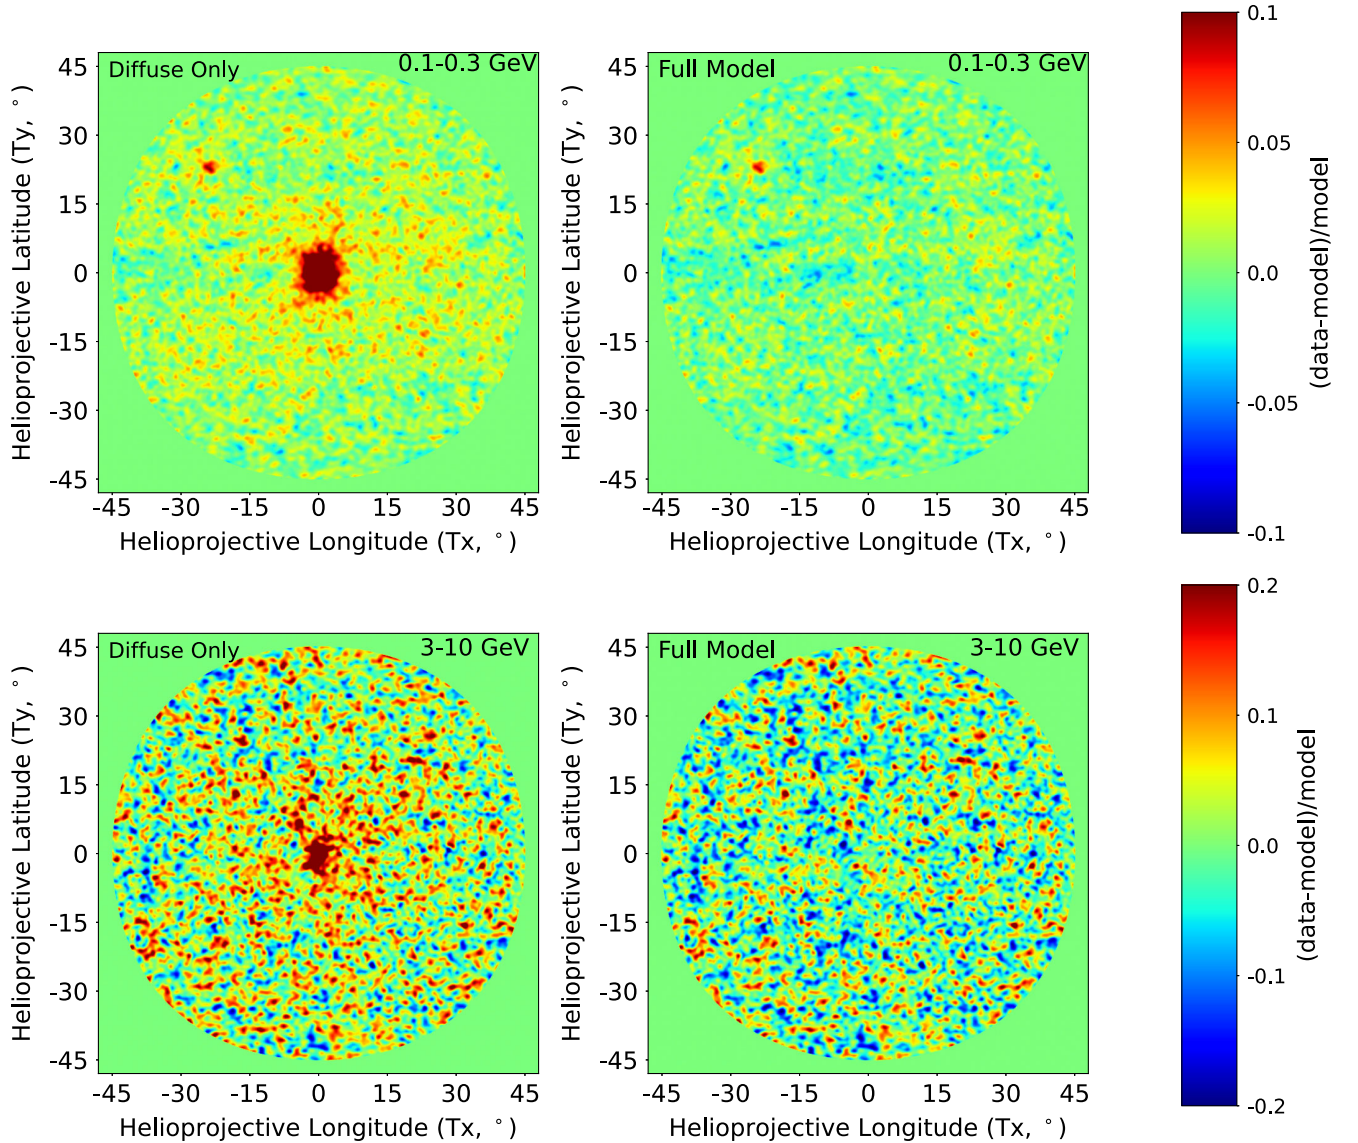


FIG. 14. Left: γ -ray residual (smoothed with a 1° Gaussian) in models where we only subtract the best-fitting diffuse emission model, but do not include model contributions from the Moon, solar disk, and solar ICS emission at energies of 0.1–0.3 GeV (top) and 3–10 GeV (bottom). Right: The same residual map after we have subtracted all model components, including the Moon, solar disk, and solar ICS emission. We find that the largest residual, by more than an order of magnitude, corresponds to the solar emission, implying that errors in our diffuse background modeling will not affect our analysis of the solar disk flux or spectrum.

3 and 10 GeV, due to the decreased photon statistics at higher energies.

Most importantly, we stress that the residual at the solar position oversaturates this image at both energy scales. When only diffuse emission is subtracted, the residual at the central pixel has a value of the (data model)/model of 0.75 between 0.1–0.3 GeV and 10.75 between 3 and 10 GeV (with a larger value owing to the much better angular resolution at high energies). This exceeds the level of background fluctuations by more than an order of magnitude in our lowest-energy bin (within a single pixel) and nearly two orders of magnitude at our higher energy bin—proving our assertion that background fluctuations cannot affect our determination of the solar disk emission.

APPENDIX G: COMPARISON WITH THE SATURATION MODEL

Comparisons of the relative log-likelihood between models that include or do not include a solar disk component provide a powerful method to determine the strength, spectrum, and temporal characteristics of the solar disk signal. However, calculations of the delta-log likelihood cannot, by themselves, determine whether the given model provides a quantitatively good fit to the γ -ray data.

In order to accomplish this task, we compare the fit of our model to the “saturated model,” a hypothetical model of the binned γ -ray data, which includes an independent degree of freedom for each individual spatial and energy bin. Thus, the saturated model includes a predicted number of counts that exactly matches the observed photon count in each bin. Because the saturated model is identical to the data, the likelihood of the fit of the saturated model to the data can be calculated as follows:

$$\ln(\mathcal{L}_{\text{Sat}}) = \sum_i \frac{\exp[-X_i] X_i^{X_i}}{X_i!}, \quad (\text{G1})$$

which is, in general, not equal to 0, even though the fit to the data is perfect. Using this value, the residual deviance can be calculated as follows:

$$\text{Res. Dev.} = -2[\ln(\mathcal{L}_{\text{BF}}) - \ln(\mathcal{L}_{\text{Sat}})], \quad (\text{G2})$$

where \mathcal{L}_{BF} is the likelihood of our best fit model, and the -2 allows the result to be interpreted as a χ^2 statistic with a number of degrees of freedom equal to the difference between the degrees of freedom in the best-fit and saturated models, which in all cases here will be a large number.

One notable issue with this comparison is that while the comparison of $\Delta LG(\mathcal{L})$ is a meaningful statistic, even when the number of photons per bin is very small, the calculation of the relevant number of degrees of freedom between the saturated model and the best-fit model becomes difficult to define when the number of bins exceeds the number of photons (as tends to occur at high energies in our analysis). Thus, the interpretation of the residual deviation using χ^2 statistics becomes tricky. Thus, we chose to calculate the residual deviation in two regimes. In the first, we utilize the maximum angular resolution in each individual energy bin (up to our maximum resolution of $n_{\text{side}} = 512$), which maintains an average of at least one count per bin. In the second case, we standardize our model to a common value of $n_{\text{side}} = 64$. We stress that both scenarios differ from what are utilized in every other fit in this paper—where the value $n_{\text{side}} = 512$ is used throughout. This is due to the fact that the $\Delta LG(\mathcal{L})$ is still robust in scenarios with fewer than one account per bin in regimes where the number of degrees of freedom is much smaller than the number of bins.

In Table II we show the results of this fit in each energy bin, including the $\chi^2/\text{d.o.f.}$ at both the maximal resolution of each simulation, as well as at a common resolution of $n_{\text{side}} = 64$. We find that the $\chi^2/\text{d.o.f.}$ is relatively small in both cases. While the models at high resolution (or when approaching the regime with one count per bin) do not provide formally good fits to the data, we certainly do not expect such an accuracy as there are a number of instrumental uncertainties, systematic effects from flaring sources, and mismodeling of the ICS emission, which are guaranteed to affect our analysis.

However, we stress that the degree of fit between our best-fit model is surprisingly good—far surpassing any *Fermi*-LAT analysis that uses models for the diffuse emission (rather than the data driven technique employed here) to model the γ -ray background. While no study of the *Fermi*-LAT diffuse modeling has (to our knowledge) produced a formal comparison using a saturated model, we compare the $\sim 1\%$ residuals obtained by our analysis with the $\sim 10\text{--}20\%$ residuals obtained in diffuse model fits to the *Fermi*-LAT sky [51]. In fact, we find that in our $n_{\text{side}} = 64$ models, a total residual deviation below 7350 would translate into a p value larger than 0.05, indicating that our models shockingly provide a formally good fit to the data. We find that this is true in 15 of our 24 energy bins—indicating the immense power of this model to accurately model background emission.

TABLE II. A comparison of the log-likelihood of our best-fitting model (with 10 d.o.f.) to the saturated model (where each individual pixel includes its own degree of freedom) in each energy bin of our analysis. We report the log-likelihood of the best-fit model, the saturated model, as well as the residual deviation calculated as $-2[LG(\mathcal{L}_{\text{BF}}) - LG(\mathcal{L}_{\text{Sat}})]$, while $\chi^2/\text{d.o.f.}$ interprets the residual deviation as a χ^2 statistic based on the difference in the number of degrees of freedom between the models. The color scheme represents the nside of the HEALPix binning used at each energy in order to guarantee that the expected number of events in the average pixel exceeds 1 (only for this test), including nside = 512 (green, 458702 d.o.f.), nside = 256 (blue, 114708 d.o.f.), nside = 128 (yellow, 28652 d.o.f.), nside = 64 (orange, 7164 d.o.f.), and nside = 32 (red, 1784 d.o.f.). On the right half, we report the same values at a common nside of 64 for all energy bins.

Energy (GeV)	$LG(\mathcal{L}_{\text{BF}})$	$LG(\mathcal{L}_{\text{Sat}})$	Res. dev.	$\chi^2/\text{d.o.f.}$	$LG(\mathcal{L}_{\text{BF}})$ (64)	$LG(\mathcal{L}_{\text{Sat}})$ (64)	Res. dev. (64)	$\chi^2/\text{d.o.f.}$ (64)
0.10–0.13	–1108720	–873094	471252	1.027	–32304	–28701	7204	1.007
0.13–0.18	–1129309	–894524	469570	1.023	–32676	–29018	7316	1.023
0.18–0.24	–1127444	–892568	469572	1.023	–32717	–28989	7457	1.042
0.24–0.32	–1107820	–871847	471946	1.029	–32403	–28683	7438	1.040
0.32–0.42	–1077005	–840866	472278	1.030	–31939	–28229	7419	1.037
0.42–0.56	–1037740	–799047	477386	1.041	–31253	–27628	7249	1.013
0.56–0.75	–987622	–745625	483994	1.055	–30481	–26879	7202	1.007
0.75–1.00	–927682	–679963	495438	1.080	–29665	–25995	7338	1.026
1.00–1.33	–858852	–605164	507376	1.106	–28669	–25020	7298	1.020
1.33–1.78	–770870	–508460	524820	1.144	–27452	–23792	7320	1.023
1.78–2.37	–672229	–406465	531528	1.159	–26065	–22458	7215	1.009
2.37–3.16	–232117	–170383	123468	1.076	–24587	–21042	7091	0.991
3.16–4.22	–205347	–141317	128060	1.116	–23167	–19524	7286	1.019
4.22–5.62	–176172	–110258	131828	1.149	–21533	–17911	7244	1.013
5.62–7.50	–59272	–43863	30818	1.076	–19926	–16292	7266	1.016
7.50–10.0	–52348	–36315	32066	1.119	–18366	–14666	7399	1.034
10.0–13.3	–45108	–28418	33380	1.165	–16732	–12967	7530	1.053
13.3–17.8	–37661	–21302	32718	1.142	–14898	–11235	7325	1.024
17.8–23.7	–13210	–9219	7982	1.116	–13210	–9219	7982	1.116
23.7–31.6	–11724	–7491	8467	1.184	–11724	–7491	8467	1.184
31.6–42.2	–10103	–5896	8414	1.176	–10103	–5896	8414	1.176
42.2–56.2	–3471	–2497	1948	1.092	–8308	–4343	7930	1.108
56.2–75.0	–3054	–2013	2082	1.167	–6650	–3027	7245	1.013
75.0–100.	–2582	–1570	2024	1.135	–5198	–2104	6187	0.865

- [1] M. A. Hapgood, *Adv. Space Res.* **47**, 2059 (2011).
[2] R. J. Murphy, C. D. Dermer, and R. Ramaty, *Astrophys. J. Suppl. Ser.* **63**, 721 (1987).
[3] E. J. Schneid, D. L. Bertsch, B. L. Dingus, J. A. Esposito, C. E. Fichtel, R. C. Hartman, S. D. Hunter, G. Kanbach, D. A. Kniffen, Y. C. Lin, H. A. Mayer-Hasselwander, P. F. Michelson, C. von Montigny, R. Mukherjee, P. L. Nolan, P. Sreekumar, and D. J. Thompson, *Astron. Astrophys. Suppl. Ser.* **120**, 299 (1996).
[4] M. Ackermann, M. Ajello, A. Albert, A. Allafort, L. Baldini, G. Barbiellini, D. Bastieri, K. Bechtol (Fermi LAT Collaboration), *Astrophys. J.* **787**, 15 (2014).
[5] M. Ajello, A. Albert, A. Allafort, L. Baldini, G. Barbiellini, D. Bastieri, R. Bellazzini, E. Bissaldi *et al.*, *Astrophys. J.* **789**, 20 (2014).
[6] M. Pesce-Rollins, N. Omodei, V. Petrosian, W. Liu, F. R. da Costa, A. Allafort, and Q. Chen, *Astrophys. J. Lett.* **805**, L15 (2015).
[7] N. Omodei, M. Pesce-Rollins, F. Longo, A. Allafort, and S. Krucker, *Astrophys. J. Lett.* **865**, L7 (2018).
[8] I. V. Moskalenko, T. A. Porter, and S. W. Digel, *Astrophys. J. Lett.* **652**, L65 (2006); **664**, L143(E) (2007).
[9] E. Orlando and A. Strong, *Astrophys. Space Sci.* **309**, 359 (2007).
[10] A. A. Abdo, M. Ackermann, M. Ajello, L. Baldini, J. Ballet, G. Barbiellini, D. Bastieri, K. Bechtol *et al.*, *Astrophys. J.* **734**, 116 (2011).
[11] J. F. Dolan and G. G. Fazio, *Rev. Geophys. Space Phys.* **3**, 319 (1965).
[12] L. E. Peterson, D. A. Schwartz, R. M. Pelling, and D. McKenzie, *J. Geophys. Res.* **71**, 5778 (1966).
[13] E. Orlando, *arXiv:0907.0557*.
[14] B. Zhou, K. C. Y. Ng, J. F. Beacom, and A. H. G. Peter, *Phys. Rev. D* **96**, 023015 (2017).
[15] G. Zhu, J. Beacom, T. Linden, A. Peter, and B. Zhou (to be published).

- [16] D. Seckel, T. Stanev, and T. K. Gaisser, *Astrophys. J.* **382**, 652 (1991).
- [17] M. Mazziotta, P. De La Torre Luque, L. Di Venere, A. Fassò, A. Ferrari, F. Loparco, P. Sala, and D. Serini, *Phys. Rev. D* **101**, 083011 (2020).
- [18] Z. Li, K. C. Y. Ng, S. Chen, Y. Nan, and H. He, [arXiv:2009.03888](#).
- [19] E. Orlando and A. W. Strong, *Astron. Astrophys.* **480**, 847 (2008).
- [20] K. C. Y. Ng, J. F. Beacom, A. H. G. Peter, and C. Rott, *Phys. Rev. D* **94**, 023004 (2016).
- [21] T. Linden, B. Zhou, J. F. Beacom, A. H. G. Peter, K. C. Y. Ng, and Q.-W. Tang, *Phys. Rev. Lett.* **121**, 131103 (2018).
- [22] Q.-W. Tang, K. C. Y. Ng, T. Linden, B. Zhou, J. F. Beacom, and A. H. G. Peter, *Phys. Rev. D* **98**, 063019 (2018).
- [23] A. Albert *et al.* (HAWC Collaboration), *Phys. Rev. D* **98**, 123011 (2018).
- [24] M. Nisa, J. Beacom, S. BenZvi, R. Leane, T. Linden, K. Ng, A. Peter, and B. Zhou, [arXiv:1903.06349](#).
- [25] W. Atwood *et al.* (Fermi-LAT Collaboration), [arXiv:1303.3514](#).
- [26] W. T. Barnes, *Astrophys. J.* **890**, 68 (2020).
- [27] E. Orlando and A. W. Strong, *Nucl. Phys. B, Proc. Suppl.* **239–240**, 266 (2013).
- [28] E. Orlando and A. Strong, [arXiv:1307.6798](#).
- [29] D. Bisschoff, M. Potgieter, and O. Aslam, *Astrophys. J.* **878**, 59 (2019).
- [30] L. J. Gleeson and W. I. Axford, *Astrophys. J.* **154**, 1011 (1968).
- [31] Z. Fujii and F. McDonald, *Adv. Space Res.* **35**, 611 (2005).
- [32] T. Linden, B. Zhou, J. Beacom, and A. Peter, [arXiv:21XX.XXXX](#).
- [33] H. Dembinski *et al.*, [10.5281/zenodo.3949207](#) (2020).
- [34] W. B. Atwood, A. A. Abdo, M. Ackermann, W. Althouse, B. Anderson, M. Axelsson, L. Baldini, J. Ballet, D. L. Band *et al.*, *Astrophys. J.* **697**, 1071 (2009).
- [35] I. Cholis, D. Hooper, and T. Linden, [arXiv:2007.00669](#).
- [36] T. P. Robitaille, E. J. Tollerud, P. Greenfield, M. Droettboom, E. Bray, T. Aldcroft, M. Davis, A. Ginsburg, A. M. Price-Whelan *et al.* (Astropy Collaboration), *Astron. Astrophys.* **558**, A33 (2013).
- [37] K. C. Y. Ng, J. F. Beacom, A. H. G. Peter, and C. Rott, *Phys. Rev. D* **96**, 103006 (2017).
- [38] W. H. Press and D. N. Spergel, *Astrophys. J.* **296**, 679 (1985).
- [39] A. H. Peter, *Phys. Rev. D* **79**, 103532 (2009).
- [40] S. Sivertsson and J. Edsjö, *Phys. Rev. D* **85**, 129905 (2012).
- [41] R. K. Leane, K. C. Y. Ng, and J. F. Beacom, *Phys. Rev. D* **95**, 123016 (2017).
- [42] J. Edsjo, J. Elevant, R. Enberg, and C. Niblaeus, *J. Cosmol. Astropart. Phys.* **06** (2017) 033.
- [43] A. Albert *et al.* (HAWC Collaboration), *Phys. Rev. D* **98**, 123012 (2018).
- [44] C. Niblaeus, A. Beniwal, and J. Edsjo, *J. Cosmol. Astropart. Phys.* **11** (2019) 011.
- [45] A. Cuoco, P. De La Torre Luque, F. Gargano, M. Gustafsson, F. Loparco, M. Mazziotta, and D. Serini, *Phys. Rev. D* **101**, 022002 (2020).
- [46] M. Mazziotta, F. Loparco, D. Serini, A. Cuoco, P. De La Torre Luque, F. Gargano, and M. Gustafsson, *Phys. Rev. D* **102**, 022003 (2020).
- [47] A. Vittino, C. Evoli, and D. Gaggero, *Proc. Sci., ICRC2017* (**2018**) 024 [[arXiv:1707.09003](#)].
- [48] M. J. Boschini, S. Della Torre, M. Gervasi, G. La Vacca, and P. G. Rancoita, *Adv. Space Res.* **64**, 2459 (2019).
- [49] I. Cholis, D. Hooper, and T. Linden, *Phys. Rev. D* **93**, 043016 (2016).
- [50] M. Kuhlen and P. Mertsch, *Phys. Rev. Lett.* **123**, 251104 (2019).
- [51] F. Acero *et al.* (Fermi-LAT Collaboration), *Astrophys. J. Suppl. Ser.* **223**, 26 (2016).