

MIT Open Access Articles

*DECIPHERING CONTRIBUTIONS TO THE EXTRAGALACTIC
GAMMA-RAY BACKGROUND FROM 2 GeV TO 2 TeV*

The MIT Faculty has made this article openly available. *Please share*
how this access benefits you. Your story matters.

Citation: Lisanti, Mariangela; Mishra-Sharma, Siddharth; Necib, Lina and Safdi, Benjamin R. "DECIPHERING CONTRIBUTIONS TO THE EXTRAGALACTIC GAMMA-RAY BACKGROUND FROM 2 GeV TO 2 TeV." *The Astrophysical Journal* 832, no. 2 (November 2016): 117 © 2016 The American Astronomical Society

As Published: <http://dx.doi.org/10.3847/0004-637x/832/2/117>

Publisher: IOP Publishing

Persistent URL: <http://hdl.handle.net/1721.1/109857>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.





DECIPHERING CONTRIBUTIONS TO THE EXTRAGALACTIC GAMMA-RAY BACKGROUND FROM 2 GeV TO 2 TeV

MARIANGELA LISANTI¹, SIDDHARTH MISHRA-SHARMA¹, LINA NECIB², AND BENJAMIN R. SAFDI²

¹Department of Physics, Princeton University, Princeton, NJ 08544, USA

²Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Received 2016 July 25; revised 2016 September 22; accepted 2016 September 22; published 2016 November 22

ABSTRACT

Astrophysical sources outside the Milky Way, such as active galactic nuclei and star-forming galaxies, leave their imprint on the gamma-ray sky as nearly isotropic emission referred to as the extragalactic gamma-ray background (EGB). While the brightest of these sources may be individually resolved, their fainter counterparts contribute diffusely. In this work, we use a recently developed analysis method, called the Non-Poissonian Template Fit, on up to 93 months of publicly available data from the *Fermi* Large Area Telescope to determine the properties of the point sources (PSs) that comprise the EGB. This analysis takes advantage of photon-count statistics to probe the aggregate properties of these source populations below the sensitivity threshold of published catalogs. We measure the source-count distributions and PS intensities, as a function of energy, from ~ 2 GeV to 2 TeV. We find that the EGB is dominated by PSs, likely blazars, in all seven energy sub-bins considered. These results have implications for the interpretation of IceCube's PeV neutrinos, which may originate from sources that contribute to the non-blazar component of the EGB. Additionally, we comment on implications for future TeV observatories such as the Cherenkov Telescope Array. We provide sky maps showing locations most likely to contain these new sources at both low ($\lesssim 50$ GeV) and high ($\gtrsim 50$ GeV) energies for use in future observations and cross-correlation studies.

Key words: BL Lacertae objects: general – dark matter – galaxies: active – galaxies: statistics – gamma rays: diffuse background – neutrinos

1. INTRODUCTION

The extragalactic gamma-ray background (EGB) is the nearly isotropic all-sky emission that arises from sources outside of the Milky Way. The OSO-3 (Clark et al. 1968; Kraushaar et al. 1972) and SAS-2 satellites (Fichtel et al. 1975, 1978) were the first to see hints of the EGB and have since been followed by EGRET (Sreekumar et al. 1998; Strong et al. 2004) and, most recently, the *Fermi* Large Area Telescope³ (Ackermann et al. 2015b, 2016b). The origin of the EGB remains an open question. The dominant contributions are likely due to blazars (Stecker et al. 1993; Stecker & Salamon 1996; Muecke & Pohl 1999; Narumoto & Totani 2006; Dermer 2007; Pavlidou & Venters 2008; Ajello et al. 2009; Fermi-LAT 2010; Abazajian et al. 2011; Stecker & Venters 2011; Ajello et al. 2012, 2014; Singal et al. 2012; Di Mauro et al. 2014c; Ajello et al. 2015; Ackermann et al. 2016a), star-forming galaxies (SFGs; Soltan & Juchniewicz 1999; Pavlidou & Fields 2002; Ando & Pavlidou 2009; Bhattacharya et al. 2009; Fields et al. 2010; Makiya et al. 2011; Ackermann et al. 2012b; Chakraborty & Fields 2013; Lacki et al. 2014; Tamborra et al. 2014), and misaligned active galactic nuclei (mAGNs; Stawarz et al. 2006; Inoue 2011; Massaro & Ajello 2011; Di Mauro et al. 2014a; Hooper et al. 2016). Understanding the relative contributions of these source components to the EGB has taken on a new sense of importance in light of IceCube's observation of ultra-high-energy extragalactic neutrinos (Aartsen et al. 2013a, 2013b, 2015a, 2015b), the origin of which still remains a mystery. For instance, the same sources that dominate the extragalactic neutrino background at \sim PeV energies may also contribute significantly to the EGB from \sim GeV–TeV energies (Murase et al. 2013; Tamborra et al. 2014; Hooper et al. 2016). In

addition, the EGB may harbor the imprints of more exotic physics such as dark-matter annihilation or decay (Bengtsson et al. 1990; Bergstrom et al. 2001; Ullio et al. 2002; Bottino et al. 2004; Bertone et al. 2005; Bringmann & Weniger 2012; Ackermann et al. 2015a; Ajello et al. 2015; Di Mauro & Donato 2015), as well as contributions from truly diffuse processes such as propagating ultra-high-energy cosmic-rays (Loeb & Waxman 2000; Kalashev et al. 2009; Ahlers & Salvado 2011; Murase et al. 2012; Taylor et al. 2015) and structure formation shocks in clusters of galaxies (Murase et al. 2008; Zandanel et al. 2015). Given the potential wealth of information that can be extracted from the EGB, deciphering its constituents remains a high priority.

Most recently, *Fermi* presented a measurement of the EGB intensity from 100 MeV to 820 GeV (Ackermann et al. 2015b). The total EGB intensity is the sum of all resolved point sources (PSs) and smooth isotropic emission. The smooth emission, referred to as the isotropic gamma-ray background (IGRB), arises from PSs that are too faint to be resolved individually as well as other truly diffuse processes. It is also important to note that both the EGB and IGRB may be contaminated by cosmic-rays that are mis-identified as gamma-rays; this emission is expected to be smoothly distributed across the sky. Of the known gamma-ray emitting PSs at high latitudes, which are captured by *Fermi*'s 3FGL (Acero et al. 2015) catalog from 0.1 to 300 GeV and the more recent 2FHL (Ackermann et al. 2016a) catalog from 50 to 2000 GeV, the dominant source class is blazars.

In this work, we use a novel analysis method, called Non-Poissonian Template Fitting (NPTF), to study the source populations that contribute to the EGB in a data-driven manner. The method relies on photon-count statistics to illuminate the aggregate properties of a source population, even when its constituents are not individually resolvable (Malyshev &

³ <http://fermi.gsfc.nasa.gov/>

Hogg 2011; Lee et al. 2015, 2016). This allows us to constrain the contribution of PSs to the EGB whose flux is too dim to be detected individually. While at very low fluxes the NPTF also loses the ability to distinguish PSs from smooth emission, the threshold for PS detection is lower for the NPTF than it is for other techniques that rely on finding individually significant sources. This is because the NPTF only measures the aggregate properties of a PS population.

Using the NPTF, we are able to recover, for the first time, the source-count distribution (e.g., flux distribution) for isotropically distributed PSs at high Galactic latitudes, as a function of energy from 1.89 GeV to 2 TeV. This builds on previous studies that use related methods to obtain the source-count distributions in single energy bins from ~ 2 –12 GeV (Lee et al. 2016; Zechlin et al. 2016b) and from 50–2000 GeV (Ackermann et al. 2016b).

The source-count distribution for a given astrophysical population convolves information about its cosmological evolution. For a flat, non-expanding universe, a uniformly distributed population of galaxies has a differential source-count distribution $dN/dF \propto F^{-5/2}$, where F is the source flux at Earth and dN is the differential number of sources (Sandage et al. 1995). This is the well-known Euclidean limit. However, the power-law index changes when one takes the standard Λ CDM cosmology and more realistic assumptions for the redshift evolution of source-dependent observables such as luminosity. Therefore, the features of the source-count distribution—especially, its power-law indices and/or flux breaks—encode information about the number of source classes contributing to the EGB as well as their cosmological evolution.

These source-count distributions provide the keys for interpreting the GeV–TeV sky. For example, they enable us to obtain the intensity spectrum for PSs, down to a certain flux threshold, as a function of energy. We find that while the EGB is dominated by PSs, likely blazars, in the entire energy range from 1.89 to 2000 GeV, there is also room for other source classes, which contribute flux more diffusely, to produce a sizable fraction of the EGB. Our findings may therefore leave open the possibility that IceCube’s PeV neutrinos (Aartsen et al. 2013a, 2013b, 2015a, 2015b) can be explained by pp hadronic interactions in, e.g., SFGs (Murase et al. 2013; Tamborra et al. 2014; Ando et al. 2015) or mAGNs (Hooper 2016), which—as we show in Section 3—show up as smooth isotropic emission under the NPTF. Additionally, the high-energy source-count distributions allow us to make predictions for the number of blazars, which dominate the high-energy data, that will be resolved by upcoming TeV observatories such as the Cherenkov Telescope Array (CTA; CTA Consortium 2011; Dubus et al. 2013). While our analysis does not let us conclusively identify the locations of these sources, we provide maps showing the locations on the sky where, statistically, there are most likely to be PSs.

This paper is organized as follows. We begin in Section 2 by reviewing the analysis methods. Section 3 then applies these methods to simulated sky maps. We cannot stress the importance of these simulated-data studies enough; they are crucial for proving the stability of the analysis methods and laying the foundation for the data results that follow. Our data study is divided into two separate analyses for low (1.89–94.9 GeV) and high (50–2000 TeV) energies, described in Sections 4 and 5, respectively. The global fits to the full energy range, as well as their implications, are discussed in Section 6. Further details on the creation of the simulated-data

maps are provided in Appendix A, and supplementary analysis plots in Appendices B and C. The main results of this work are summarized in a few key figures. In particular, the source-count distributions for the low- and high-energy analyses are shown in Figures 5, 8, and 10, respectively, while Figure 11 presents a spectral fit to the PS intensity from 2 GeV to 2 TeV.

2. METHODOLOGY

In this paper, we make use of both Poissonian and Non-Poissonian Template Fitting techniques. Poissonian template fitting is a standard tool in astrophysics for decomposing a sky map into component “templates” with different spatial morphologies. The NPTF builds upon this technique by allowing for the addition of templates whose spatial morphology traces the distribution of a PS population, even if the exact position of the sources that make up that population are not known. More precisely, in both template-fitting procedures, one starts with a data set d that consists of counts n_p in each pixel p .⁴ One then fits a model \mathcal{M} with parameters θ to the data by calculating the likelihood function

$$p(d|\theta, \mathcal{M}) = \prod_p p_{n_p}^{(p)}(\theta), \quad (1)$$

where $p_{n_p}^{(p)}(\theta)$ denotes the probability of observing n_p photons in pixel p with model parameters θ .

In Poissonian template fits, the probabilities $p_{n_p}^{(p)}(\theta)$ are Poisson distributions, with the model parameters θ only determining the means of the distributions. That is, the mean expected number of photon counts at each pixel p may be written as

$$\mu_p(\theta) = \sum_{\ell} \mu_{p,\ell}(\theta), \quad (2)$$

where the sum is over template components and $\mu_{p,\ell}(\theta)$ denotes the mean of the ℓ th component for model parameters θ . The θ may parameterize, for example, the overall normalization of the templates or the shapes of the templates. Then, the probability $p_{n_p}^{(p)}(\theta)$ is simply given by the Poisson distribution with mean μ_p .

In the NPTF, the situation is more complicated because we do not know where the PSs are. As a result, if we want to calculate the probability of observing n_p photons in a given pixel p , we must first calculate the probability that a PS (or a collection of PSs) exists in the vicinity of the pixel p , with a given flux (or set of fluxes). Then, for that PS population, we calculate the probability of n_p photons being produced in pixel p . Convolving these two calculations together leads to distinctly non-Poissonian probabilities. In particular, the probability distributions in the presence of unresolved PSs tend to be broader than Poisson distributions, if both distributions have the same mean expected number of photon counts. The intuition behind this fact is that, relative to a diffuse source, a collection of PSs leads to more “hot” pixels with many photons (where there are PSs) and more “cold” pixels with very few photons (where there are no PSs).

⁴ We will only work with a single energy bin at a time for simplicity, though, in principle, model parameters may be shared between energy bins. In this case, the likelihood function over the full energy range may be written as the product of the likelihood functions in the energy sub-bins.

Table 1
Parameters and Associated Prior Ranges for the Templates Used in the NPTF

Parameter	Prior Range	Parameter	Prior Range	Parameter	Prior Range
A_{diff}	[0, 2]	$\log_{10} A_{\text{iso}}^{\text{PS}}$	[-10, 20]	n_1	[2.05, 5]
A_{bub}	[0, 2]	$S_{b,3}$	[0.1, 1] ph	n_2	[1.0, 3.5]
A_{iso}	[0, 2]	$S_{b,2}$	[1, 30] ph	n_3	[1.0, 3.5]
		$S_{b,1}$	[30, $2 \times S_{b,\text{max}}$] ph	n_4	[-1.99, 1.99]

Note. The priors on the breaks $S_{b,1,2,3}$ are given in terms of counts, defined relative to the mean exposure ($\bar{\mathcal{E}}^{(p)}$) in the ROI. $S_{b,\text{max}}$ is the maximum number of photons in the 3FGL (Acero et al. 2015; 2FHL, Ackermann et al. 2016a) catalog in the energy bin of interest for the low (high)-energy analysis. Note that all prior distributions are linear-flat, except for that of $A_{\text{iso}}^{\text{PS}}$, which is log-flat. The baseline normalizations of the A_ℓ are described in the text.

2.1. The Templates

We include three Poissonian templates for (1) diffuse gamma-ray emission in the Milky Way, assuming the *Fermi* p8r2 (*gll_iem_v06.fits*) foreground model, (2) uniform emission from the *Fermi* bubbles (Su et al. 2010), and (3) smooth isotropic emission. Each of these templates is associated with a single model parameter describing its overall normalization. Variations to the choice of foreground model and bubbles template will be discussed in Section 4.2.

The model parameters specific to the isotropic-PS population enter into the source-count distribution dN/dF , which we characterize as a triply broken power law:

$$\frac{dN}{dF} = A_{\text{iso}}^{\text{PS}} \begin{cases} \left(\frac{F}{F_{b,3}}\right)^{-n_4} & F < F_{b,3} \\ \left(\frac{F}{F_{b,3}}\right)^{-n_3} & F_{b,3} \leq F < F_{b,2} \\ \left(\frac{F_{b,2}}{F_{b,3}}\right)^{-n_3} \left(\frac{F}{F_{b,2}}\right)^{-n_2} & F_{b,2} \leq F < F_{b,1} \\ \left(\frac{F_{b,2}}{F_{b,3}}\right)^{-n_3} \left(\frac{F_{b,1}}{F_{b,2}}\right)^{-n_2} \left(\frac{F}{F_{b,1}}\right)^{-n_1} & F_{b,1} \leq F \end{cases} \quad (3)$$

In particular, there are three breaks, $F_{b,1,2,3}$, along with four indices, $n_{1,2,3,4}$, and the overall normalization, $A_{\text{iso}}^{\text{PS}}$.⁵ The justification for a triply broken power law is that $F_{b,1}$ designates the high-flux loss of sensitivity, beyond which dN/dF cannot be probed because no sources exist with such high flux. The break $F_{b,3}$ designates the low-flux sensitivity, below which PS emission cannot be distinguished from smooth emission. This leaves $F_{b,2}$ to probe any physical break in the source-count distribution in the flux region where the NPTF can constrain it. We have verified, however, that the results do not change significantly if the source-count distribution is fit by a doubly broken power law.

It is important to stress that the photon-count probabilities are non-Poissonian in the presence of unresolved PSs because their locations are unknown. Once we know where a PS is, we can fix its location and describe it through a Poissonian template with a free parameter for the overall normalization of the source. However, even resolved sources with known locations may be characterized by the non-Poissonian template if we do not also put down Poissonian templates at their locations. This is the approach that we take throughout this work; that is, we model both the resolved (in the 3FGL and 2FHL catalogs) and

unresolved PS populations through a single dN/dF distribution, without individually specifying the locations of any sources.

The point-spread function (PSF) must be properly accounted for in the template-fitting procedure. The diffuse models are smoothed according to the PSF using the *Fermi* Science Tools routine `gtsrcmaps`. The bubbles template is smoothed with a Gaussian approximation to the PSF, with width set to give the correct 68% containment radius in each energy bin. We follow the prescription developed in Malyshev & Hogg (2011) to account for the PSF in the calculation of the non-Poissonian photon-count probabilities; for this, we use the King function parameterization of the PSF provided with the instrument response function for the given data set. In Section 4.2, however, we show that consistent results are obtained when using a Gaussian approximation to the PSF instead.

2.2. Bayesian Fitting Procedure

The formalism developed in Malyshev & Hogg (2011) and Lee et al. (2015, 2016; see also Zechlin et al. 2016b and Linden et al. 2016) is used to calculate the photon-count probability distributions in each pixel as a function of the Poissonian and non-Poissonian model parameters θ . Then, Bayesian techniques are used to construct a posterior distribution $p(\theta|d, \mathcal{M})$ for the parameters θ and the likelihood function in (1). We construct the posterior distribution numerically using the `MultiNest` package (Feroz et al. 2009; Buchner et al. 2014) with 700 live points, importance nested sampling and constant efficiency mode disabled, and sampling efficiency set for model-evidence evaluation.

All prior distributions are taken to be flat except for $A_{\text{iso}}^{\text{PS}}$, which is taken to be log-flat. The prior ranges for the model parameters are shown in Table 1. These prior ranges successfully reconstruct the source-count distributions of simulated data sets, as discussed in Section 3. Variations to the prior ranges in Table 1 are considered in Section 4.2.

In Table 1, the parameter A_ℓ denotes the normalization of the ℓ th template, which is defined in terms of a baseline value. The baseline value is obtained by first performing a Poissonian template fit over 17 (10) log-spaced energy sub-bins between 1.89 and 94.9 GeV (50 and 2000 GeV) for the low-(high)-energy analysis. When this procedure is applied to the low-energy analysis where the known PSs are very bright, we mask the 300 brightest and most variable 3FGL sources, at 95% containment. At both high and low energies, we include a PS model constructed from the 3FGL catalog.⁶ The fitting procedure then allows us to recover the normalizations for the diffuse background, bubbles, and isotropic templates in each energy sub-bin.

⁵ Note that the NPTF can also handle PS templates with non-trivial spatial distribution, as was done in the Inner Galaxy analyses in Lee et al. (2016) and Linden et al. (2016); though, in this work we will only consider the isotropic-PS template.

⁶ Importantly, we do not include the PS model or mask any PSs in the NPTF analyses.




Figure 1. Source-count distribution of the isotropic-PS population obtained by running the NPTF on simulated data in which the EGB arises from the Blazar-1 model (Ajello et al. 2015). Results are presented for the four energy bins considered. The source-count distribution of the input blazar model (dashed red) matches the posterior for the isotropic PSs (68% and 95% credible intervals, constructed pointwise, shaded in red) well at fluxes corresponding to counts above ~ 1 photon (vertical, dotted-dashed black). The vertical dotted green lines indicate the fluxes at which 90%, 50%, and 10% of the flux is accounted for, on average, by sources with larger flux (from left to right, respectively). The red points show the histogram of the simulated PSs, with 68% Poisson error bars (vertical). Note that the NPTF loses sensitivity to sources contributing less than ~ 1 photon; as a result, the NPTF result does not match the simulated data well below the dotted-dashed black line.

The actual energy bins used for the NPTF studies presented in this work are larger than the sub-bins described above. Therefore, the baseline normalizations used to define the NPTF priors in the energy range $[E_{\min}, E_{\max}]$ are found by applying the best-fit Poissonian normalizations from the individual sub-bins to the corresponding templates, which are then combined.⁷ Therefore, $A_\ell = 1$ in the NPTF analysis implies that the

⁷ In practice, however, this prescription for combining the templates between energy sub-bins does not significantly affect our results.

normalization of the ℓ th template is the same as that computed from the Poissonian scans. The benefit of this approach is that it allows one to keep track of how the individual Poissonian templates react to the addition of non-Poissonian ones. For example, the normalization of the diffuse-background template should remain consistent between a standard template analysis, where PSs are accounted for by the 3FGL model, and the NPTF analysis, where PSs are accounted for by the non-Poissonian template; indeed, we find that is the case in all of the analyses we perform.




Figure 2. Energy spectra for the isotropic and isotropic-PS templates in each energy bin considered; the 68% and 95% credible intervals, constructed from the posterior distributions, are shown in blue and red, respectively. The top row represents the results for simulated data, with the *ultraleanveto* PSF3 instrument response function, in which the EGB consists of only Blazar-1 sources (Ajello et al. 2015; top left) or Blazar-2 sources (Ajello et al. 2012, 2014; top right). The bottom row shows the same results, except when SFGs (Tamborra et al. 2014) are also included in the simulation. The simulated spectrum for blazars (SFGs) is shown in dashed red (blue). For the Blazar-1 model, the isotropic-PS template absorbs almost the entirety of the flux. For the Blazar-2 model, both smooth and PS isotropic components absorb flux, but their sum (EGB, purple band) is consistent with the input. When SFGs are also included, more emission is absorbed by the smooth isotropic template; however, the total emission absorbed by the smooth and PS isotropic templates is consistent with the expected total of SFG and blazar intensities. The spectrum for Galactic diffuse emission is shown by the green line in each panel (median only). The sum of all template emission (yellow band) agrees with the total spectrum of the simulated data. Note that the energy spectrum of the bubbles template is not shown.

2.3. Exposure Correction

While the source-count distribution dN/dF is defined in terms of flux, F , with units of $\text{ph}/\text{cm}^2/\text{s}$, the priors for the breaks in Table 1 are written in terms of counts, $S_{b,1\dots3}$. To convert from flux to counts, we multiply by the exposure of the instrument, with units of $\text{cm}^2 \text{ s}$. However, the relation between

flux and counts is complicated by the fact that the exposure of the instrument varies both with energy and position in the sky. Below, we describe how we deal with both complications, starting first with the energy dependence.

The exposure map in the i th energy sub-bin is given by $\mathcal{E}_i^{(p)}$. To construct the exposure map $\bar{\mathcal{E}}^{(p)}$ in the larger energy range from




Figure 3. Same as Figure 1, except for the Blazar-2 model (Ajello et al. 2012, 2014).

$[E_{\min}, E_{\max}]$, which contains multiple energy sub-bins, we average over the $\mathcal{E}_i^{(p)}$ of the individual sub-bins, weighted by a power-law spectrum $dN/dE \sim E^{-2.2}$ because this is generally consistent with the isotropic spectrum over most of our energy range. This procedure introduces a source of systematic uncertainty in going from counts to flux, because not all source components have an energy spectrum consistent with this spectrum. However, we have checked that variations to this procedure—such as weighting the exposures in the sub-bins by power laws of the form E^{-n} , with n varying between 1 and 3—do not significantly change the results.⁸

⁸ We have also checked that weighting the exposures in the sub-bins by the intensities computed from the Poissonian template scans gives consistent results.

The weighting procedure is most important at very high energies, on the order of hundreds of GeV, where the exposure map varies strongly across the energy sub-bins.

The breaks $S_{b,1\dots 3}$ in Table 1, with units of counts, are defined relative to the mean exposure $\langle \bar{\mathcal{E}}^{(p)} \rangle$, averaged over all pixels in the region of interest (ROI). Because the NPTF is performed at the level of counts and not flux, we must also convert the source-count distribution dN/dF to a distribution $dN^{(p)}/dS$, which is unique to each pixel p :

$$\frac{dN^{(p)}}{dS}(S) = \frac{1}{\bar{\mathcal{E}}^{(p)}} \frac{dN}{dF} \Big|_{F=S/\bar{\mathcal{E}}^{(p)}}. \quad (4)$$




Figure 4. Best-fit energy spectra for the NPTF analysis using Pass 8 *ultracleanveto* data and the p8r2 foreground model. The left (right) panel shows the PSF3 (PSF1–3) results. The 68% and 95% credible intervals, constructed from the posterior distributions in each energy bin, are shown for the isotropic-PS and smooth isotropic templates in red and blue, respectively. The median intensity for the foreground model is also shown (green). The sum of all the components (yellow band) agrees with the total spectrum of the *Fermi* data (black). The *Fermi* bubbles contribution is subdominant (averaged over the full region of interest) and is thus not plotted. For comparison, the spectrum of the 3FGL sources is shown in dashed black. We caution the reader that, at higher energies, the 3FGL spectra are driven by extrapolations from low energies where the statistics are better. The systematic uncertainties associated with this extrapolation are difficult to quantify and are not shown here.

Table 2

Best-fit Intensities for All Templates Used in the NPTF Analysis of Pass 8 *ultracleanveto* PSF3 Data and the p8r2 Foreground Model

Energy (GeV)	I_{EGB}	$I_{\text{iso}}^{\text{PS}}$	I_{iso} ($\text{cm}^{-2} \text{s}^{-1} \text{sr}^{-1}$)	I_{diff}	I_{bub}
1.89–4.75	$1.38^{+0.05}_{-0.04} \times 10^{-7}$	$9.00^{+0.66}_{-0.54} \times 10^{-8}$	$4.82^{+0.43}_{-0.52} \times 10^{-8}$	$3.22^{+0.02}_{-0.02} \times 10^{-7}$	$2.90^{+0.67}_{-0.69} \times 10^{-8}$
4.75–11.9	$5.46^{+0.24}_{-0.22} \times 10^{-8}$	$2.68^{+0.26}_{-0.21} \times 10^{-8}$	$2.77^{+0.18}_{-0.21} \times 10^{-8}$	$7.38^{+0.15}_{-0.16} \times 10^{-8}$	$1.44^{+0.39}_{-0.39} \times 10^{-8}$
11.9–30.0	$1.76^{+0.10}_{-0.09} \times 10^{-8}$	$7.17^{+0.99}_{-0.76} \times 10^{-9}$	$1.04^{+0.08}_{-0.08} \times 10^{-8}$	$1.63^{+0.07}_{-0.07} \times 10^{-8}$	$5.18^{+2.35}_{-2.23} \times 10^{-9}$
30.0–94.9	$5.74^{+0.46}_{-0.41} \times 10^{-9}$	$2.40^{+0.48}_{-0.38} \times 10^{-9}$	$3.30^{+0.39}_{-0.42} \times 10^{-9}$	$3.73^{+0.31}_{-0.33} \times 10^{-9}$	$1.46^{+1.25}_{-0.92} \times 10^{-9}$

Note. Note that the *Fermi* bubbles template intensity is defined relative to the interior of the bubbles, while the intensities of the other templates are computed with respect to the region $|b| \geq 30^\circ$. The best-fit EGB intensity, which is the sum of the smooth and PS isotropic contributions, is also shown. The posterior distributions for the template intensities are provided in Figures 17–20.

Then, the photon-count probability distribution must be computed uniquely at each pixel. In practice, however, it is numerically expensive to perform this procedure for every pixel in the ROI. Instead, we follow Zechlin et al. (2016b) and break the ROI up into N_{exp} regions by exposure. Within each region, we assume that all pixels have the same exposure, which is taken to be the mean over all pixels in the sub-region. The likelihood function is then computed uniquely in each exposure region, and the total likelihood function for the ROI is the product of the likelihoods across exposure regions. In practice, we find that our results are convergent for $N_{\text{exp}} \geq 10$. We will take $N_{\text{exp}} = 15$ throughout this work, though we have checked that our main results are consistent with those found using $N_{\text{exp}} = 25$.

2.4. Data Samples

We run the NPTF analysis, as described above, on *Fermi* data, considering low (1.89–94.9 GeV) and high (50–2000 GeV)

energies separately. The former is discussed in Section 4, while the latter is the focus of Section 5. The primary difference between the data sets used in these studies is the data-quality cuts; moving to higher energies requires loosening these criteria to avoid being limited by statistics. The overlap in energy between the two studies allows us to compare the consistency of the results when transitioning between analyses.

The low-energy study uses the Pass 8 *Fermi* data from ~2008 August 4 to 2015 June 3. The primary studies use the top quartile of the *ultracleanveto* event class (PSF3) as ranked by angular resolution; though, the top-three quartiles (PSF1–3) are also studied separately.⁹ As a systematic check, we also consider the top-three quartiles of source data. The *ultra-cleanveto* event class is the cleanest event class released with the Pass 8 data and is recommended for studies of the EGB. However, the source event class has an enhanced exposure and thus may be advantageous at high energies where statistics

⁹ The PSF quartiles indicate the quality of the reconstructed photon direction, with “PSF3” being the best and “PSF0” being the worst.




Figure 5. Best-fit source-count distribution, as a function of energy, for the isotropic-PS population obtained by the NPTF analysis of Pass 8 *ultracleanveto* PSF3 data with the `p8r2` foreground model. The median (red line) and 68% and 95% credible intervals (shaded red bands) are shown. The vertical dotted–dashed black line denotes the ~ 1 photon boundary, below which the NPTF begins to lose sensitivity. The vertical dotted red lines indicate the fluxes at which 90%, 50%, and 10% of the flux is accounted for, on average, by sources of larger flux (from left to right, respectively). The black points correspond to the *Fermi* 3FGL sources, with 68% statistical error bars (vertical). The NPTF is expected to be sensitive down to the ~ 1 photon limit, extending the reach to sources below the 3FGL detection threshold. This is most apparent in the lowest energy bin, where the apparent 3FGL flux threshold is ~ 10 times higher than that for the NPTF. We caution the reader that, at higher energies, the 3FGL spectra are driven by extrapolations from low energies where the statistics are better. The systematic uncertainties associated with this extrapolation are difficult to quantify and are not included in the source counts shown here.

become limited. On the other hand, we expect the source data to have additional cosmic-ray contamination relative to the *ultracleanveto* data.

The recommended¹⁰ event quality cuts are applied, requiring that all photons have a zenith angle of less than 90° and satisfy

¹⁰ http://fermi.gsfc.nasa.gov/ssc/data/analysis/documentation/Cicerone/Cicerone_Data_Exploration/Data_preparation.html

“DATA_QUAL==1 && LAT_CONFIG==1 && ABS(ROCK_ANGLE) < 52.” A HEALPix (Gorski et al. 2005) pixelation is used with $n_{\text{side}} = 128$, which corresponds to pixels roughly 0.5° to a side. We consider four separate energy bins: [1.89, 4.75], [4.75, 11.9], [11.9, 30], and [30, 94.9] GeV.

In the low-energy analysis with *ultracleanveto* PSF3 data, the means of the weighted exposure maps in the four increasing energy bins are $[5.78 \times 10^{10}, 5.40 \times 10^{10}, 5.18 \times 10^{10},$

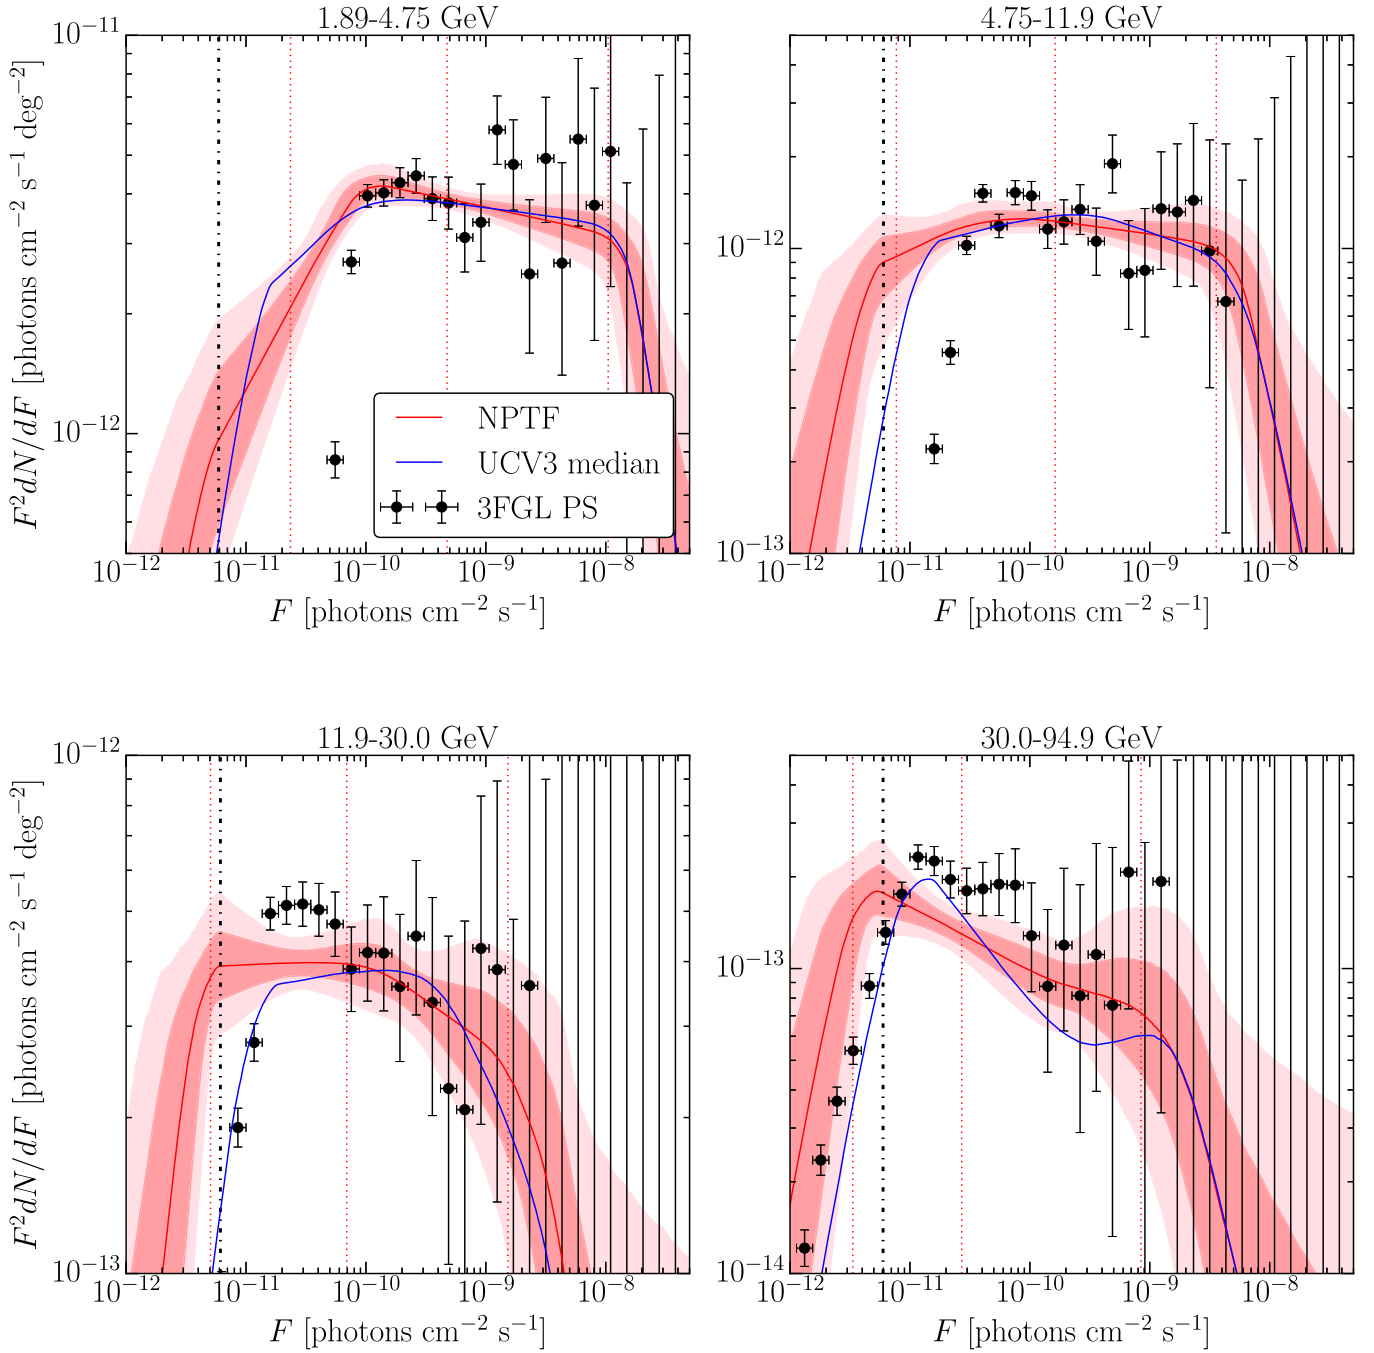


Figure 6. Same as Figure 5, except using the top three quartiles (PSF1–3) of the Pass 8 *ultracleanveto* data. The median source-count distribution for the PSF3 analysis is shown in blue. The best-fit values for the source-count distributions are provided in Table 6 of the Appendix.

Table 3
Best-fit Parameters for the Source-count Distributions Recovered for Each Energy Bin; the Flux Breaks $F_{b,i}$ and Indices n_i Are Labeled from Highest to Lowest ($F_{b,i} > F_{b,i+1}$)

Energy (GeV)	n_1	n_2	n_3	n_4	$F_{b,3}$	$F_{b,2}$ ($\text{cm}^{-2} \text{s}^{-1}$)	$F_{b,1}$
1.89–4.75	$3.96^{+0.68}_{-0.80}$	$2.04^{+0.05}_{-0.05}$	$1.74^{+0.19}_{-0.37}$	$-0.40^{+1.18}_{-1.05}$	$1.13^{+0.39}_{-0.52} \times 10^{-11}$	$1.22^{+2.00}_{-0.56} \times 10^{-10}$	$1.43^{+0.51}_{-0.46} \times 10^{-8}$
4.75–11.9	$3.84^{+0.78}_{-0.86}$	$2.13^{+0.15}_{-0.13}$	$1.91^{+0.09}_{-0.12}$	$-0.44^{+1.21}_{-1.03}$	$1.16^{+0.47}_{-0.51} \times 10^{-11}$	$2.95^{+1.80}_{-1.79} \times 10^{-10}$	$5.52^{+2.66}_{-2.06} \times 10^{-9}$
11.9–30.0	$3.54^{+0.96}_{-0.91}$	$2.42^{+0.41}_{-0.32}$	$1.97^{+0.11}_{-0.13}$	$-0.14^{+1.13}_{-1.15}$	$1.11^{+0.52}_{-0.50} \times 10^{-11}$	$3.47^{+1.56}_{-1.76} \times 10^{-10}$	$2.83^{+1.34}_{-1.34} \times 10^{-9}$
30.0–94.9	$3.63^{+0.89}_{-0.98}$	$1.83^{+0.52}_{-0.47}$	$2.51^{+0.29}_{-0.21}$	$-0.20^{+1.15}_{-1.16}$	$1.02^{+0.47}_{-0.46} \times 10^{-11}$	$2.48^{+1.86}_{-1.36} \times 10^{-10}$	$1.68^{+0.68}_{-0.65} \times 10^{-9}$

Note. These values correspond to the NPTF analysis of Pass 8 *ultracleanveto* PSF3 data with the p8r2 foreground model. The median and 68% credible intervals are recovered from the posterior distributions, which are provided in Figures 17–20.

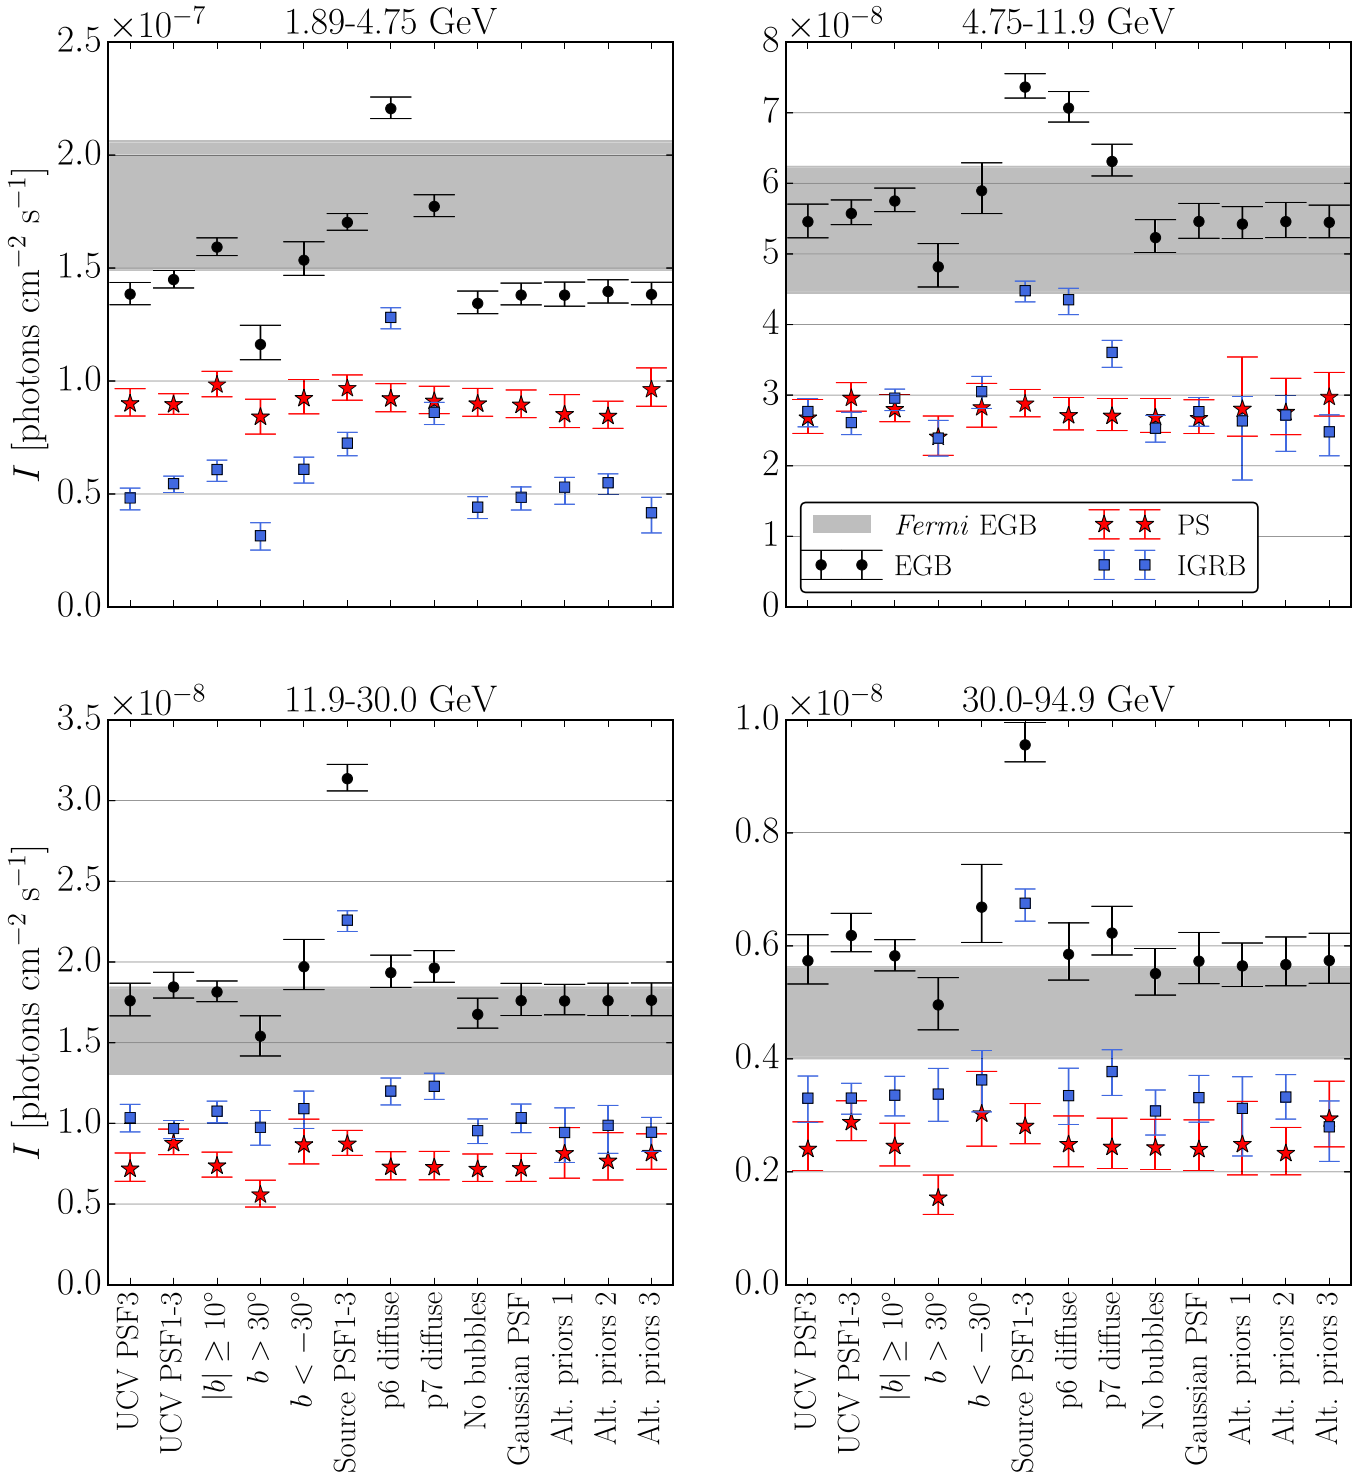


Figure 7. Comparison of the EGB (black circles), IGRB (blue squares), and PS (red stars) intensities recovered by the NPTF for the various systematic tests described in Section 4.2 and Appendix B.2. Note that “UCV” is shorthand for *ultracleanveto*. The gray band is meant to indicate the systematic uncertainty associated with the measured *Fermi* EGB (Ackermann et al. 2015b; see the text for more details).

$5.38 \times 10^{10} \text{ cm}^2 \text{ s}$ over the ROI with $|b| \geq 30^\circ$. The 68% containment radii for the PSF, averaged over the isotropic spectra in the energy sub-bins, are [0.20, 0.11, 0.06, 0.04] degrees. Going to PSF1–3 data, the exposures increase to [1.69×10^{11} , 1.66×10^{11} , 1.63×10^{11} , 1.67×10^{11}] $\text{cm}^2 \text{ s}$, while the 68% containment radii of the PSF degrade to [0.32, 0.16, 0.10, 0.08] degrees. Going to source data with PSF1–3, the exposures ([2.10×10^{11} , 2.07×10^{11} , 2.07×10^{11} ,

2.15×10^{11}] $\text{cm}^2 \text{ s}$) increase further, while the 68% containment radii ([0.32, 0.16, 0.10, 0.08] degrees) are essentially the same as in the *ultracleanveto* case.

The high-energy analysis uses the Pass 8 *Fermi* data from ~ 2008 August 4 to 2016 May 2 and all PSF quartiles of either the *ultracleanveto* or *source* event class. The ROI is also extended to $|b| > 10^\circ$. We include more data in the high-energy analysis because there are far fewer photons than at

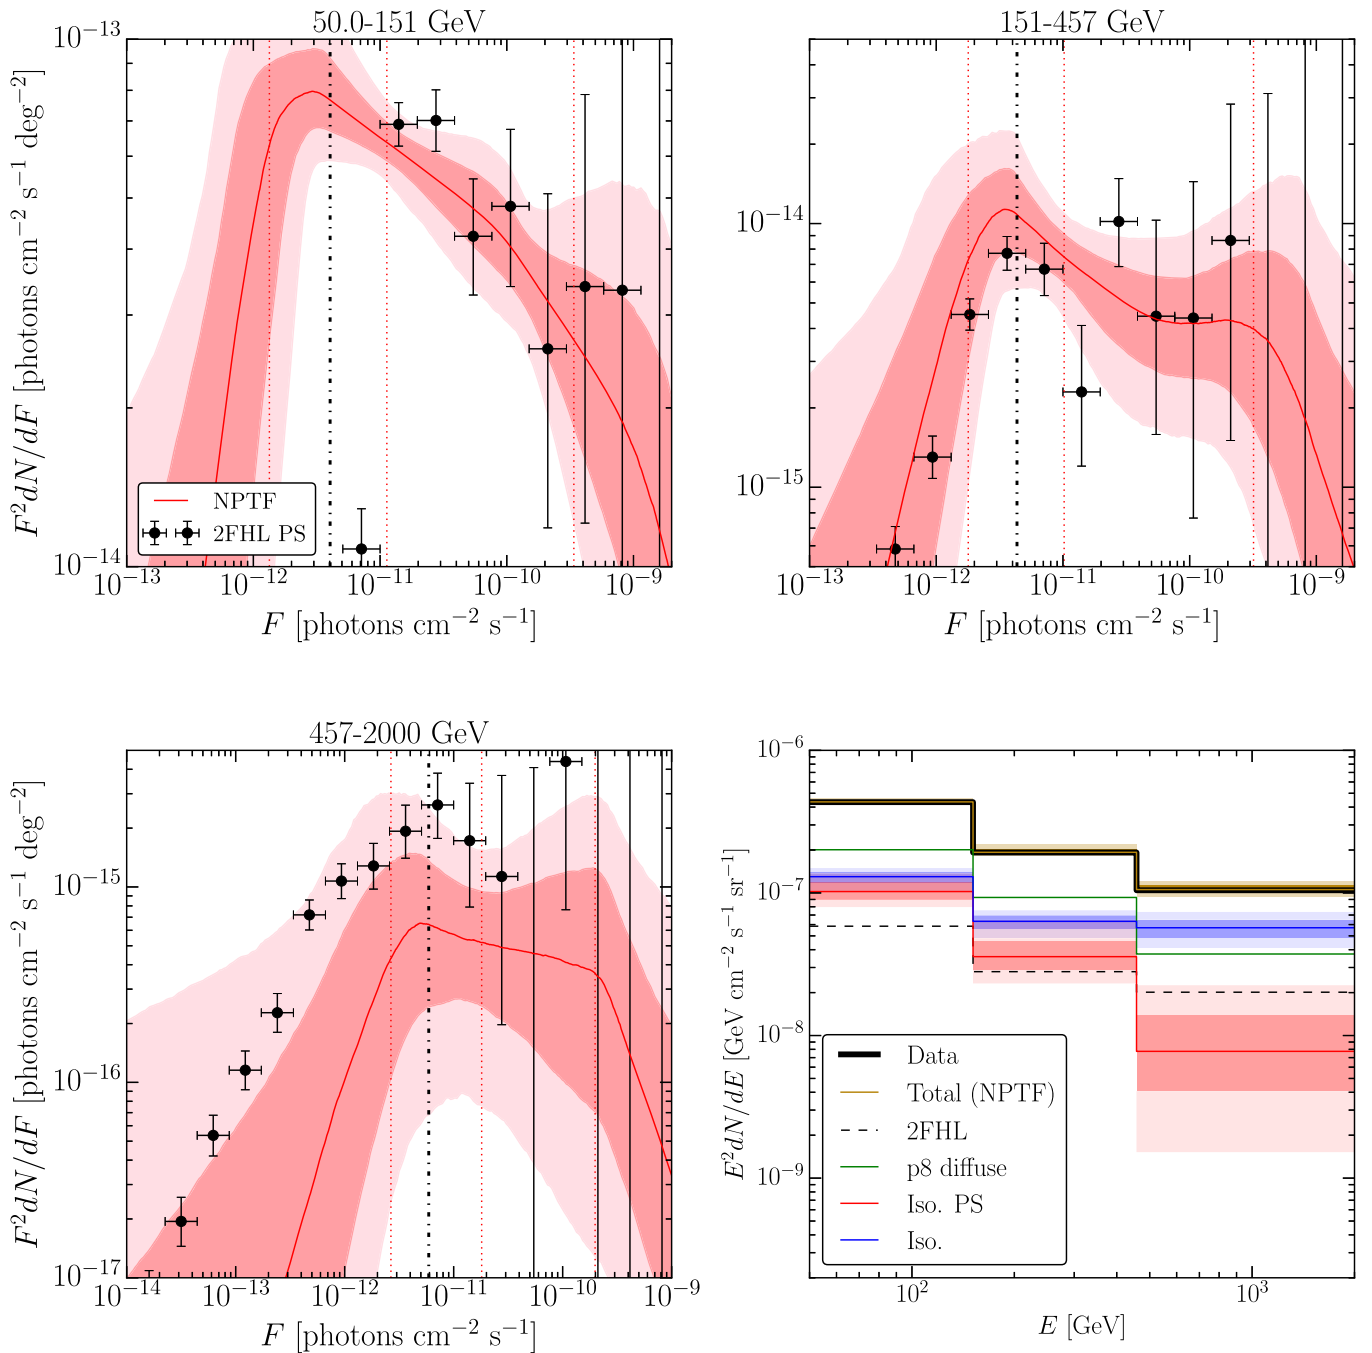


Figure 8. NPTF results for the high-energy analysis of all quartiles of Pass 8 *ultracleanveto* data. Top row and bottom left panel: the best-fit source-count distribution for the isotropic-PS population, for each separate energy bin, is shown using the same format conventions as in Figure 5. The black points correspond to the *Fermi* 2FHL sources (Ackermann et al. 2016a), with 68% statistical error bars (vertical). Bottom right panel: best-fit energy spectrum. The 68% and 95% credible intervals are shown for the isotropic-PS and smooth isotropic templates in red and blue, respectively. The median intensity for the foreground is also included (green). The sum of all the components (yellow band) agrees with the total spectrum of the *Fermi* data (black). The spectrum of the 2FHL sources is provided in dashed black. Best-fit values for the intensities and source-count distributions in each energy bin are provided in Tables 7 and 8 of the Appendix. Note that, as for the 3FGL case, the spectra of 2FHL sources are driven at the high end by extrapolations from lower energies; the associated uncertainties are not shown here.

lower energies. We employ the recommended event-quality cuts as in the low-energy analysis and also choose $n_{\text{side}} = 128$ HEALPix pixelation. Results are presented for the three energy bins [50, 151], [151, 457], and [457, 2000] GeV. With *ultracleanveto* data, the weighted exposures in the energy bins are $[2.48 \times 10^{11}, 2.31 \times 10^{11}, 1.69 \times 10^{11}] \text{ cm}^2 \text{ s}$, while with source data the exposures become $[3.23 \times 10^{11}, 3.20 \times 10^{11}, 2.87 \times 10^{11}] \text{ cm}^2 \text{ s}$. For both data sets, the 68% containment radii are approximately [0.14, 0.12, 0.11] degrees. We will also

discuss results of analyses performed over a single wide-energy bin from [50, 2000] GeV.

3. SIMULATED-DATA STUDIES

To study the behavior of the NPTF, we apply it to simulated data sets of the gamma-ray sky. These results are crucial both for understanding systematics associated with the NPTF as well

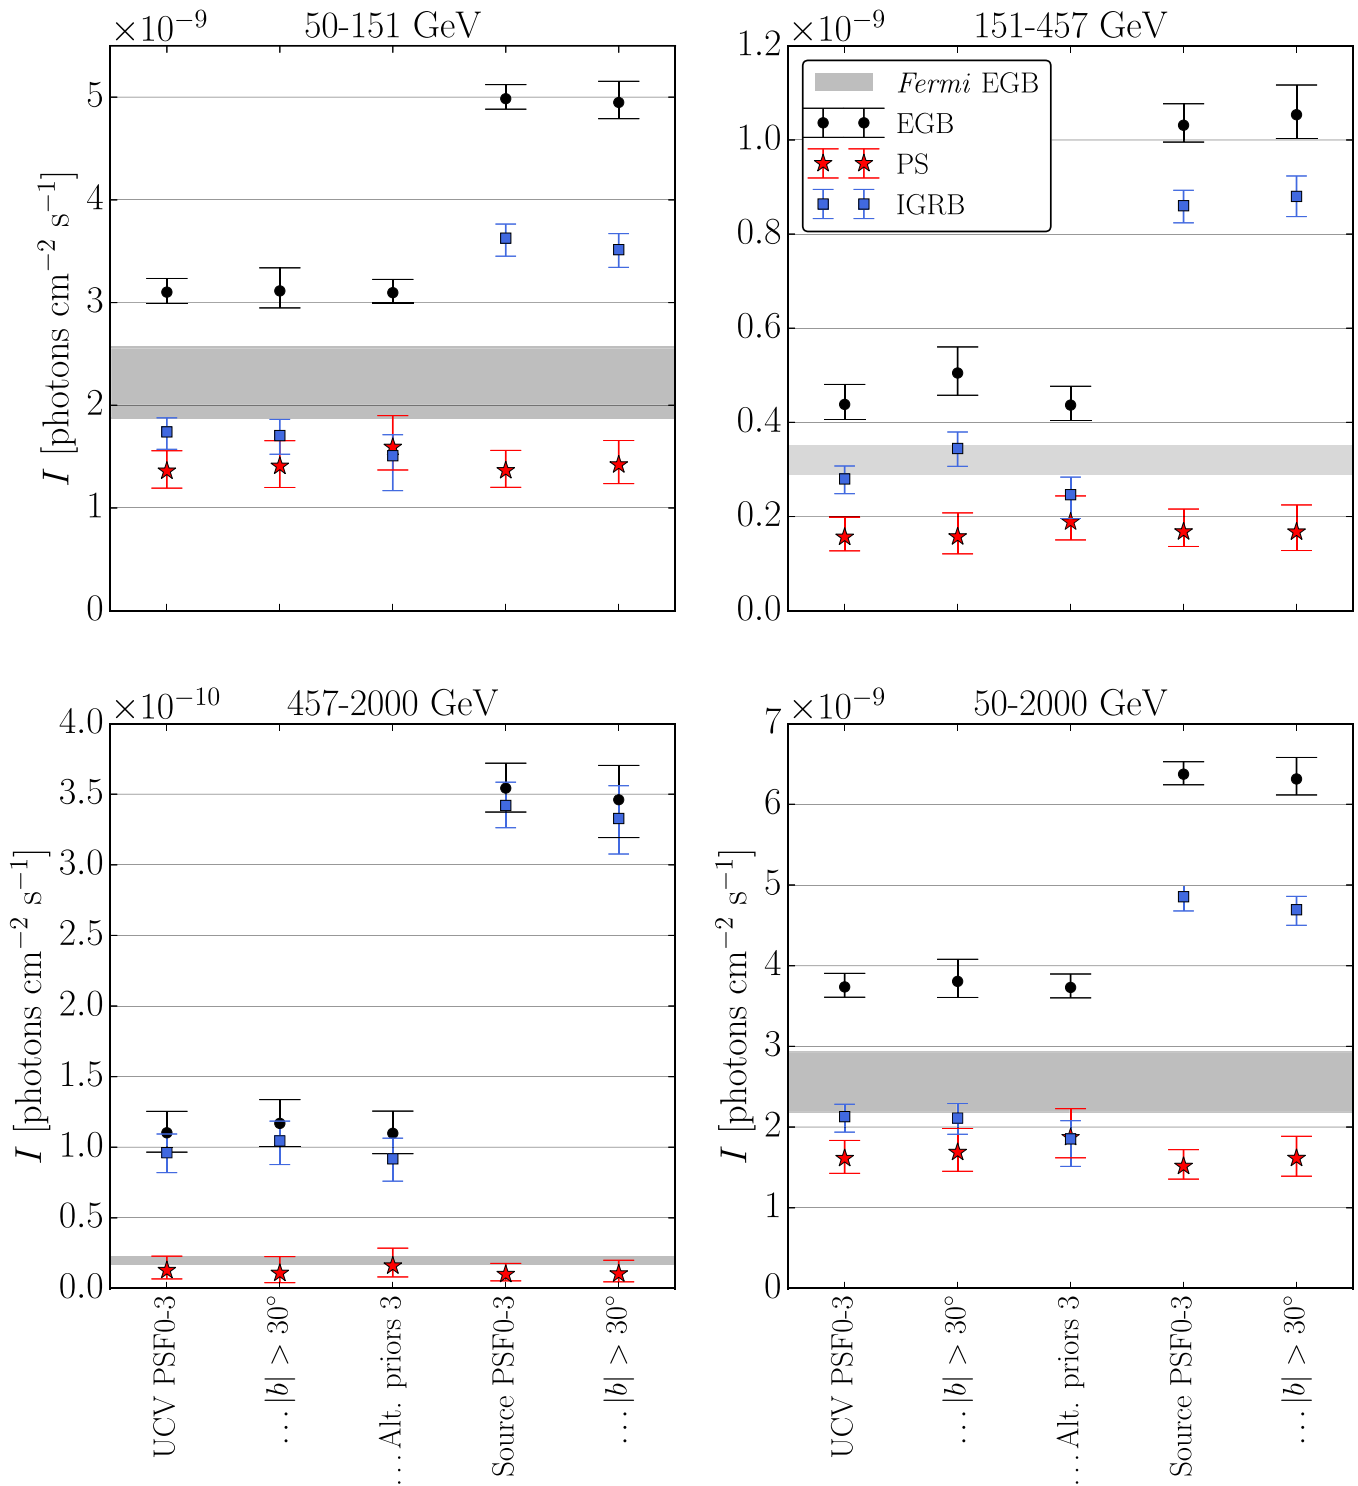


Figure 9. Comparison of the EGB (black circles), IGRB (blue squares), and PS (red stars) intensities recovered by the NPTF for the various systematic tests specific to high energies. The gray band indicates the systematic uncertainty associated with the measured *Fermi* EGB (Ackermann et al. 2015b). The associated source-count distributions for these analyses are provided in Appendix C.

as for interpreting the results of the NPTF in terms of evidence for or against the existence of these source populations.

A simulated-data map can be created starting from a particular source population that contributes to the EGB. Using a theory model for the energy spectrum and luminosity function (LF), the source-count distribution for that population can be derived in a specified energy range—see Appendix A.1 for further details on this procedure. The appropriate number of

sources is then drawn from this function and randomly distributed across the sky, with counts chosen to follow the intensity spectrum. Sources are then smeared with the appropriate Gaussian PSF to mimic the desired *Fermi* data set bin-wise in energy, and Poisson counts are drawn to obtain the simulated map for the population. This is then combined with the simulated contribution of the p_{8r2} foreground model and the *Fermi* bubbles, whose normalizations are determined




Figure 10. (Left) Best-fit source-count distribution in the wide-energy bin from 50 to 2000 GeV using all quartiles of Pass 8 *ultracleanveto* data. The black points indicate the 2FHL sources, and the blue line denotes the best-fit source-count from Ackermann et al. (2016b) that corresponds to the same energy bin. (Right) A comparison of the cumulative source-count distribution for the same analysis.

from the Poissonian template fits to the real data, as described in Section 2.

For most of this section, we simulate data corresponding to the PSF3 event type (best PSF quartile) of the *ultracleanveto* event class and focus on the following four energy bins: [1.89, 4.75], [4.75, 11.9], [11.9, 30], and [30, 94.9] GeV.¹¹ However, we also simulate data corresponding to the PSF1–3 (top 3 PSF quartiles) instrument response function to illustrate potential advantages in going to the more inclusive data set, albeit with a slightly worse PSF. Once the simulated-data maps are created, we run them through the NPTF analysis pipeline. First, we analyze the case where either blazars or SFGs fully account for the EGB, and then we analyze a perhaps more realistic scenario where both populations contribute significantly to the flux. The particular blazar and SFG models used here are merely meant for illustration. They are chosen as examples that span the range of possibilities between smooth and PS isotropic contributions. Because mAGNs are fainter and more numerous than blazars, they likely act similarly to SFGs in the context of the NPTF and so we do not consider them separately here. A detailed analysis of how the NPTF responds to the broader class of theoretical models for these source classes is beyond the scope of this work.

3.1. Blazars

Active galactic nuclei (AGNs) are the highly luminous central regions of galaxies where emission is dominated by accretion onto a supermassive black hole (Urry & Padovani 1995). If the black hole is spinning, then relativistic jets may also form. Blazars are a subclass of AGNs in which the jet is oriented within 14° of the line of sight (Angel & Stockman

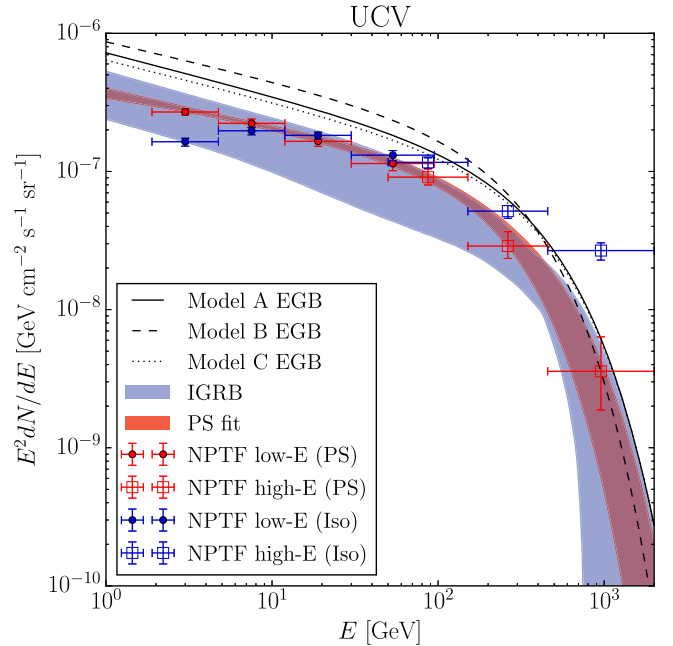


Figure 11. Global fit to the PS intensity spectrum recovered by the NPTF. The results of the NPTF low-energy analysis on *ultracleanveto* PSF1–3 data and the high-energy analysis on *ultracleanveto* PSF0–3 data are shown (filled red circles and open red boxes, respectively). The red band indicates the best-fit (68% credible interval) to a power law with exponential cut-off. For comparison, the best-fit *Fermi* EGB spectra from Ackermann et al. (2015b) are shown for three different diffuse-background models (Models A–C). The blue band indicates the estimated IGRB spectrum, obtained by subtracting the PS spectrum from the *Fermi* EGB; the spread includes the statistical uncertainty from the PS intensity as well as the systematic uncertainty on the EGB. We also plot the best-fit smooth isotropic spectrum recovered by the NPTF (filled blue circles and open blue boxes). The results are in good agreement with the estimated IGRB result (blue band) below ~ 100 GeV, but overestimate the result at higher energies due to cosmic-ray contamination.

¹¹ Potential systematic issues arising from going to lower energies are discussed in Appendix A.2.

Table 4
PS Fractions ($I_{\text{PS}}/I_{\text{EGB}}$) for the Low (PSF1–3) and High-energy (PSF0–3) Analyses, Using *ultracleanveto* Data, with Energy Sub-bins in Units of GeV

I_{EGB}	Low-energy Analysis				High-energy Analysis			
	1.89–4.75	4.75–11.9	11.9–30	30–94.9	50–151	151–457	457–2000	50–2000
Scenario A	$0.62^{+0.04}_{-0.02}$	$0.53^{+0.03}_{-0.03}$	$0.48^{+0.03}_{-0.03}$	$0.47^{+0.05}_{-0.04}$	$0.44^{+0.06}_{-0.05}$	$0.36^{+0.08}_{-0.06}$	$0.12^{+0.09}_{-0.06}$	$0.43^{+0.05}_{-0.04}$
Scenario B	$0.54^{+0.03}_{-0.03}$	$0.60^{+0.04}_{-0.03}$	$0.61^{+0.06}_{-0.05}$	$0.66^{+0.09}_{-0.07}$	$0.67^{+0.10}_{-0.09}$	$0.51^{+0.13}_{-0.09}$	$0.58^{+0.45}_{-0.27}$	$0.68^{+0.09}_{-0.08}$

Note. The first row (“Scenario A”) uses the EGB intensity obtained in this work using foreground model `p8r2`; however, this scenario likely overestimates the I_{EGB} at energies above ~ 100 GeV due to cosmic-ray contamination. The second row shows the PS fractions calculated with respect to the *Fermi* EGB intensity from Ackermann et al. (2015b), with foreground model A (“Scenario B”). Although the *Fermi* analysis uses a different foreground model, it takes advantage of a dedicated event selection above ~ 100 GeV that mitigates effects of additional contamination.

Table 5
Same as Table 2, Except Using the Top Three Quartiles (PSF1–3) of the Pass 8 *ultracleanveto* Data

Energy (GeV)	I_{EGB}	$I_{\text{iso}}^{\text{PS}}$	I_{iso} ($\text{cm}^{-2} \text{s}^{-1} \text{sr}^{-1}$)	I_{diff}	I_{bub}
1.89–4.75	$1.47^{+0.05}_{-0.04} \times 10^{-7}$	$8.99^{+0.53}_{-0.45} \times 10^{-8}$	$5.72^{+0.30}_{-0.36} \times 10^{-8}$	$3.28^{+0.01}_{-0.01} \times 10^{-7}$	$3.34^{+0.39}_{-0.41} \times 10^{-8}$
4.75–11.9	$5.60^{+0.21}_{-0.17} \times 10^{-8}$	$2.85^{+0.23}_{-0.18} \times 10^{-8}$	$2.74^{+0.12}_{-0.16} \times 10^{-8}$	$7.69^{+0.09}_{-0.08} \times 10^{-8}$	$1.52^{+0.22}_{-0.23} \times 10^{-8}$
11.9–30.0	$1.85^{+0.08}_{-0.07} \times 10^{-8}$	$8.54^{+0.85}_{-0.71} \times 10^{-9}$	$9.92^{+0.53}_{-0.63} \times 10^{-9}$	$1.68^{+0.04}_{-0.04} \times 10^{-8}$	$5.59^{+1.35}_{-1.31} \times 10^{-9}$
30.0–94.9	$6.19^{+0.42}_{-0.29} \times 10^{-9}$	$2.85^{+0.39}_{-0.33} \times 10^{-9}$	$3.36^{+0.25}_{-0.30} \times 10^{-9}$	$3.77^{+0.17}_{-0.19} \times 10^{-9}$	$1.10^{+0.72}_{-0.61} \times 10^{-9}$

1980). The spectral energy distribution of these objects is bimodal with a peak in the ultraviolet due to synchrotron radiation of electrons in the jet, and another peak in the gamma band from inverse Compton scattering of the same electrons (Fossati et al. 1998; Ghisellini et al. 1998, 2010; Boettcher et al. 2013). There is also the possibility of a hadronic contribution to blazar gamma-ray spectra; though, this is likely to be subdominant (Tavecchio 2014; Cerruti et al. 2015; Zdziarski & Böttcher 2015). Blazars may be further classified as either BL Lacertae (BL Lacs) or flat spectrum radio quasars (FSRQs), which are characterized by the absence or presence of broad optical/ultraviolet emission lines, respectively.

Before *Fermi*, few blazars had been identified in gamma-rays, and to estimate the size of this population, one had to extrapolate based on those observed at lower frequencies. However, *Fermi* brought the discovery of many more blazars in the gamma-ray band, making it possible to study their properties directly (Fermi-LAT 2010; Ajello et al. 2012, 2014; Di Mauro et al. 2014c). Most recently, 403 blazars (with $|b| > 15^\circ$) from the First LAT AGNs Catalog (Abdo 2010) were studied (Ajello et al. 2015). FSRQs and BL Lacs were considered together in the same sample to improve statistics. We use the best-fit luminosity and spectral energy distributions given in Ajello et al. (2015) (specifically, the luminosity-dependent density evolution, or LDDE, scenario) to model the blazar component in our simulated data and refer to it as the “Blazar-1” model. Alternatively, we also consider BL Lacs and FSRQs separately, adding up their respective contributions using the LDDE1 model from Ajello et al. (2014) and the LDDE model from Ajello et al. (2012), which we refer to as the “Blazar-2” model. This model predicts a much flatter source-count distribution below the *Fermi* detection threshold, with more low-flux sources. The two source-count models approximately bracket the current theoretical uncertainty in the faint-end slope of blazars, and we use them to study the response of different blazar models to the NPTF; though, this is meant to be purely illustrative and by no means exhaustive.

Figure 1 shows the best-fit source-count distributions recovered when the NPTF analysis is run on the Blazar-1 simulated-data map, assuming the PSF3 instrument response function. In each panel, the dark (light) red band is the 68% (95%) credible interval for the isotropic-PS source-count distribution as recovered from the posterior and constructed pointwise in flux. The red line shows the median source-count distribution, constructed in the same way. The dashed red curve, on the other hand, indicates the source-count distribution of the blazar model used to generate the simulated-data. A flux histogram of the simulated PSs for the particular realization shown here is given by the red points, with vertical error bars indicating the 68% credible interval associated with Poisson counting statistics on the number of sources in that bin. Notice that these error bars become large at high fluxes because there are very few sources per flux bin.

In general, the reconstructed source-count distribution is in good agreement with the input source-count distribution at intermediate fluxes, with uncertainties becoming large at low and high fluxes. At high flux, this is due to the fact that it is unlikely to draw a bright source from the underlying source-count distribution. At low fluxes, it is difficult to distinguish PS emission from genuinely isotropic emission. To illustrate this point, we also mark the flux that corresponds to a single photon on average (in the particular energy range, ROI, and event class) with the vertical dotted–dashed black line. At fluxes corresponding to counts near or below ~ 1 photon, it is difficult to distinguish PS emission from smooth emission with the NPTF, as evidenced by the growing uncertainties. In this low-flux regime, we do not expect that the NPTF will be able to fully recover the properties of the input source-count distribution.

The vertical dotted green lines in Figure 1 correspond to the fluxes above which 90%, 50%, and 10% (from left to right) of the photon counts are accounted for, on average, by sources with larger flux. Note that in the lowest energy bin, 90% of the flux arises from sources that contribute more than one photon. Moving toward higher energies, a larger fraction of the flux




Figure 12. Full-sky maps showing the value (clipped at 20) of $-\log \epsilon_p$ in each pixel p . The larger the value of $-\log \epsilon_p$, the more likely the pixel contains a point source. (Top) Results using *ultracleanveto* data (PSF3) for energies of 1.89–94.9 GeV. *Fermi* 3FGL sources are indicated by the white circles, with radii weighted by the predicted number of photon counts for a given source. (Bottom) Results using all quartiles of *ultracleanveto* data for 50–2000 GeV. Circles now represent *Fermi* 2FHL sources. The data for this figure is available upon request.

arises from sources that contribute less than one photon. In all energy bins, more than 50% of the flux is accounted for by sources that contribute more than a single photon each.

The corresponding energy spectra for the various templates are shown on the top left panel of Figure 2. As is evident, these blazars show up as PSs under the NPTF; indeed, the smooth isotropic flux (blue) is subdominant in each energy bin. Overlaid in dashed red is the spectrum for the simulated Blazar-1 sources. The sum of the smooth and PS isotropic components—which is simply the EGB intensity—is consistent with the simulated spectrum for the blazar model. The green curve shows the median of the posterior for the galactic diffuse model spectrum. The energy spectrum of the diffuse model is softer than that for blazars, so that the diffuse model dominates more at low energies than at high. The sum of the components

(yellow band) is consistent with the total flux in the simulated data (black lines) at 68% confidence.

As a contrasting example, we also simulate the Blazar-2 model, which predicts more low-flux sources than the previous example we considered. The best-fit source-count distributions for the Blazar-2 simulated maps are shown in Figure 3. Once again, we see good agreement between the input data and the recovered source-count distribution above the single-photon sensitivity threshold. In this case, however, the reference model predicts a larger fraction of flux coming from sources below this threshold. For example, about 50% of the flux comes from sub-single-photon sources in the second energy bin, and this fraction only increases further at higher energies. The corresponding energy spectrum is shown in the top right panel of Figure 2. As expected, an increasing amount of flux is

absorbed by the Poissonian isotropic template. However, the EGB spectrum, shown by the purple band, is still consistent with the input spectrum for the Blazar-2 model.

To further quantify the ability of the NPTF to reconstruct the blazar flux as PS emission, it is convenient to consider the ratios $I_{\text{iso}}^{\text{PS}}/I_{\text{blazar-sim}}$ in each energy bin, where $I_{\text{iso}}^{\text{PS}}$ is the PS intensity found by the NPTF and $I_{\text{blazar-sim}}$ is the blazar intensity in the simulation. For the Blazar-1 model, we find¹²

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.94_{-0.04}^{+0.05}, 0.88_{-0.05}^{+0.07}, 0.86_{-0.07}^{+0.08}, 0.64_{-0.07}^{+0.08}]$$

in each of the four respective energy bins, while for the Blazar-2 model, we find

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.74_{-0.04}^{+0.06}, 0.64_{-0.05}^{+0.07}, 0.53_{-0.06}^{+0.07}, 0.51_{-0.07}^{+0.09}],$$

for the particular Monte Carlo realizations shown.¹³ For the Blazar-2 scenario, more flux goes into smooth isotropic emission, which is why the PS fractions are correspondingly smaller in each energy bin. Note that, in both scenarios, the fraction of the blazar flux absorbed by the PS template decreases at higher energies, where the photon counts become less numerous and a higher fraction of the blazar flux is generated by sub-threshold sources. As a result, the intensities $I_{\text{iso}}^{\text{PS}}$ should be interpreted as lower bounds on the blazar flux; this intuition is validated by the fact that the ratios $I_{\text{iso}}^{\text{PS}}/I_{\text{blazar-sim}}$ tend to be less than unity.

Next, we explore whether including more quartiles of the *ultracleanveto* data, as ranked by PSF, increases our ability to reconstruct the blazar flux as PSs under the NPTF. When including more quartiles of data, there are two competing effects that determine our ability to constrain the PS flux: on the one hand, we increase the effective area, but on the other hand, we worsen the angular resolution of the data set. We investigate these effects by repeating the Monte Carlo tests described above using the PSF1–3 instrument response function. The results of the PSF1–3 tests are described in Appendix A.3, and here we simply quote the fractions

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.78_{-0.05}^{+0.06}, 0.81_{-0.06}^{+0.07}, 0.72_{-0.06}^{+0.06}, 0.57_{-0.05}^{+0.06}]$$

for a generic realization of the Monte Carlo simulations for the Blazar-2 model. The PSF1–3 event type increases our ability to distinguish between the blazar emission and smooth emission compared to the PSF3 event type.

3.2. Star-forming Galaxies

SFGs like our own Milky Way are individually fainter, though much more numerous, than blazars. The modeling of SFGs in the gamma-ray band is highly uncertain because *Fermi* has only detected eight SFGs thus far (Fornasa & Sánchez-Conde 2015). However, SFGs could still contribute a sizable fraction of the total flux observed by *Fermi*. Even though SFGs are PSs, their flux is expected to be dominated by a large population of dim sources degenerate with smooth isotropic

emission. Under the NPTF, therefore, we expect that the majority of their emission will be absorbed by the smooth isotropic template. To illustrate this point, we simulate SFGs using the LF and energy spectrum from Tamborra et al. (2014). In that work, input from infrared wavelengths was used to construct a model for the infrared flux from SFGs. Then, a scaling relation was used to convert from infrared to gamma-ray luminosities. The contributions from quiescent and starburst SFGs were considered separately, along with SFGs that host an AGN. Note, however, that other models predict less emission from SFGs than this particular case—see, e.g., Makiya et al. (2011), Inoue (2011), Ackermann et al. (2012b).

The results of the SFG-only simulations are described in Appendix A.4. We find that while the NPTF does detect a small PS component in the first few energy bins, as the result of a few SFGs above the sensitivity threshold of the NPTF in those energy bins, by far, most of the SFG emission is detected as smooth isotropic emission, with the ratio $I_{\text{iso}}^{\text{PS}}/I_{\text{iso}} \lesssim 1/100$ in all energy bins, where I_{iso} is the intensity of smooth isotropic emission. Moreover, the intensity I_{iso} is consistent with the simulated EGB (SFG flux) in all energy bins, at 68% confidence.

3.3. Blazar and SFG Combination

A perhaps more realistic scenario for testing the NPTF is to consider a scenario where both SFGs and blazars contribute to the EGB. Therefore, we create simulated maps that include both components and test them on the NPTF. The recovered energy spectra for the SFG + Blazar-1 (Blazar-2) example is shown in the bottom left (right) panel of Figure 2. In both cases, the PS spectrum is consistent with that found in the blazar-only simulations, which are shown in the top panels in that figure. The reconstructed source-count distributions for these examples are not shown, as they are consistent with those found in the blazar-only cases.

In the case of the Blazar-1 model, the spectra of the smooth isotropic emission and the PS emission trace the spectra of the input SFG population and blazar population, respectively. In the case of the Blazar-2 model, the PS flux is further below the input blazar spectrum, as was found in the blazar-only simulations. However, the smooth isotropic emission is further above the simulated SFG spectrum. In both cases, the sum of the smooth isotropic emission and PS emission (EGB) is consistent with the simulated blazar plus SFG flux.

There is, in fact, a subtle difference between the PS distribution recovered with and without the addition of an SFG population. The difference becomes noticeable when comparing the fractions

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.97_{-0.05}^{+0.06}, 1.00_{-0.09}^{+0.11}, 0.87_{-0.07}^{+0.09}, 0.72_{-0.09}^{+0.12}]$$

for SFG + Blazar-1 and

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.80_{-0.06}^{+0.08}, 0.59_{-0.06}^{+0.07}, 0.59_{-0.06}^{+0.08}, 0.43_{-0.05}^{+0.06}]$$

for SFG + Blazar-2 to the corresponding values for the blazar-only simulations. In the simulations with SFGs, the fractions $I_{\text{iso}}^{\text{PS}}/I_{\text{blazar-sim}}$ are generally higher and have larger uncertainties. The reason for this is that the SFG emission is degenerate with an enhanced sub-threshold component to the PS source-count distribution.

¹² Throughout this work, best-fit values indicate the 16th, 50th, and 84th percentiles of the appropriate posterior probability distributions.

¹³ Different Monte Carlo realizations are found to induce variations consistent with the quoted statistical uncertainties, generally on the order of 5%.




Figure 13. Best-fit source-count distribution for a simulated map containing both Blazar-1 and SFG sources in the 0.475–0.753 (left) and 0.753–1.89 (right) GeV energy bins. (Formatted as in Figure 1.)

Simulating data with the PSF1–3 instrument response function, we find that the ratios $I_{\text{iso}}^{\text{PS}}/I_{\text{blazar-sim}}$ are somewhat closer to unity than in the PSF3 case. In particular, for the SFG + Blazar-2 model simulations,

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [1.03_{-0.13}^{+0.20}, 0.73_{-0.05}^{+0.06}, 0.66_{-0.06}^{+0.07}, 0.57_{-0.06}^{+0.07}].$$

The improved exposure allows the NPTF to probe lower fluxes and to therefore recover a larger fraction of the isotropic-PS emission.

4. LOW-ENERGY ANALYSIS: 1.89–94.9 GeV

The findings from the previous section illustrate that the NPTF procedure is able to set strong constraints on the PS (e.g., blazar) and smooth Poissonian (e.g., SFGs and mAGNs) contributions to the EGB. In this section, we focus on the energy range from 1.89 to 94.9 GeV, and begin by presenting the results of our benchmark analysis on the real *Fermi* data. This is followed by a detailed discussion of potential systematic uncertainties and their effects on the conclusions.

4.1. Pass 8 Ultracleanveto Data

4.1.1. Top PSF Quartile

We begin by analyzing the *ultracleanveto* PSF3 data for $|b| \geq 30^\circ$, using the `p8r2` foreground model. This is referred to as the “benchmark analysis” throughout the text. Table 2 provides the best-fit intensities for each spectral component, as a function of energy, and the best-fit spectra are plotted in the left panel of Figure 4. The `p8r2` diffuse model is shown in green (median only), while the smooth isotropic and isotropic-PS posteriors are shown by the blue and red bands, respectively. The best-fit spectrum for PSs with $|b| > 30^\circ$ in the 3FGL catalog (Acero et al. 2015) is shown by the dashed black line in Figure 4; the spectrum as plotted should be treated

with care as systematic uncertainties are not properly accounted for. In particular, the 3FGL catalog includes sources between 0.1 and 300 GeV. At the high end of this range, the spectrum is driven to a large extent by extrapolations from lower energies, where the statistics are better. The potential errors associated with such extrapolations are difficult to quantify and are not shown in Figure 4. As a result, a direct comparison between the 3FGL spectrum and our results is difficult to make, especially in the highest energy bins. For this reason, we have a dedicated NPTF study for energies greater than 50 GeV in Section 5. Those results are compared to the *Fermi* 2FHL catalog (Ackermann et al. 2016b), which is explicitly constructed at higher energies and is likely a more faithful representation of above-threshold PSs in this regime.

The source-count distributions reconstructed from the NPTF are shown in Figure 5, with best-fit parameters provided in Table 3. For comparison, the binned 3FGL source-count distributions are also plotted; the vertical error bars represent 68% statistical uncertainties and do not account for systematic uncertainties. A few trends are clearly visible. First, each flux break tends to have large uncertainties. This may be a reflection of the fact that the real source-count distribution is not a simple triply broken power law, but rather a more complicated function, as in the blazar simulations of Section 3. Therefore, the best-fit values for each of these parameters, when viewed independently, may be somewhat deceptive. As is evident in Figure 5, the posteriors for the breaks and indices are distributed in such a way as to describe a smooth concave function for $F^2 dN/dF$.

At very high and very low flux, the uncertainties on the indices (n_1 and n_4 , respectively) become large. At high flux, this is simply due to the fact that there are very few sources, so the source-count distribution falls off rapidly. At low flux, the large uncertainties on n_4 arise from the difficulty in distinguishing the isotropic-PS contribution from its smooth counterpart. Indeed, below the single-photon boundary (dotted–dashed black line),




Figure 14. Best-fit source-count distribution, as a function of energy, obtained from simulating the Blazar-2 model with the top three quartiles of *ultracleanveto* data. (Formatted as in Figure 1.)

the NPTF analysis starts to lose sensitivity. The posterior distributions for the slopes above (below) the highest (lowest) break are highly dependent on the priors and so the quoted values in Table 3 should be treated with care.

The presence of any distinctive breaks encodes information about the number of source populations as well as their evolutionary properties. In all energy bins, we see that the NPTF places the lowest break, $F_{b,3}$, close to the one-photon sensitivity threshold and the highest break, $F_{b,1}$, in the vicinity of the highest-flux 3FGL source (see Table 3 for the exact values). The evidence for an additional break, $F_{b,2}$, at intermediate fluxes varies depending on the energy bin. From 1.89 to 4.75 GeV, there is strong indication for a break at fluxes

$\sim 10^{-10}$ $\text{ph cm}^{-2} \text{s}^{-1}$, with the index $n_2 \approx 2.04$ above the break hardening to $n_3 \approx 1.74$ below the break. In the two subsequent energy bins, up to ~ 30 GeV, we also find evidence that the source-count distribution hardens as we move from high fluxes to below the second break, with the index n_3 below the second break ~ 1.9 – 2.0 in both cases. In the last bin, the uncertainties are too large to determine if the source-count distribution changes slope at any flux above the lowest break $F_{b,1}$.

4.1.2. Top Three PSF Quartiles

The benchmark analysis described in the previous section used only the top quartile (PSF3) of the Pass 8 *ultracleanveto* data set. This restriction selects events with the best angular




Figure 15. Best-fit source-count distribution for a simulated map containing SFGs (Tamborra et al. 2014). (Formatted as in Figure 1.)

resolution, but at the price of reducing the total photon count. In Section 3, we showed that including the top three quartiles of the Pass 8 *ultracleanveto* data may help constrain the source-count distribution at low fluxes. With that in mind, we now investigate how the results of the benchmark analysis change when using the PSF1–3 *ultracleanveto* data set.

In general, the best-fit intensities for the individual spectral components are consistent within uncertainties with those obtained using only the top quartile of data. (See Table 5 for specific values.) The PS flux does increase slightly in going from PSF3 to PSF1–3 in the upper energy bins due to the increased exposure. More specifically, the ratios of the median PS intensities measured with *ultracleanveto* PSF1–3 data to those measured with PSF3 data are [1.00, 1.06, 1.19, 1.19] in

the four increasing energy bins. This can also be seen in the associated spectral intensity plot (right panel of Figure 4), where the red bands are further above the 3FGL line in the last energy bins than in the corresponding plot for the PSF3 analysis (left panel). The intensity of the EGB is seen to increase slightly, in all energy bins, when going from PSF3 to PSF1–3 data, potentially suggesting additional cosmic-ray contamination with the looser photon-quality cuts, though the increases in EGB intensities are within statistical uncertainties.

The best-fit source-count distributions recovered by the NPTF with PSF1–3 data are shown in Figure 6. For reference, the blue curve shows the best-fit for the PSF3-only analysis. The most important difference between the PSF3 and PSF1–3 results is that the source-count distributions extend to lower flux with




Figure 16. Energy spectra for a simulated map containing SFGs (Tamborra et al. 2014). (Formatted as in Figure 2.)

PSF1–3 data. This is due to the fact that the exposure in each energy bin, averaged over the ROI, is larger for the top three quartiles compared to the top quartile alone. As a result, the flux corresponding to single-photon detection is lower (compare the vertical dotted–dashed line in Figure 6 with that in Figure 5), which improves the NPTF reach. Thus, the PSF1–3 analysis is sensitive to more sub-threshold sources. Note that the same trend was observed in the simulation tests in Section 3 in going from PSF3 to PSF1–3 data sets.

Other than the location of the lowest break, which is lower due to the increased exposure, all other source-count distribution parameters are consistent, within uncertainties, between analyses. At the lowest energy, the break at $F_{b,2} \sim 10^{-10}$ photons $\text{cm}^{-2} \text{s}^{-1}$ is even more pronounced, with an index of $n_2 \sim 2.10$ above the break and $n_3 \sim 1.75$ below the break. In the highest energy bin, the structure observed in the source-count distribution for the benchmark analysis has smoothed out.

4.2. Systematic Tests

The previous subsection illustrated how the results of the NPTF change when additional *ultracleanveto* PSF quartiles are included in the analysis. We also tested the stability of our analysis to variations in the ROI, *Fermi* event class, foreground modeling, *Fermi* bubbles, PSF modeling, and choice of priors. These systematic tests are described in detail in Appendix B.2.

Figure 7 briefly summarizes the results. The EGB intensity as measured by *Fermi* is shown by the gray band. To obtain this band, we use the best-fit power-law spectrum with the exponential cut-off provided in Ackermann et al. (2015b); the width of the gray band is found by varying between best-fit values for the three foreground models considered in that paper (Models A/B/C) and does not include statistical uncertainties, which become increasingly important at high energies. The smooth isotropic intensity, and thus the intensity of the EGB, is subject to large systematic uncertainties. As expected, the

variation in smooth isotropic intensity is most pronounced when using the source event class, which contains more cosmic-ray contamination. However, the spectrum of emission from PSs as captured by the NPTF appears robust to all the systematic effects considered here. This is the primary conclusion of this subsection.

5. HIGH-ENERGY ANALYSIS: 50–2000 GeV

We now consider the NPTF results at high energies from 50 to 2000 GeV. The number of photons available decreases when moving to higher energies, so we loosen the restrictions on the PSF quartiles to maximize the sensitivity potential of the NPTF. In this section, the majority of the analyses are done using all quartiles of the *ultracleanveto* data, though we also show results using all quartiles of *source* data. For the same reason, we widen the ROI to $|b| > 10^\circ$ rather than 30° ; though, the results are not sensitive to this cut, as we will show.

The best-fit energy spectra recovered by the NPTF analysis for the high-energy study of *ultracleanveto* data is shown in the bottom right panel of Figure 8.¹⁴ The fit results are compared with the best-fit energy spectrum for sources in *Fermi*'s 2FHL catalog (Ackermann et al. 2016a; dashed black line). This recently published catalog is based on 80 months of data and focuses on hard sources in the range from 50 to 2000 GeV. Statistical and systematic uncertainties are not accounted for in the determination of the 2FHL spectrum in Figure 8; these are likely non-negligible, especially at the highest energies.

The best-fit source-count distributions for the three energy bins are also shown in Figure 8, in the top row and bottom left panel. The black points in those panels denote the 2FHL source-count distributions, with vertical error bars indicating 68% Poisson errors. The statistical errors on the 2FHL sources are large due to the fact that there are not many sources. In all energy bins, the NPTF places the lowest break close to the single-photon sensitivity threshold (vertical dotted–dashed line) and the highest break in the vicinity of the brightest 2FHL source, just as in the low-energy analysis. Most notably in the 50–151 GeV bin, the NPTF probes unresolved sources with fluxes nearly an order-of-magnitude below the apparent 2FHL threshold. We find no evidence for an additional intermediate-flux break in any of the energy bins; though, it is difficult to make conclusive statements due to the large uncertainties in the individual source-count distributions.

We have completed a number of systematic tests of the high-energy analyses that include looking at all quartiles of the source data, requiring $|b| > 30^\circ$ for both event classes, and using the third alternate prior choice, with $n_4 > 1$. The results are summarized in Figure 9, and the source-count distributions for each of the systematic tests are given in Appendix C. Importantly, the isotropic-PS intensity is consistent across all the tests. However, the EGB intensities recovered by the NPTF are, in general, higher than those measured by *Fermi*. This discrepancy is likely due to increased cosmic-ray contamination above ~ 100 GeV, as suggested by the high IGRB intensities recovered by the NPTF at these energies. Indeed, the *Fermi* EGB study on Pass 7 data (Ackermann et al. 2015b) used dedicated event classes with specific data cuts to minimize such contributions. Such an analysis is beyond the scope of this work, as our primary focus is on the PS populations. We

¹⁴ The intensities and best-fit values for the individual model parameters are given in Appendix C.

Table 6
Same as Table 3, Except Using the Top Three Quartiles (PSF1–3) of the Pass 8 *ultracleanveto* Data

Energy (GeV)	n_1	n_2	n_3	n_4	$F_{b,1}$	$F_{b,2}$ ($\text{cm}^{-2} \text{s}^{-1}$)	$F_{b,3}$
1.89–4.75	$3.97^{+0.69}_{-0.83}$	$2.10^{+0.04}_{-0.03}$	$1.75^{+0.09}_{-0.14}$	$-0.45^{+1.10}_{-1.01}$	$8.16^{+6.26}_{-3.00} \times 10^{-9}$	$6.63^{+3.45}_{-2.31} \times 10^{-11}$	$3.85^{+1.38}_{-1.73} \times 10^{-12}$
4.75–11.9	$3.94^{+0.68}_{-0.85}$	$2.08^{+0.08}_{-0.05}$	$1.94^{+0.07}_{-0.19}$	$-0.46^{+1.07}_{-1.01}$	$3.85^{+2.69}_{-1.65} \times 10^{-9}$	$8.05^{+7.34}_{-5.66} \times 10^{-11}$	$4.00^{+1.36}_{-1.64} \times 10^{-12}$
11.9–30.0	$3.70^{+0.82}_{-0.90}$	$2.20^{+0.17}_{-0.14}$	$2.02^{+0.07}_{-0.08}$	$-0.37^{+1.05}_{-1.04}$	$2.69^{+1.37}_{-1.34} \times 10^{-9}$	$1.19^{+0.45}_{-0.63} \times 10^{-10}$	$3.78^{+1.33}_{-1.50} \times 10^{-12}$
30.0–94.9	$3.54^{+0.94}_{-0.90}$	$2.15^{+0.24}_{-0.28}$	$2.21^{+0.13}_{-0.11}$	$-0.09^{+1.06}_{-1.20}$	$1.59^{+0.75}_{-0.73} \times 10^{-9}$	$1.07^{+0.49}_{-0.59} \times 10^{-10}$	$3.37^{+1.51}_{-1.48} \times 10^{-12}$

simply caution the reader that the derived intensity for the smooth isotropic component in the high-energy analyses is subject to potentially large contamination.

It is possible to make stronger statements about the best-fit source-count distribution at high energies if we consider the wide-energy bin from 50 to 2000 GeV. The results are shown in the left panel of Figure 10. Due to the improved statistics, the uncertainties on the source-count distribution are smaller than those for the three sub-bins. Other than the low-flux sensitivity break, the NPTF finds no preference for an additional break. The intermediate-flux break, $F_{b,2}$, is essentially unconstrained as a result, and the power-law slope above (below) it are consistent within uncertainties: $n_2 = 2.28^{+0.28}_{-0.22}$ and $n_3 = 2.17^{+0.12}_{-0.09}$, respectively. We compare this result to the best-fit source-count distribution (blue line) published by *Fermi* for sources in this same energy range (Ackermann et al. 2016b). There are important differences between the two analyses. In the *Fermi* study, simulated maps were created using several different source-count distributions, parametrized as singly broken power laws. The histogram of the photon-count distribution for each of these maps, averaged over the full ROI, was compared to the actual data, and a fit was done to select the simulated maps that most closely resembled the data. This method is related to but in many ways distinct from the NPTF. The NPTF considers the difference between Poissonian and non-Poissonian photon probability distributions at the pixel-by-pixel level, instead of averaging the distributions over the full region. Moreover, in our analysis, we rely on semi-analytic techniques to calculate the photon-count probability distributions as we scan over the space of model parameters, instead of relying on Monte Carlo samples to numerically construct these distributions. As a result, we are able to consider source-count distributions with additional degrees of freedom and also scan over the normalizations of all of the background templates, which tend to be well determined given the pixel-by-pixel nature of the fit. In contrast, the intensity of all Poissonian models in Ackermann et al. (2016b), including the smooth isotropic emission, was kept fixed while scanning over the source-count distribution degrees of freedom.

The cumulative source-count plot is provided in the right panel of Figure 10. Our result is in good agreement with the 2FHL sources above the catalog sensitivity threshold $\sim 10^{-11} \text{ ph cm}^{-2} \text{ s}^{-1}$. In the first few flux bins above this threshold, there appear to be more 2FHL sources than what is predicted by the NPTF; though, the results are still consistent within uncertainties. This may be due to the Eddington bias (Eddington 1913) where extra sources are observed above threshold due to upward statistical fluctuations from sources immediately below.

Based on the results in Figure 10, we can project the expected number of these sources that may be observed by the CTA (CTA Consortium 2011; Dubus et al. 2013). For

energies above 50 GeV, the CTA flux sensitivity is $\sim 2.93 \times 10^{-12} \text{ cm}^{-2} \text{ s}^{-1}$ for 50 hr of observation per field of view (5σ detection).¹⁵ For 250 hr total of observation time, this covers $\sim 190 \text{ deg}^2$ of the sky, assuming a 7° field of view. As shown in Figure 10, the NPTF predicts a density of $0.029^{+0.008}_{-0.005} \text{ deg}^{-2}$ for sources above this threshold. This translates to $5.51^{+1.52}_{-0.95}$ detected sources, more than double what had previously been estimated for similar observing parameters (Dubus et al. 2013). Relaxing the observing time per source and assuming, as in Ackermann et al. (2016b), that a quarter of the sky is surveyed in 240 hr at 5m Crab sensitivity, then the NPTF predicts 161^{+30}_{-20} sources. This is lower, and in slight tension, with the 200 ± 45 sources predicted by the *Fermi* study using the blue source-count distribution illustrated in Figure 10.

6. DISCUSSION AND CONCLUSIONS

The primary focus of this paper is to characterize the properties of the PSs contributing to the EGB in a data-driven manner. To achieve this, we use a novel analysis method, referred to as NPTF, which takes advantage of photon-count statistics to distinguish diffuse and PS contributions to gamma-ray maps with non-trivial spatial variations. We presented the NPTF results on *Fermi* Pass 8 data at low (1.89–94.9 GeV, $|b| > 30^\circ$) and high (50–2000 GeV, $|b| > 10^\circ$) energies. For the first time, the intensity and source-count distributions for the isotropic PSs have been obtained as a function of energy, up to 2 TeV. The best-fit source-count distributions probe fluxes below the current detection threshold for the *Fermi* 3FGL and 2FHL catalogs, providing information on the unresolved populations.

Through extensive studies of how the NPTF responds to simulated populations, we have shown that the analysis procedure reproduces the properties of input source classes. Therefore, the features of the best-fit source-count distributions obtained from the data provide a potential wealth of information about the source populations of the EGB. While a detailed interpretation of the source-count distributions in terms of particular theoretical models is beyond the scope of this paper, several important trends were observed.

In this work, the source-count distributions are parametrized as triply broken power laws in the NPTF. At all energies, a break is fit at low (high) fluxes, below (above) which the analysis method loses sensitivity. Of particular interest is whether an additional break, $F_{b,2}$, is preferred at intermediate flux. We find a break in the lowest energy bin (1.89–4.75 GeV) at $1.22^{+2.00}_{-0.56} \times 10^{-10} \text{ cm}^{-2} \text{ s}^{-1}$ with slope $2.04^{+0.05}_{-0.05}$ above and $1.74^{+0.19}_{-0.37}$ below. In the subsequent two energy bins, 4.75–11.9 GeV and 11.9–30.0 GeV, there is a mild indication

¹⁵ <https://portal.cta-observatory.org/Pages/CTA-Performance.aspx>




Figure 17. Triangle plot for the 1.89–4.75 GeV bin. The posterior distributions correspond to the NPTF analysis for Pass 8 *ultracleanveto* PSF3 data using the `p8r2` foreground model. The *Fermi* bubbles intensity is defined relative to the interior of the bubbles, while the intensities of the other templates are computed with respect to the region $|b| \geq 30^\circ$.

that the source-count distribution hardens below the intermediate-flux break, though the change in slope is not as robust and significant as in the lowest energy bin. At higher energies, above ~ 30 GeV, there is no indication that the source-count distribution changes slope at the intermediate break. This trend is in line with the expectations from the blazar simulations in Section 3. For example, in both Figures 1 and 3 (see also Figure 14), which show the results of the NPTF run on simulated data with the Blazar-1 and Blazar-2 models, we find evidence for curvature in the source-count distribution at

intermediate fluxes in the lowest energy bins, while at higher energies the recovered source-count distribution appears as a single power law at fluxes above the sensitivity threshold of the NPTF. In the energy bin from 50 to 2000 GeV the best-fit value for $F_{b,2}$ is essentially unconstrained and the slopes above and below it are consistent within uncertainties: $2.28^{+0.28}_{-0.09}$ and $2.17^{+0.12}_{-0.22}$.

The NPTF also provides the best-fit intensities for the isotropic-PS populations as a function of energy. Figure 11 illustrates this spectrum for analyses done using the *ultracleanveto* event class.




Figure 18. Same as Figure 17, except for 4.75–11.9 GeV.

The filled red circles (open red boxes) show the results for the dedicated low (high)-energy analysis, with PSF1–3 data used at low energies and PSF0–3 data at high energies. For comparison, the *Fermi* EGB spectrum is shown by the black line (Ackermann et al. 2015b). This corresponds to the best-fit intensity using the Model A diffuse background from that study. To illustrate the systematic uncertainty on this curve, we also plot the spectra for diffuse models B and C (dashed and dotted, respectively).

The PS fraction, defined as $I_{\text{PS}}/I_{\text{EGB}}$, is provided in Table 4 for each energy bin. While using the EGB intensity derived in this work (“Scenario A”) is the most self-consistent comparison, this may underestimate the PS contribution above

~ 100 GeV, where the NPTF appears to recover too much smooth isotropic emission due to increased cosmic-ray contamination in the data sets used, as already discussed. Therefore, we also show the PS fractions calculated relative to the *Fermi* EGB intensity from Ackermann et al. (2015b) for diffuse model A (“Scenario B”). The comparison to the EGB as measured in Ackermann et al. (2015b) is not fully self-consistent, since, for example, the foreground modeling and data sets in Ackermann et al. (2015b) differ from those used in this work to measure I_{PS} . However, the advantage of this comparison is that the *Fermi* analysis uses special event-quality cuts to mitigate contamination, and thus their measure of I_{EGB}




Figure 19. Same as Figure 17, except for 11.9–30.0 GeV.

is likely more faithful than that presented in this work. These results are shown in the second row of Table 4. For the low-energy analysis, the PS fractions are consistent, within uncertainties, when I_{EGB} is taken from our work or *Fermi*'s.¹⁶ The substantial differences occur at high energies, where our result is systematically lower than the fractions based on *Fermi*'s EGB intensity.

¹⁶ For “Scenario B,” the quoted uncertainties only include those measured in this work for I_{PS} . For I_{EGB} , we use the best-fit value given in Ackermann et al. (2015b).

In general, we find that approximately 50%–70% of the EGB consists of PSs in the energy ranges considered. To interpret these results, we use the ratios $I_{\text{ISO}}^{\text{PS}}/I_{\text{blazar-sim}}$ obtained in the simulation studies of Section 3. In that section, we showed that the efficiency for the NPTF to recover the flux for the Blazar-2 model (with PSF1–3) is $\sim 100\%$ in the first energy bin and drops to $\sim 60\%$ in the fourth energy bin. For the Blazar-1 model, the efficiencies are consistently higher than the Blazar-2 scenario. These two blazar models are meant to illustrate extreme scenarios, with the Blazar-1 model having a significant fraction of the total flux arising from high-flux sources, while low-flux sources dominate instead in the Blazar-2 case. The




Figure 20. Same as Figure 17, except for 30.0–94.9 GeV.

high efficiency of the NPTF to recover the blazar component at low energies, combined with the PS fractions observed in the data (Table 4), clearly suggests that there is a substantial non-blazar component of the EGB up to energies of ~ 30 GeV. The interpretation of the results in the energy bin from 30.0 to 94.9 GeV is less clear. A proper interpretation of the results at higher energies in terms of evidence for or against a non-blazar component of the EGB requires dedicated blazar simulations, which we leave to future work.

Our results tend to predict fewer PSs (and photons from PSs), where we do overlap with previous studies. For example,

a similar photon-count analysis was used by Zechlin et al. (2016b) to study 1–10 GeV energies in the Pass 7 Reprocessed data. They found an $\sim 80\%$ PS fraction at these energies. At the lowest energies that we probe—which admittedly do not extend down as low as ~ 1 GeV—we only find a $\sim 54\%$ PS fraction (relative to Model A). Systematic uncertainties, as shown in Figure 7, can affect the recovered PS intensities at the $\mathcal{O}(10\%)$ level, which can partially alleviate the tension between our results.

Above 50 GeV, the NPTF procedure predicts that $0.68^{+0.09}_{-0.08}$ of the EGB consists of PSs, with systematic uncertainties




Figure 21. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and `p8r2` foreground model, but with $|b| > 10^\circ$. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

estimated at approximately $\pm 10\%$. This fraction is smaller, and in slight tension, with the predicted value $0.86^{+0.16}_{-0.14}$ obtained in previous work (Ackermann et al. 2016b). The fact that our results suggest that there is more diffuse isotropic emission at high energies may help alleviate the tension between Ackermann et al. (2016b) and the hadronuclear (pp) interpretation of IceCube’s PeV neutrinos (Murase et al. 2013). Some models suggest, for example, that these very-high-energy neutrinos are produced in hadronuclear interactions, along with high-energy gamma-rays that would contribute to the IGRB (Murase et al. 2013; Tamborra et al. 2014; Ando et al. 2015; Hooper et al. 2016). If the smooth isotropic gamma-ray

spectrum (i.e., the non-blazar spectrum) is suppressed above 50 GeV in the *Fermi* data, it could put such scenarios in tension with the data (Bechtol et al. 2015; Murase et al. 2016); however, that does not necessarily appear to be the case given the results of our analysis. With that said, and as already mentioned, dedicated blazar simulations at high energies are needed to properly interpret our results at these energies.

The PS spectrum in Figure 11 is well-modeled ($\chi^2 = 1.18$) as a power law with an exponential cut-off:

$$\frac{dN}{dE} = C \left(\frac{E}{0.1 \text{ GeV}} \right)^{-\gamma} \exp \left(-\frac{E}{E_{\text{cut}}} \right), \quad (5)$$

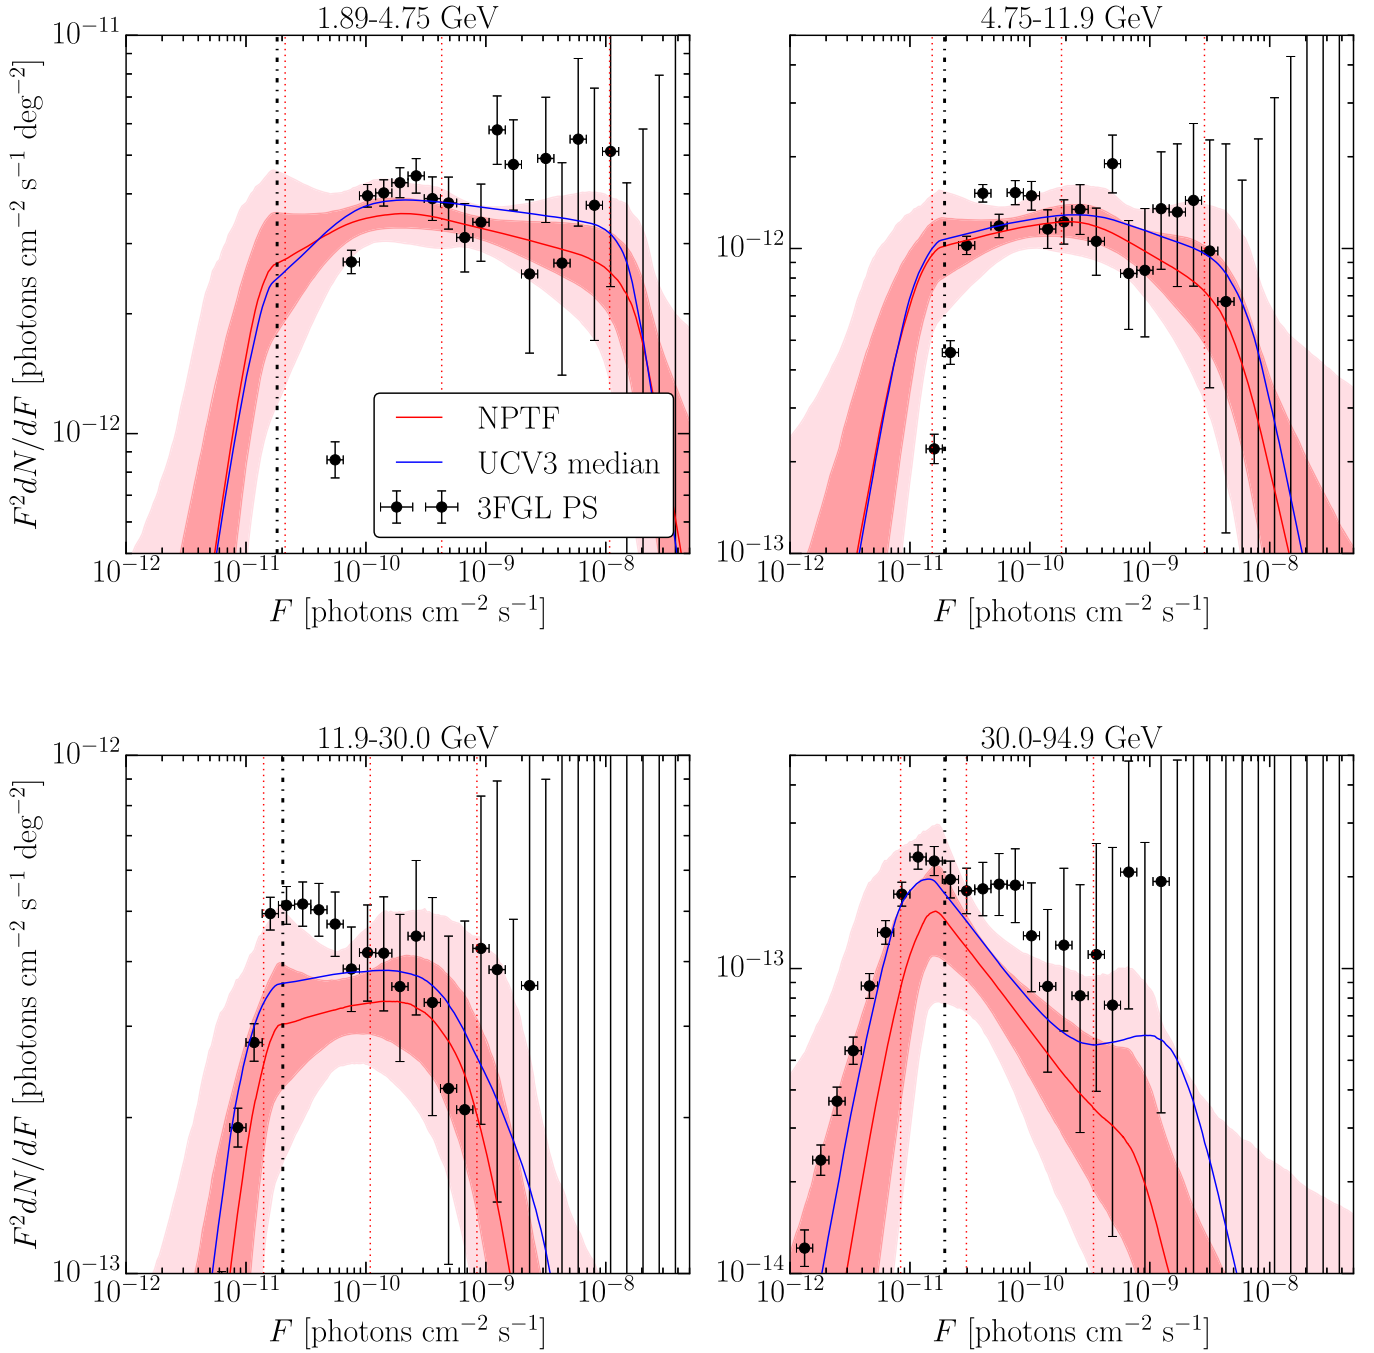


Figure 22. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and $p8r2$ foreground model, but with $b > 30^\circ$. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

where $C = 6.91^{+1.44}_{-1.29} \times 10^{-5} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$, $\gamma = 2.26^{+0.05}_{-0.05}$, and $E_{\text{cut}} = 289^{+127}_{-86.3} \text{ GeV}$ are the best-fit parameters.¹⁷ Note that the fit is done taking into account the uncertainties on the PS intensities in the energy sub-bins. The global fit for the PS spectrum is shown in Figure 11 by the red band, which denotes the 68% credible interval. Interestingly,

¹⁷ Repeating the fit using the results from the NPTF analyses with source data returns similar results, though the PS spectrum is slightly enhanced relative to the *ultracleanveto* result. In particular, with source data, we find $C = 7.98^{+1.58}_{-1.40} \times 10^{-5} \text{ GeV}^{-1} \text{ cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$, $\gamma = 2.29^{+0.04}_{-0.05}$, and $E_{\text{cut}} = 325^{+117}_{-78.1} \text{ GeV}$, with $\chi^2 = 0.93$.

the index γ and cut-off E_{cut} that we extract from the fit are very similar to the values found in Ackermann et al. (2015b), which used the same functional form to fit the EGB spectrum. Subtracting our PS spectrum from the EGB spectral fits gives the blue band in Figure 11. The band includes statistical uncertainties from our global fit as well as systematic uncertainties associated with varying between Models A–C. The blue band is an estimate of the IGRB spectrum and we compare it to the smooth isotropic spectrum recovered by the NPTF (blue points). Note that the two are consistent, within the large uncertainties, below $\sim 100 \text{ GeV}$; above this energy, our IGRB value is expectedly high.




Figure 23. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and $p8r2$ foreground model, but with $b < -30^\circ$. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

The NPTF allows us to make statistical statements about the properties of source populations contributing to the EGB, but at the expense of identifying the precise locations of these sources. However, it is still possible to make probabilistic statements about these locations. To do so, we compare the observed photon count in a given pixel, n_p , to the mean expected value, μ_p , without accounting for PSs. To determine μ_p we include the diffuse background, smooth isotropic emission, and the *Fermi* bubbles templates, with normalizations as determined from the NPTF. The pixel-dependent survival function is defined as

$$\epsilon_p \equiv 1 - \text{CDF}[\mu_p, n_p], \quad (6)$$

where CDF is the Poisson cumulative distribution function. The smaller the value of ϵ_p (or, conversely, the larger the value of $-\log \epsilon_p$), the more probable it is that the pixel contains a PS. Figure 12 shows full-sky maps of $-\log \epsilon_p$ for both low (1.89–94.9 GeV) and high (50–2000 GeV) energies.¹⁸ The white circles indicate the presence of a 3FGL (2FHL) source for the low-(high-)energy map, with the radii proportional to the predicted photon counts for the sources. There is good correspondence between the hottest pixels, as determined by $-\log \epsilon_p$, and the brightest resolved sources. Pixels that are

¹⁸ Digital versions of these maps are available upon request.




Figure 24. Best-fit source-count distribution using the top three quartiles of Pass 8 source data and the `p8r2` foreground model. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

correspondingly less “hot” tend to be associated with less-bright 3FGL (or 2FHL) sources. Of particular interest are the hot pixels not already identified by the published catalogs. In the region $|b| \gtrsim 30^\circ$ ($|b| \gtrsim 10^\circ$) in the low-(high)-energy analysis, these are likely the sources lending the most weight to the NPTF below the catalog sensitivity thresholds. While more sophisticated algorithms are needed to further refine the candidate source locations, Figure 12 provides a starting point for identifying the spatial locations of potential new sources to help guide, for example, future TeV gamma-ray observations and cross-correlations with other data sets, such as the IceCube ultra-high-energy neutrinos.

Deciphering the constituents of the EGB remains an important goal in the study of high-energy gamma-ray astrophysics, with broad implications extending from the production of PeV neutrinos to signals of dark-matter annihilation or decay. The *Fermi*-LAT has already played an important role in the discovery of many new sources in the GeV sky. By taking advantage of the statistical properties of unresolved populations, our results provide a glimpse at the aggregate properties of the sources that lie below the detection threshold of these published catalogs and suggest a wealth of detections for future observatories.

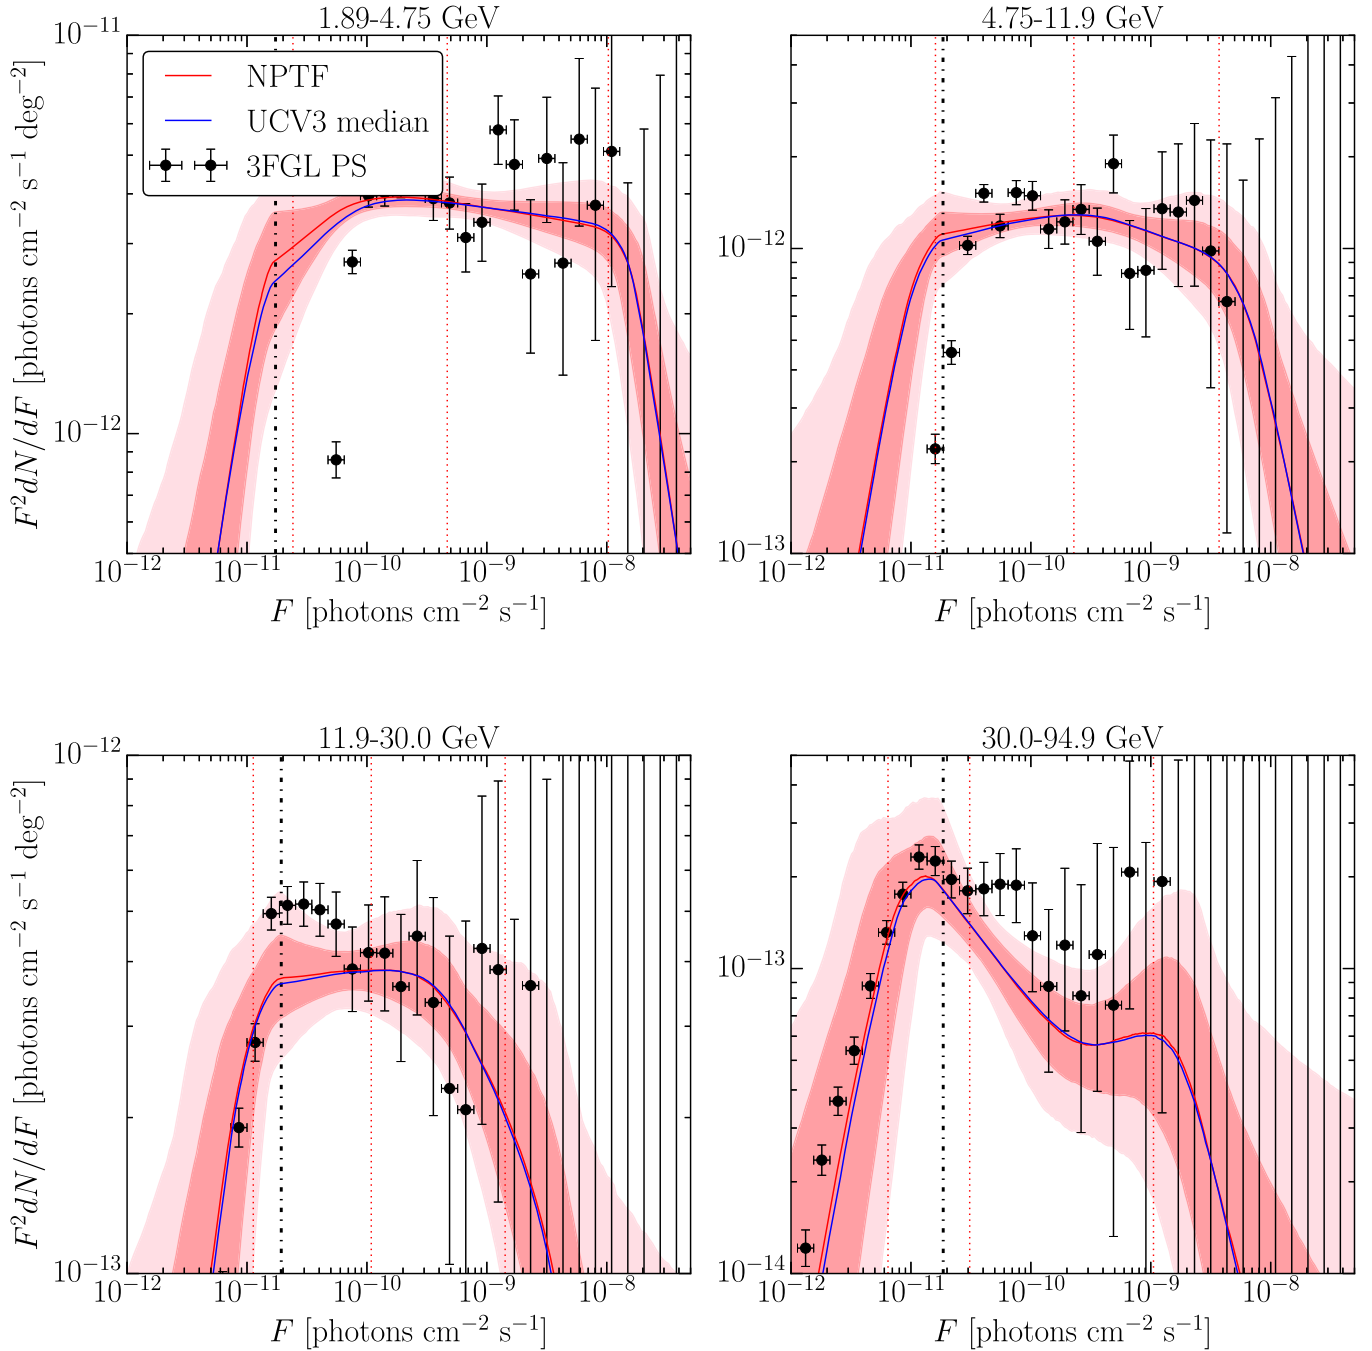


Figure 25. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and `p6v11` foreground model. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

We thank the *Fermi* Collaboration for the use of the *Fermi* public data and the Fermi Science Tools. We also thank J. Balkind, K. Bechtol, K. Blum, M. Di Mauro, N. Rodd, and T. Slatyer for helpful conversations. L.N. is supported by the U.S. Department of Energy under grant Contract Number DE-SC00012567. M.L. is supported by the DoE under grant Contract Number DE-SC0007968. B.R.S. is supported by a Pappalardo Fellowship in Physics at MIT.

Note added in the proof. As this work was being completed, we became aware of a complementary analysis that also takes advantage

of photon-count statistics to obtain the source-count distributions in the energy range from 1 to 171 GeV (Zechlin et al. 2016a). That study used the clean event class of the Pass 7 Reprocessed data to study the source-count distributions in five energy bins from 1 to 171 GeV. A direct comparison between the two results is challenging given the different data sets, foreground models, priors, energy sub-bins, pixelation, and regions of interest that were used between the two studies. In the energy range, where the two studies overlap, we tend to find smaller PS fractions. (For a discussion of why we do not consider energies below 1.89 GeV, see Appendix A.2.) Specifically, Zechlin et al. (2016a) finds PS fractions of $0.79^{+0.04}_{-0.16}$, $0.66^{+0.20}_{-0.07}$, $0.66^{+0.28}_{-0.05}$, and

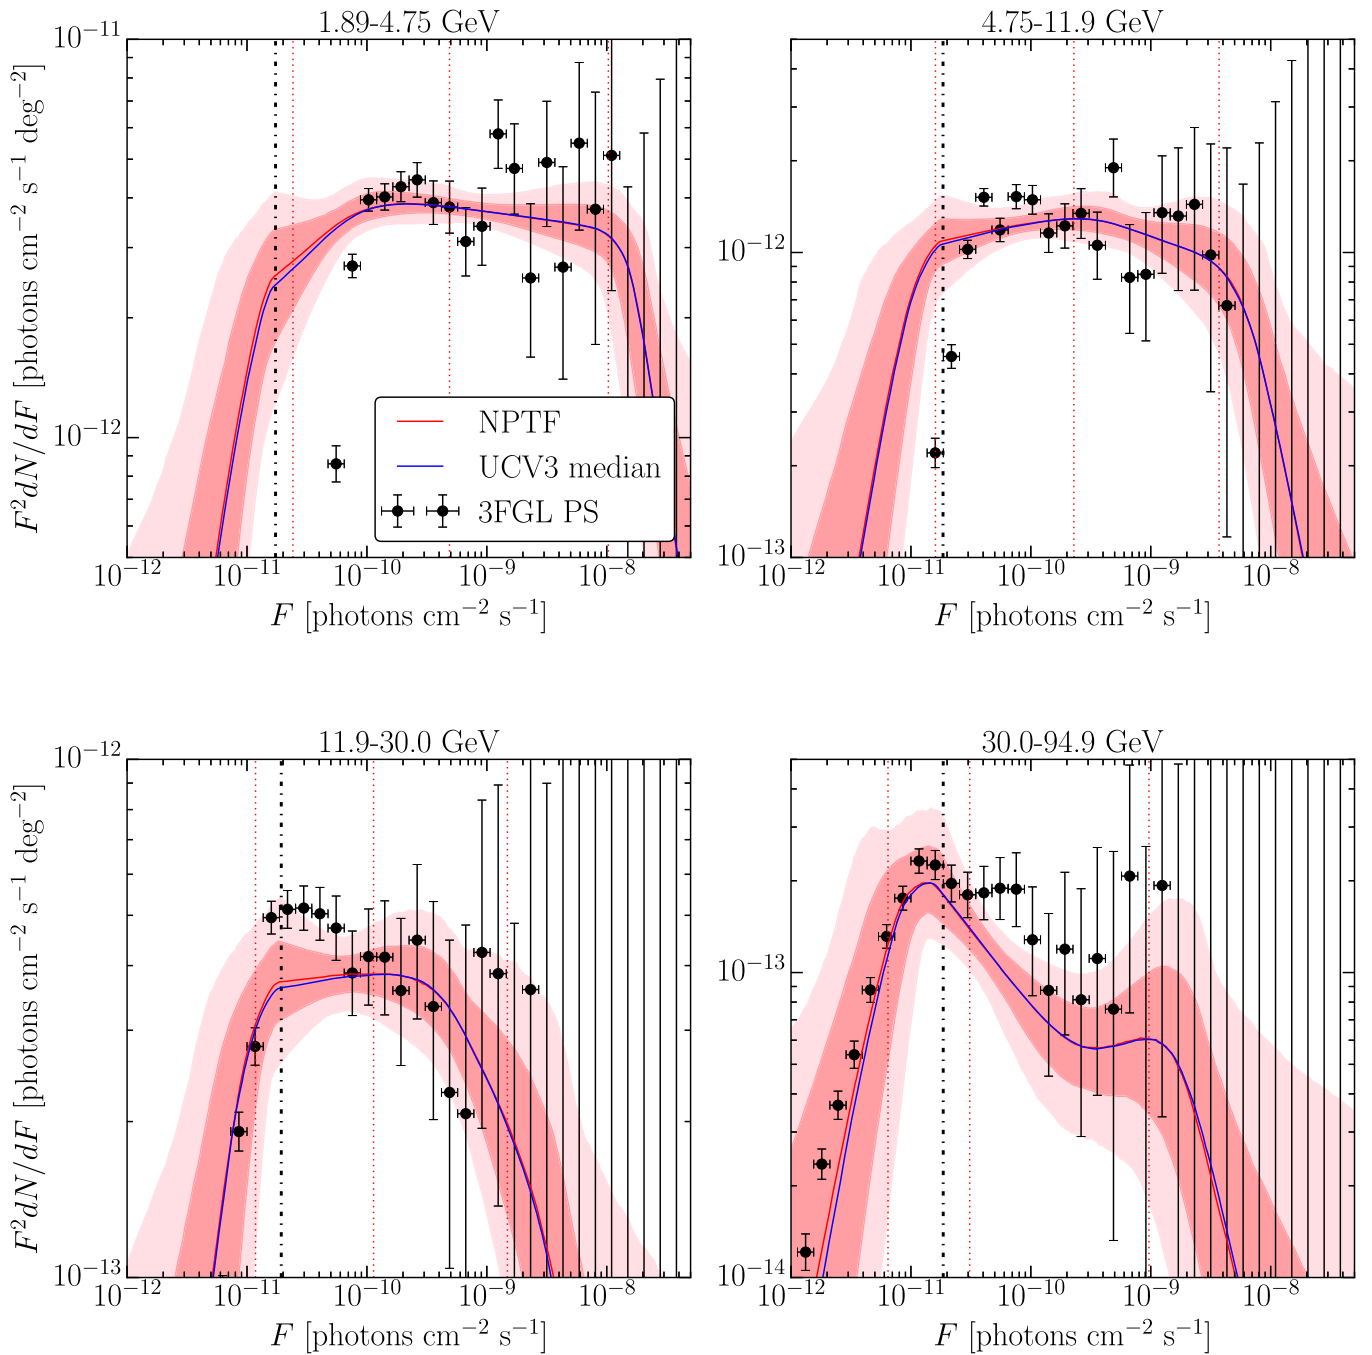


Figure 26. Best-fit source-count distribution using the Pass 8 *ultraleanveto* PSF3 data set and $p7v6$ foreground model. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

$0.81^{+0.52}_{-0.19}$ in the energy ranges [1.99, 5.0], [5.0, 10.4], [10.4, 50.0], and [50.0, 171] GeV. (PS fractions are quoted relative to the *Fermi* EGB intensity for foreground Model A; Ackermann et al. 2015b.) These numbers can be compared to our results, summarized in Table 4. The best-fit source-count distributions recovered by Zechlin et al. (2016a) closely resemble the Blazar-2 scenario that we studied in simulations (see Section 3). Our studies on simulated data show that the NPTF can successfully recover the energy spectrum and source-count distributions for this blazar population, both for the case where it singularly contributes to the EGB, as well as the case where it contributes along with SFGs, modeled according to Tamborra et al.

(2014). However, we observe different features in the best-fit source-count distributions when running the NPTF on the actual data.

APPENDIX A SIMULATIONS OF EXTRAGALACTIC GAMMA-RAY PSs

In this appendix, we provide further details about the simulations and analyses of extragalactic PSs.

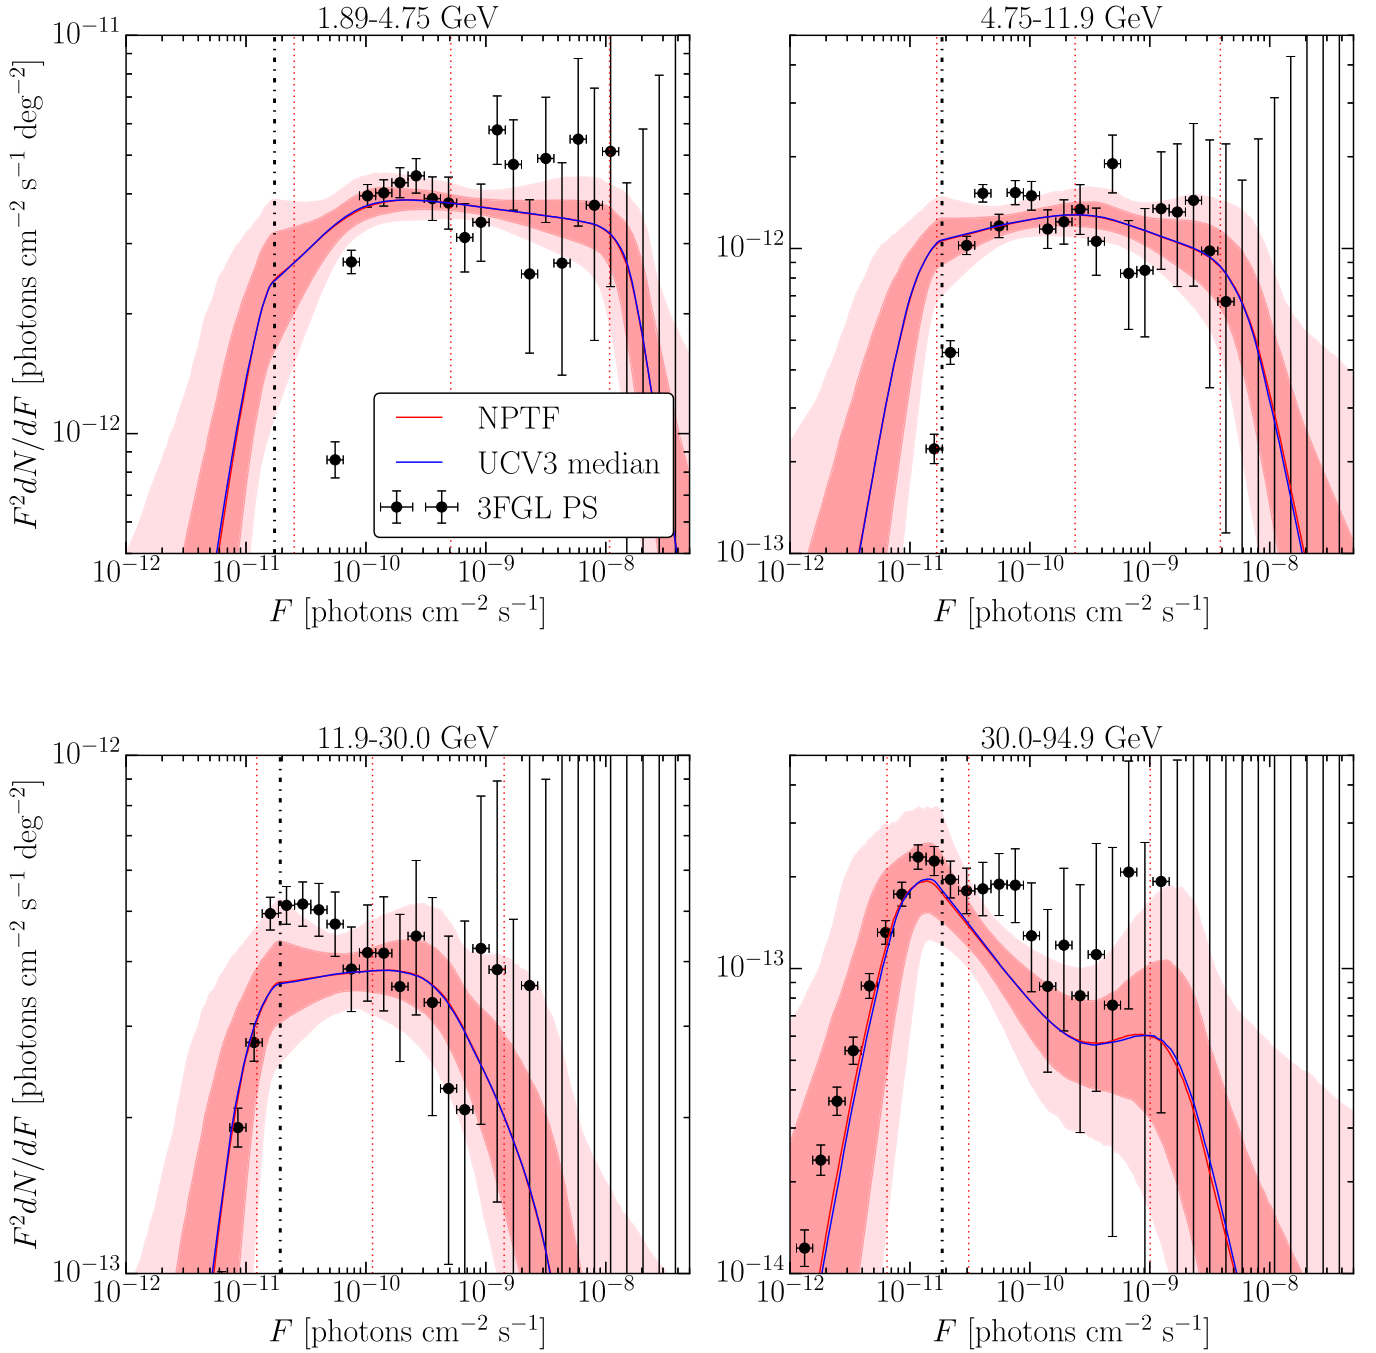


Figure 27. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and `p8r2` foreground model, but removing the *Fermi* bubbles template from the analysis. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

A.1. Energy-binned Source-count Distributions

We generate simulated maps directly from the source-count distribution dN/dF_γ . To obtain this, we need two inputs: the gamma-ray LF, $\Phi(L_\gamma, z, \Gamma)$, and the source energy spectrum, dF/dE (Di Mauro et al. 2014b). Typically, the LF is given by

$$\Phi(L_\gamma, z, \Gamma) = \frac{d^3N}{dL_\gamma dV d\Gamma}, \quad (7)$$

where V is the comoving volume, Γ is the photon spectral index, z is the redshift, N is the number of sources, and L_γ is the

rest-frame luminosity for energies from 0.1 to 100 GeV in units of GeV s^{-1} .

The photon flux in this energy range, F_γ , is defined in terms of the source energy spectrum,

$$F_\gamma(\Gamma) = \int_{E_{\min}}^{E_{\max}} \frac{dF}{dE} dE, \quad (8)$$

where the units are $\text{cm}^{-2} \text{s}^{-1}$, and $E_{\min(\max)} = 0.1(100)$ GeV.

The source-count distribution is then given by

$$\frac{dN}{dF_\gamma} = \frac{1}{4\pi} \int_{\Gamma_{\min}}^{\Gamma_{\max}} d\Gamma \int_{z_{\min}}^{z_{\max}} dz \Phi(L_\gamma, z, \Gamma) \frac{dV}{dz} \frac{dL_\gamma}{dF_\gamma}, \quad (9)$$




Figure 28. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and `p8r2` foreground model, but with a Gaussian PSF. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

which can be accurately estimated as

$$\frac{dN}{dF_\gamma} \approx \frac{1}{\Delta F_\gamma} \frac{1}{4\pi} \int_{\Gamma_{\min}}^{\Gamma_{\max}} d\Gamma \int_{z_{\min}}^{z_{\max}} dz \int_{L_\gamma(F_\gamma, \Gamma, z)}^{L_\gamma(F_\gamma + \Delta F_\gamma, \Gamma, z)} \times dL_\gamma \Phi(L_\gamma, z, \Gamma) \frac{dV}{dz}, \quad (10)$$

where 4π is the full-sky solid angle, dV/dz is the comoving volume slice for a given redshift, and ΔF_γ is sufficiently small. To calculate dN/dF_γ , we need the following expression, which

relates the luminosity to the energy flux:

$$L_\gamma(F_\gamma, \Gamma, z) = \frac{4\pi d_L^2}{(1+z)^{2-\Gamma}} \int_{E_{\min}}^{E_{\max}} E \frac{dF}{dE} dE, \quad (11)$$

where d_L is the luminosity distance. For a given F_γ and Γ , one can use (8) to solve for the normalization of dF/dE , which can be substituted into (11), along with z and Γ , to obtain the associated value of the luminosity. The photon flux, F_γ , is related to the photon count, S_γ , via the mean exposure $\langle \bar{E} \rangle$, which is averaged over 0.1–100 GeV and the ROI. This allows us to finally obtain dN/dS_γ from (10).

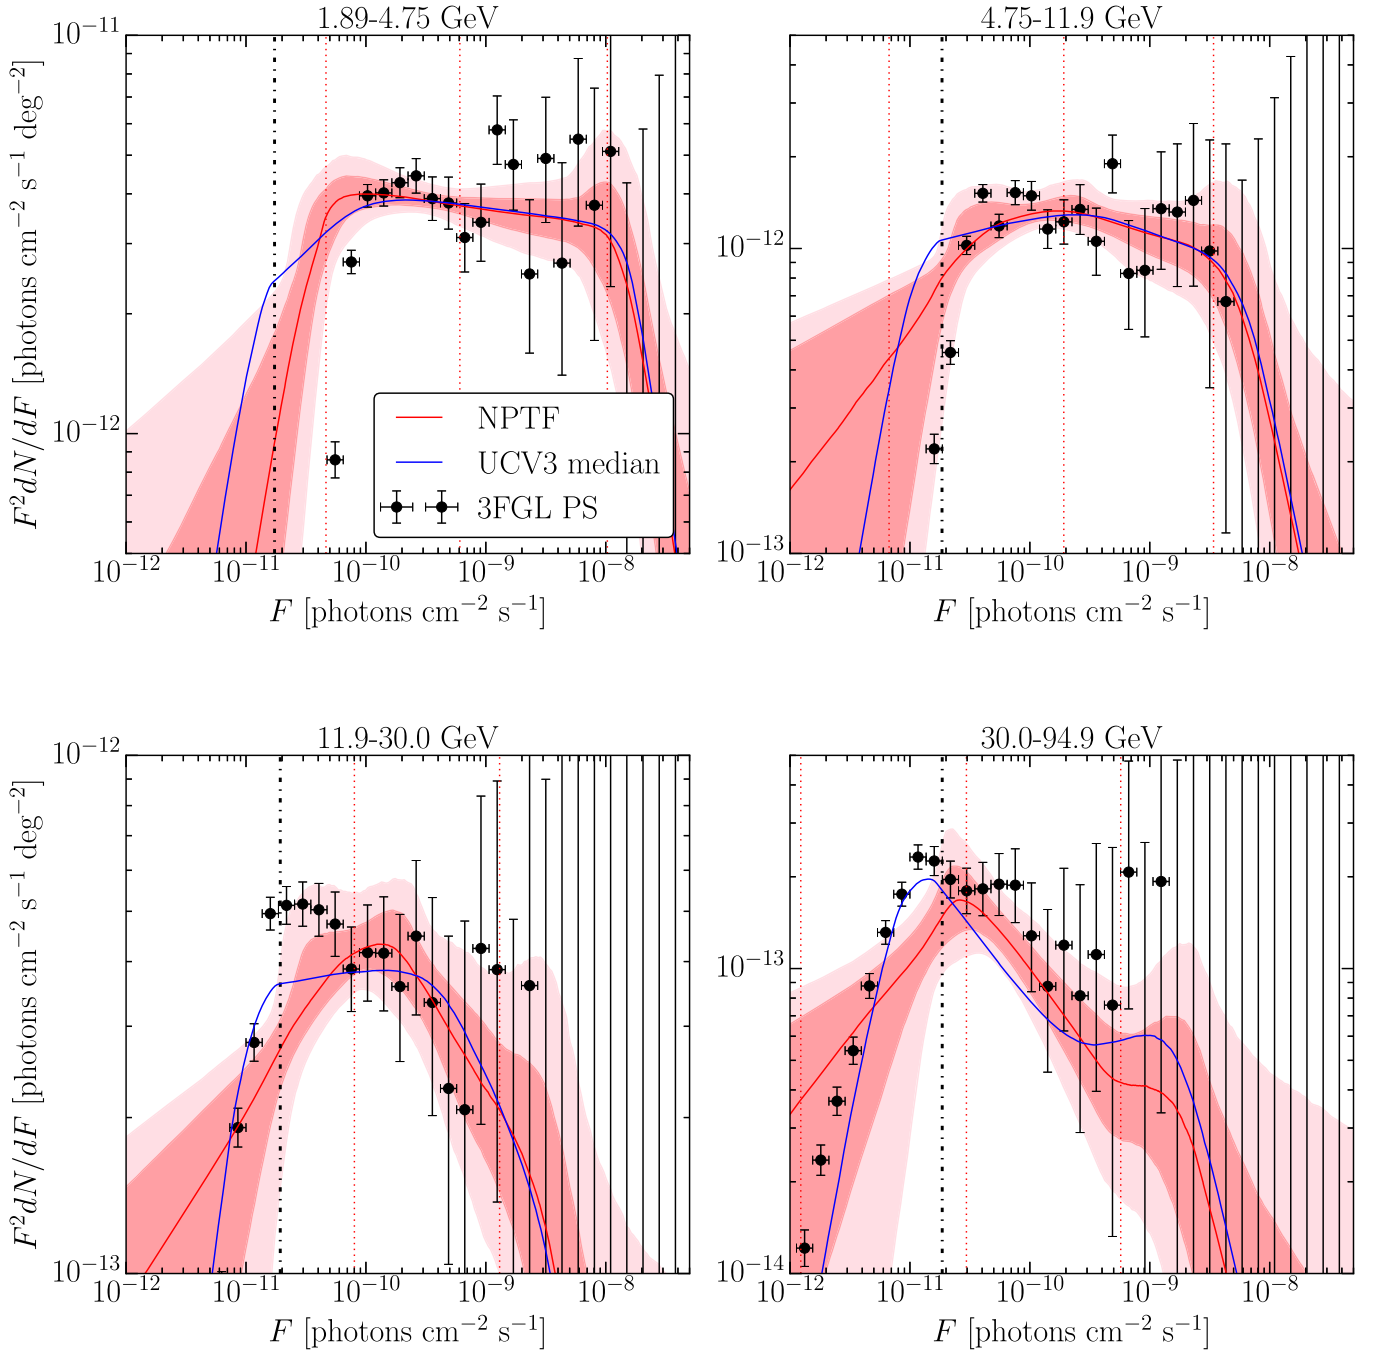


Figure 29. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and `p8r2` foreground model, but with the break priors set to $[0.1, 10]$, $[10, 40]$, $[40, 2 \times S_{b,\max}]$ ph, where $S_{b,\max}$ is the maximum number of photons in the 3FGL catalog in the energy bin of interest. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

The procedure outlined above allows one to obtain the source-count distributions based on models of LFs and spectral energy distributions provided in the literature. For the AGN and SFG examples we consider in detail in this work, the LFs correspond to photon energies from 0.1 to 100 GeV. However, we also need the source-count distributions in subset energy ranges corresponding to our energy bins of interest, with $E'_{\min,\max} \in [0.1, 100]$ GeV. We rescale the fluxes for these individual energy bins of interest to those in the provided 0.1–100 GeV range using a procedure similar to Di Mauro et al. (2014b). Denoting quantities associated with this

energy bin with a prime, we can write the new source-count distribution as

$$\begin{aligned} \frac{dN}{dF'_\gamma} &\approx \frac{1}{\Delta F'_\gamma} \frac{1}{4\pi} \int_{\Gamma_{\min}}^{\Gamma_{\max}} d\Gamma \int_{z_{\min}}^{z_{\max}} dz \int_{L_\gamma(F_\gamma(F'_\gamma, \Gamma), \Gamma, z)}^{L_\gamma(F_\gamma(F'_\gamma + \Delta F'_\gamma, \Gamma), \Gamma, z)} \\ &\quad \times dL_\gamma \Phi(L_\gamma, z, \Gamma) \frac{dV}{dz}, \end{aligned} \quad (12)$$

where $\Delta F'_\gamma$ is again sufficiently small—we set $\Delta F'_\gamma \equiv 10^{-3} F'_\gamma$, and verify that the answer is robust to this choice. Note that the

integral must still be done over L_γ (unprimed) because the LF is explicitly defined in terms of it. So, we must solve for the photon flux over the full energy, F_γ , in terms of the value in the sub-bin, F'_γ . The two are related via a proportionality relation

$$F_\gamma(F'_\gamma, \Gamma) = F'_\gamma \frac{\int_{E_{\min}}^{E_{\max}} dE \int_{L_{\gamma,\min}}^{L_{\gamma,\max}} dL_\gamma \int_{z_{\min}}^{z_{\max}} dz \Phi(L_\gamma, z, \Gamma) \frac{dV}{dz} \frac{dF}{dE} e^{-\tau_{\text{EBL}}(E,z)}}{\int_{E'_{\min}}^{E'_{\max}} dE \int_{L_{\gamma,\min}}^{L_{\gamma,\max}} dL_\gamma \int_{z_{\min}}^{z_{\max}} dz \Phi(L_\gamma, z, \Gamma) \frac{dV}{dz} \frac{dF}{dE} e^{-\tau_{\text{EBL}}(E,z)}}, \quad (13)$$

where the exponential factor accounts for the attenuation due to extragalactic background light (EBL; Gould & Schröder 1966; Fazio & Stecker 1970; Stecker et al. 1992; Franceschini et al. 2008; Ackermann et al. 2012c; Abramowski et al. 2013; Domínguez et al. 2013). It arises from pair annihilation of high-energy gamma-ray photons with other background photons in infrared, optical, and/or ultraviolet, and is described by the optical depth, τ_{EBL} . We use the EBL attenuation model from Finke et al. (2010).

Additionally, the expected gamma-ray spectrum can be calculated from the LF as

$$\frac{dN}{dE} = \frac{1}{4\pi} \int_{\Gamma_{\min}}^{\Gamma_{\max}} d\Gamma \int_{z_{\min}}^{z_{\max}} dz \int_{L_{\gamma,\min}}^{L_{\gamma,\max}} dL_\gamma \Phi(L_\gamma, z, \Gamma) \frac{dV}{dz} \frac{dF}{dE} e^{-\tau_{\text{EBL}}(E,z)}. \quad (14)$$

We use this equation to appropriately weight the number of photons per energy sub-bin for the individual sources when creating simulated maps. This ensures that the variations in PSF and exposure within the larger energy bins used in the NPTF analyses are properly accounted for in the simulation procedure.

A.2. Simulations at Energies $\lesssim 1.89$ GeV

The main analyses presented in the text do not consider energies below 1.89 GeV because we find that systematic effects may become more important at these low energies. In particular, simulations done with the set of priors presented in Section 2 show that the NPTF can overestimate the PS intensity at very low energies. That is, in simulations with both blazars and SFGs, the isotropic-PS template tends to absorb more emission than is simulated in blazars, while the smooth isotropic template absorbs less emission than is simulated in SFGs.

As an illustration, we show the results when the NPTF is run on a simulated map of Blazar-1 and SFG sources. Figure 13 shows the best-fit source-count distributions for the energy ranges 0.475–0.753 and 0.753–1.89 GeV. The PS fractions in these bins are $I_{\text{iso}}^{\text{PS}}/I_{\text{blazar-sim}} = [1.17^{+0.13}_{-0.09}, 1.26^{+0.10}_{-0.08}]$, illustrating that the PS template is absorbing more PS emission than is simulated in blazars. From Figure 13, we see that this is likely the result of the fit typically predicting more sources than it should at intermediate and low fluxes. This could be due to a variety of factors. For example, at very low energies, the angular resolution of the detector quickly degrades and the PS flux also becomes more and more subdominant compared to foreground emission. For these reasons—and out of an

abundance of caution—we do not present results of the NPTF on data below 1.89 GeV. It is certainly possible that analyses at low energies could provide a wealth of interesting observations. However, having confidence in the NPTF results at such

low-energies requires a more careful consideration of systematics, which goes beyond the scope of the present work.

A.3. Ultracleanveto PSF1–3 Simulation Analysis

The simulated-data studies in the main text used the PSF3 instrument response function for the Pass 8 *ultracleanveto* data set. Here, we show what happens when using the PSF1–3 instrument response function instead. Including the top three quartiles of data increases the total exposure, though at the cost of decreased angular resolution. Figure 14 illustrates the result for the Blazar-2 simulated data set. The best-fit source-count distribution extends to lower fluxes than the corresponding distribution for top-quartile data in Figure 3.

A.4. Simulations of SFGs

Here, we show the results of running the NPTF analysis on simulated data where the EGB arises solely from SFGs. The resulting best-fit source-count distributions are shown in Figure 15. In the first energy bin, the brightest SFGs, which contribute little more than ~ 1 photon, are detected as PSs by the NPTF. At higher energies, the best-fit source-count distributions are consistent with zero and no significant evidence for a PS population is found. The energy spectrum plot in Figure 16 shows that the SFG flux is absorbed by the smooth isotropic template. In comparison, the intensity absorbed by the isotropic-PS template is completely subdominant and is several orders of magnitude lower than its Poissonian counterpart at all energies.

APPENDIX B SUPPLEMENTARY RESULTS: LOW-ENERGY ANALYSIS

B.1. Best-fit Intensities and Posterior Distributions

This section includes supplementary information pertaining to the low-energy analysis presented in Section 4.1. In particular, Tables 5 and 6 present the best-fit intensities and source-count parameters for the NPTF analysis of the top-three quartiles of *ultracleanveto* data. Figures 17–20 show the posterior distributions, in each energy bin, for the benchmark analysis.

B.2. Systematic Tests

We now describe in detail the systematic tests that were conducted for the low-energy analysis, the results of which are summarized in Figure 7. The primary conclusion that we draw is that the PS fraction is stable under the variety of tests that we have explored.

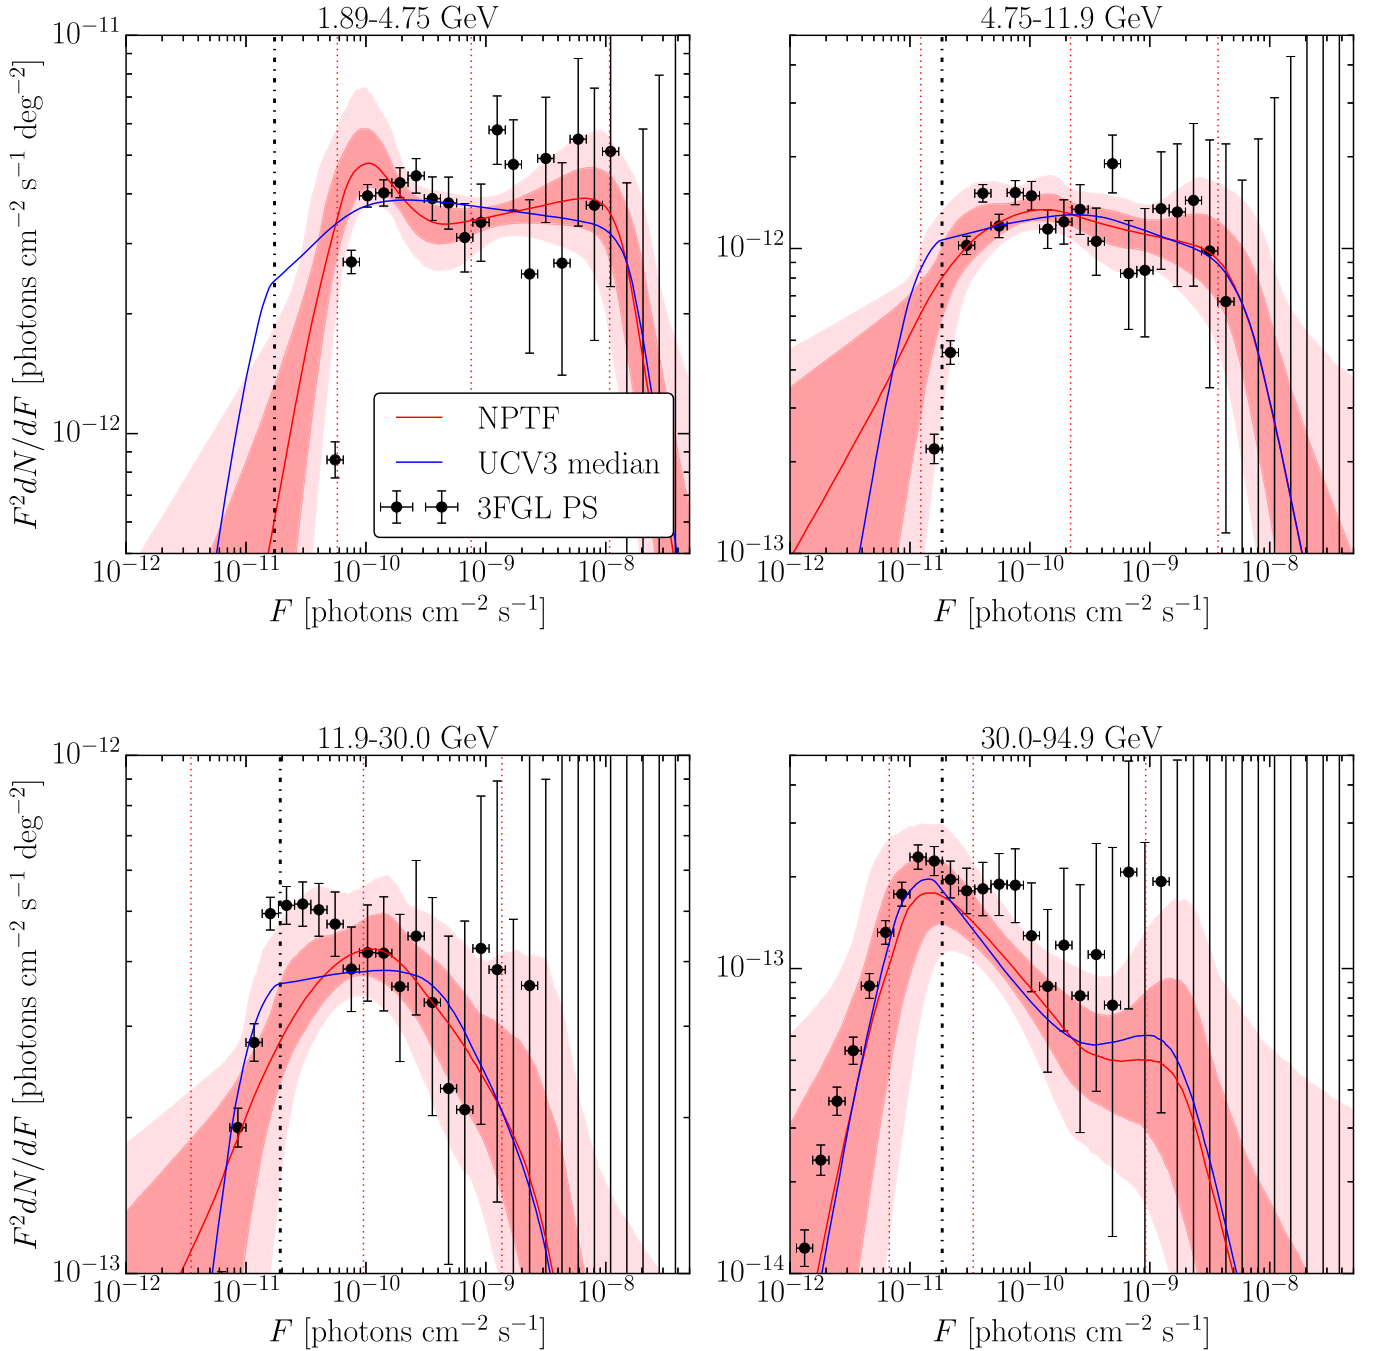


Figure 30. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and `p8r2` foreground model, but with the break priors set to $[1, 20]$, $[20, S_{b,\max}/2]$, $[S_{b,\max}/2, 2 \times S_{b,\max}]$ ph, where $S_{b,\max}$ is the maximum number of photons in the 3FGL catalog in the energy bin of interest. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

B.2.1. Region of Interest

As a first cross-check on the stability of the results presented in Section 4.1, we explore the effects of altering the ROI. While we previously defined the ROI with $|b| \geq 30^\circ$, we now loosen this constraint and consider the case of $|b| \geq 10^\circ$. Extending the ROI closer to the Galactic disk increases the amount of data being analyzed, but at the cost of potentially more contamination from diffuse foreground emission and local PSs. As shown in Figure 7, the best-fit intensities for the isotropic and isotropic-PS components are equivalent, within errors, to their counterparts in the benchmark analysis. The best-fit source counts are shown in Figure 21.

We also ran the NPTF on the northern ($b > 30^\circ$) and southern ($b < -30^\circ$) hemispheres separately. The intensities for the EGB, IGRB, and PS components are systematically lower (higher) for the northern (southern) analysis than for the benchmark case. Similar behavior is apparent in the source-count plots, shown in Figures 22 and 23.

B.2.2. Event Class

We explored the implications of broadening the *ultracleanveto* data set to include the top three quartiles in Section 4.1.2. Now, we consider the implications of repeating the NPTF analysis on the source data with PSF1–3. This event class has

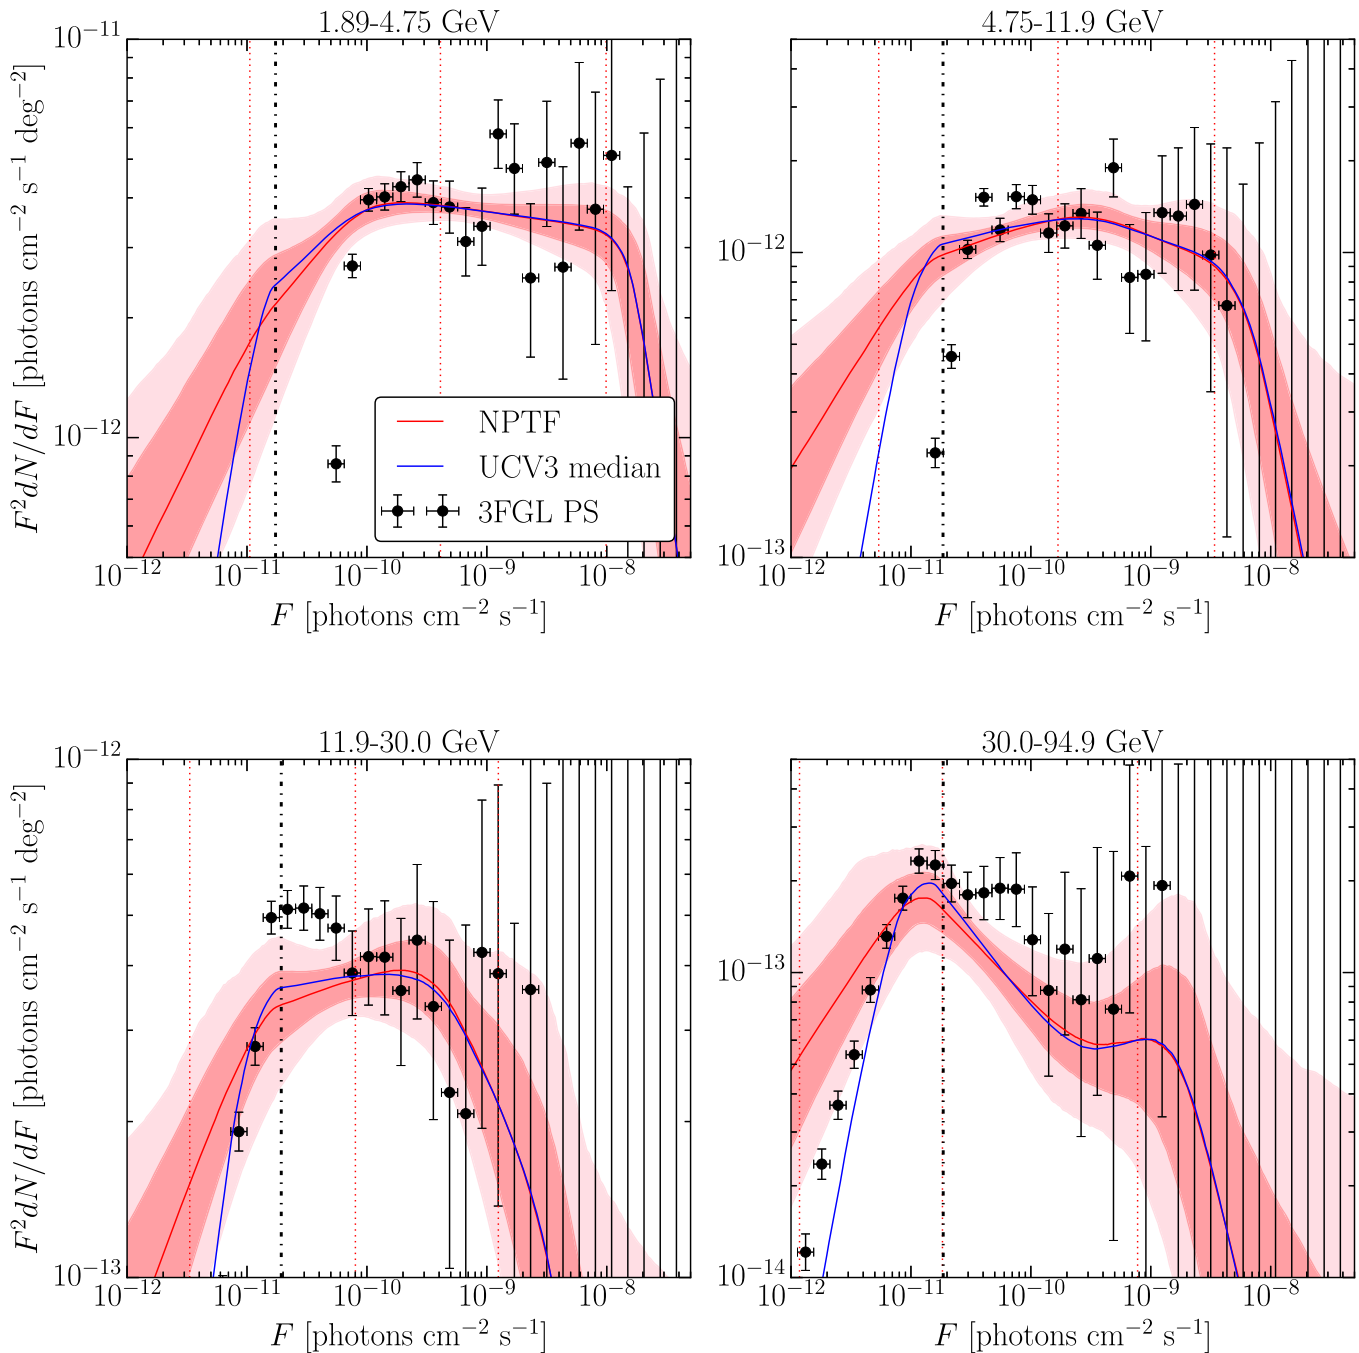


Figure 31. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF3 data set and `p8r2` foreground model, but with the prior for the lowest slope restricted to $n_4 \in [1, 2]$. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 5.)

looser photon-quality cuts, which leads to larger overall exposure, but significantly more cosmic-ray contamination. In general, it is not recommended to use source data for IGRB studies; for our purposes, however, it will be intriguing to see how the increased photon statistics affect the recovered source-count distribution for the PS component. As shown in Figure 7, the EGB intensity is far larger than that recovered by the benchmark analysis and overpredicts *Fermi*'s EGB result in most energy bins. The sharp rise in the EGB intensity can be traced to a substantial fraction of smooth isotropic emission, which is expected for this event class at most energies. Most importantly, the intensity of the isotropic-PS component is consistent, within uncertainties, with that found in the

benchmark analysis.¹⁹ This is a confirmation that the NPTF is able to successfully constrain the source-count distribution even in a data set with significantly more smooth isotropic flux.

The source-count distributions are provided in Figure 24. In general, they exhibit similar behavior to the *ultracleanveto* PSF1–3 functions, extending to lower fluxes due to the increased exposure for this event class. One potential new feature of interest in the source-data source-count distributions is that, in the second energy bin from 4.75 to 11.9 GeV, there is

¹⁹ The recovered PS intensity is slightly larger with source PSF1–3 data as compared to *ultracleanveto* PSF3 data, which is likely due to the increased exposure in the source PSF1–3 data set.

Table 7
Same as Table 2, Except Using All Quartiles (PSF0–3) of the Pass 8 *ultracleanveto* Data for the High-energy Analysis

Energy (GeV)	I_{EGB}	$I_{\text{iso}}^{\text{PS}}$	$I_{\text{iso}} (\text{cm}^{-2} \text{s}^{-1} \text{sr}^{-1})$	I_{diff}	I_{bub}
50.0–151	$3.10^{+0.13}_{-0.11} \times 10^{-9}$	$1.36^{+0.19}_{-0.16} \times 10^{-9}$	$1.74^{+0.13}_{-0.16} \times 10^{-9}$	$2.69^{+0.06}_{-0.06} \times 10^{-9}$	$5.26^{+2.60}_{-2.51} \times 10^{-10}$
151–457	$4.38^{+0.42}_{-0.32} \times 10^{-10}$	$1.56^{+0.42}_{-0.29} \times 10^{-10}$	$2.80^{+0.27}_{-0.31} \times 10^{-10}$	$4.12^{+0.21}_{-0.23} \times 10^{-10}$	$5.40^{+7.15}_{-3.81} \times 10^{-11}$
457–2000	$1.10^{+0.15}_{-0.13} \times 10^{-10}$	$1.29^{+0.99}_{-0.61} \times 10^{-11}$	$9.61^{+1.32}_{-1.40} \times 10^{-11}$	$6.29^{+1.22}_{-1.37} \times 10^{-11}$	$7.18^{+5.02}_{-4.11} \times 10^{-11}$
50.0–2000	$3.74^{+0.16}_{-0.12} \times 10^{-9}$	$1.61^{+0.22}_{-0.18} \times 10^{-9}$	$2.13^{+0.15}_{-0.19} \times 10^{-9}$	$3.29^{+0.07}_{-0.07} \times 10^{-9}$	$5.26^{+3.01}_{-2.58} \times 10^{-10}$

Note. Note that the *Fermi* bubbles template intensity is defined relative to the interior of the bubbles, while the intensities of the other templates are computed with respect to the region $|b| > 10^\circ$.

Table 8
Same as Table 3, Except Using All Quartiles (PSF0–3) of the Pass 8 *ultracleanveto* Data for the High-energy Analysis

Energy (GeV)	n_1	n_2	n_3	n_4	$F_{b,3}$	$F_{b,2} (\text{cm}^{-2} \text{s}^{-1})$	$F_{b,1}$
50.0–151	$3.56^{+0.92}_{-0.90}$	$2.34^{+0.36}_{-0.27}$	$2.18^{+0.12}_{-0.11}$	$-0.05^{+1.12}_{-1.21}$	$1.58^{+1.21}_{-0.76} \times 10^{-12}$	$7.93^{+2.78}_{-4.10} \times 10^{-11}$	$1.55^{+0.77}_{-0.73} \times 10^{-9}$
151–457	$3.56^{+0.95}_{-0.97}$	$1.87^{+0.65}_{-0.53}$	$2.42^{+0.39}_{-0.28}$	$-0.06^{+1.08}_{-1.26}$	$2.53^{+1.16}_{-1.15} \times 10^{-12}$	$6.41^{+4.34}_{-3.86} \times 10^{-11}$	$4.80^{+2.63}_{-2.29} \times 10^{-10}$
457–2000	$3.57^{+0.91}_{-0.96}$	$2.26^{+0.78}_{-0.78}$	$2.16^{+0.68}_{-0.67}$	$-0.01^{+1.25}_{-1.25}$	$3.16^{+1.74}_{-1.68} \times 10^{-12}$	$8.90^{+3.94}_{-3.73} \times 10^{-11}$	$2.35^{+0.38}_{-0.37} \times 10^{-10}$
50.0–2000	$3.63^{+0.90}_{-0.99}$	$2.28^{+0.28}_{-0.22}$	$2.17^{+0.12}_{-0.09}$	$-0.05^{+1.10}_{-1.24}$	$1.72^{+1.29}_{-0.79} \times 10^{-12}$	$7.87^{+3.16}_{-4.37} \times 10^{-11}$	$2.15^{+1.18}_{-1.06} \times 10^{-9}$

a more pronounced hardening of the source-count distribution below the second break $F_{b,2}$, as compared to the *ultracleanveto* PSF1–3 analyses.

B.2.3. Foreground Model

A potentially significant source of systematic uncertainty in the NPTF analysis is due to mis-modeling of high-energy gamma-rays produced in cosmic-ray propagation in the Milky Way (Ackermann et al. 2012a). These high-energy photons arise from bremsstrahlung of electrons off the interstellar medium, boosted pion decay, and inverse Compton (IC) emission off the interstellar radiation field. Our benchmark analysis uses the associated foreground model for the Pass 8 data set (*gll_iem_v06.fits*), denoted here as p8r2. The total diffuse emission in p8r2 is modeled as a linear combination of several sources, some of which are traced by maps of gas column densities, which serve as templates for the pion and bremsstrahlung emission. The IC component is modeled using the GALPROP package (Strong et al. 2007).²⁰ These individual templates are fit to the data, and used to identify “extended emission excesses” that are identified directly and then added back into the model (Acero et al. 2016).

To better assess the uncertainties due to the foreground modeling, we repeat the NPTF analysis using several other foreground models made available by *Fermi*. In particular, we use the *gll_iem_v02_P6_V11_DIFFUSE.fits* diffuse emission model, denoted as p6v11, which was initially developed for the Pass 6 data set.²¹ p6v11 is distinct from p8r2 in that it uses older gas and IC maps and does not include templates for large-scale structure or extended emission excesses. The Pass 7 model *gal_2yearp7v6_v0.fits*, denoted as p7v6,²² is a compromise because it uses updated gas and IC maps and

includes some large-scale extended structures, such as Loop 1 and the *Fermi* bubbles.

The NPTF results using the p6v11 and p7v6 foreground models are summarized in Figure 7, with source-count distributions shown in Figures 25 and 26, respectively. In general, we observe that the intensity of the PS components is consistent with that for the benchmark analysis in all energy bins. However, variations occur in the smooth isotropic intensity. Typically, more IGRB intensity is recovered with p6v11 and p7v6, versus p8r2. The differences are particularly dramatic in the first two energy bins and are more severe for p6v11. The net consequence is that the EGB intensity is higher than the expected range from *Fermi*. The enhancement in the isotropic component may arise from the fact that each foreground model incorporates large-scale diffuse structures differently—with p6v11 being the least inclusive and p8r2 being the most inclusive. We note, however, that the fit to data with the p8r2 foreground model, from the point of view of the Bayesian evidence, is much better than the analogous fit with the p6v11 model; the fit with the p7v6 model is intermediate.

B.2.4. The Bubbles Template

To better understand how dependent the analysis is on the details of the *Fermi* bubbles template, we simply removed the template from the analysis. This has indiscernible effects on the final results. We see in Figure 7 that the EGB, IGRB, and PS intensities are consistent, within uncertainties, to the corresponding values in the benchmark study. The source-count distributions, shown in Figure 27, are also degenerate with those found including the *Fermi* bubbles template.

B.2.5. Point-spread Function

The PSF can affect the photon-count distribution because it can redistribute photons between pixels, and must therefore be properly accounted for in the calculation of the photon-count probability distributions. For the primary analyses presented in this work, the PSF is modeled using a King function. However,

²⁰ <http://galprop.stanford.edu/>

²¹ http://fermi.gsfc.nasa.gov/ssc/data/access/lat/ring_for_FSSC_final4.pdf

²² http://fermi.gsfc.nasa.gov/ssc/data/access/lat/Model_details/Pass7_galactic.html

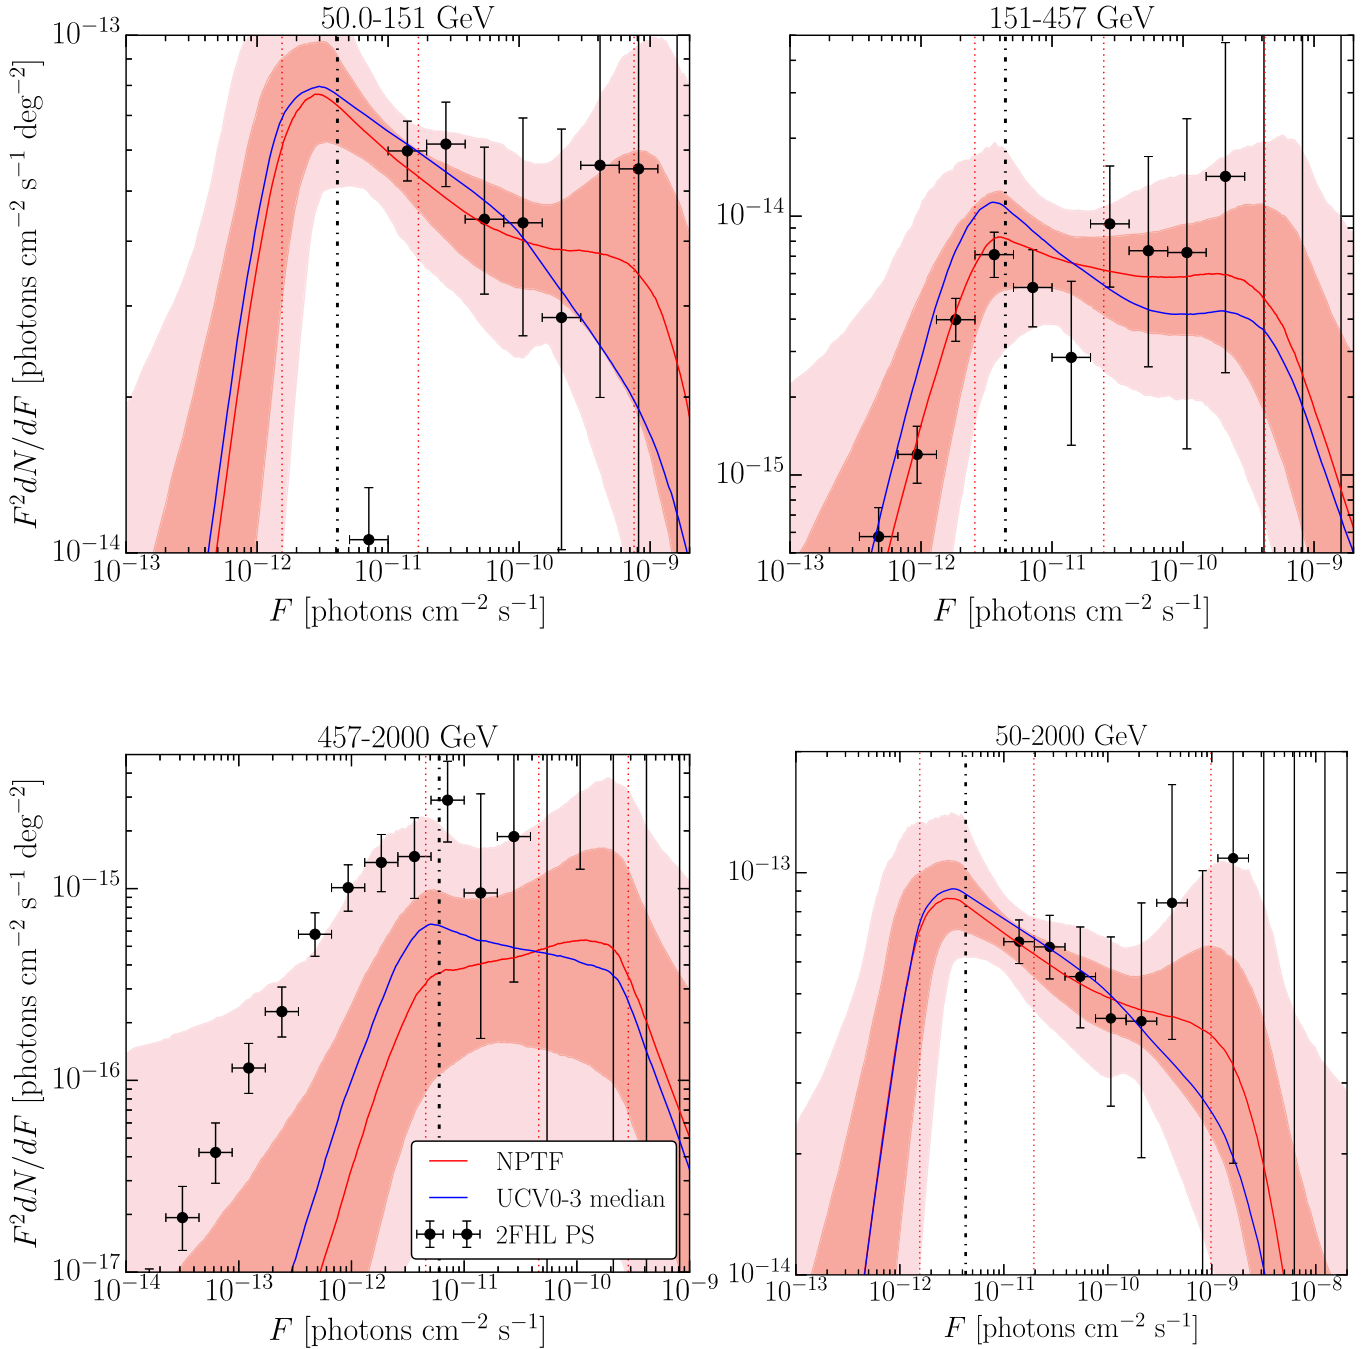


Figure 32. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF0-3 data set and $p8r2$ foreground model, but with $|b| > 30^\circ$. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 8.)

to test the sensitivity of the results to mis-modeling of the PSF, we have also repeated the NPTF analysis using a two-dimensional Gaussian in the calculation of the photon-count probability distributions, with a width set to give the correct 68% containment radius. As shown in Figure 7, the NPTF results remain unchanged with this substitution. The best-fit source-count distribution for 1.89–4.75 GeV shows some variation at the lowest fluxes, but within uncertainties (see Figure 28).

B.2.6. Priors

Our choice of priors, given in Table 1, is carefully chosen to both avoid biasing the posterior for the source-count

distribution while at the same time allowing breaks at both high and low flux. This is meant to properly account for the fact that the source-count distribution is not well constrained by the data at very high fluxes, where the mean expected number of sources over the full region is much less than unity, and at very low fluxes, where the mean photon count per source is much less than unity. Our choice of priors is further justified by the simulated-data studies, presented in Section 3, which show that the NPTF can successfully constrain the emission from blazar models. However, one may still be concerned that these particular choices of priors might bias the recovered source-count distribution in a particular way. For that reason, we have tried many variations to the priors shown in Table 1, three of




Figure 33. Best-fit source-count distribution using the Pass 8 *ultracleanveto* PSF0–3 data set and `p8r2` foreground model, but with the prior on the lowest slope restricted to $n_4 \in [1, 2]$. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 8.)

which (labeled “alternate priors 1–3”) are described below and shown in Figure 7.

1. *Alternate prior 1*: all priors are the same as in Table 1, except for those on the breaks, which are changed to $[0.1, 10]$, $[10, 40]$, and $[40, 2 \times S_{b,\max}]$ ph for $S_{b,1}$, $S_{b,2}$, and $S_{b,3}$, respectively.
2. *Alternate prior 2*: as above, except changing the priors for the breaks to $[1, 20]$, $[20, S_{b,\max}/2]$, and $[S_{b,\max}/2, 2 \times S_{b,\max}]$ ph, respectively.
3. *Alternate prior 3*: all priors are the same as in Table 1, except for that of n_4 , which is changed to $[1, 1.99]$.

The first two examples address the possibility that the break priors might artificially sculpt the source-count distribution and the recovered PS intensity, while the third example addresses how the source-count distribution is dealt with at fluxes below the lowest break, where the distribution is not well constrained by the data. In many classes of blazar models, such as those considered in Section 3, the index below the lowest break (n_4) is greater than unity, so that the total number of PSs $\sim \int_{F_{\min}} dF dN/dF$ diverges as the minimum flux cut-off F_{\min} is taken to zero.

It is useful to know if the recovered PS intensity, $I_{\text{iso}}^{\text{PS}}$, tends to under or overshoot the simulated blazar intensity, $I_{\text{blazar-sim}}$,

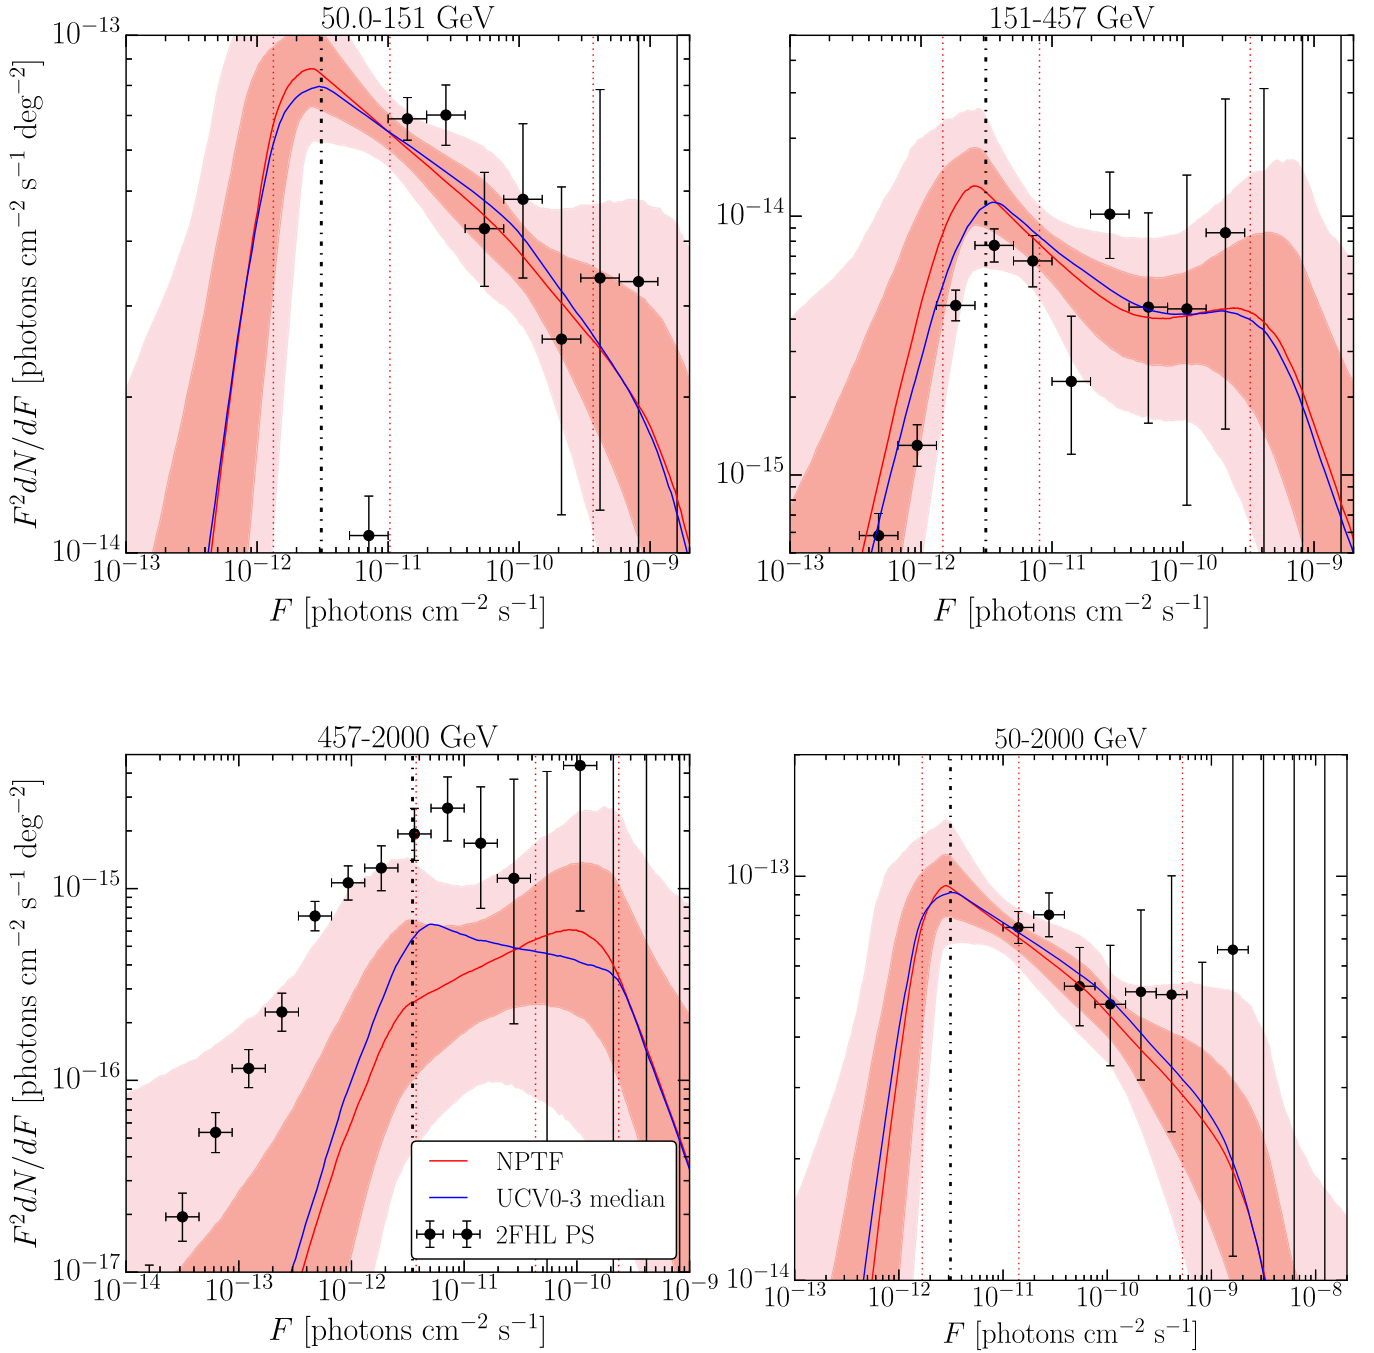


Figure 34. Best-fit source-count distribution using the Pass 8 *source* PSF0-3 data set and $p8\epsilon2$ foreground model, with $|b| > 10^\circ$. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 8.)

when using the alternate priors. With that in mind, we run the NPTF on simulated maps, as in Section 3, constructed from both the SFG + Blazar-1 model as well as the SFG + Blazar-2 model. For *Alternate prior 1*, we find that

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.87^{+0.05}_{-0.04}, 0.93^{+0.17}_{-0.08}, 0.92^{+0.23}_{-0.15}, 0.61^{+0.11}_{-0.07}]$$

and

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.68^{+0.06}_{-0.05}, 0.59^{+0.15}_{-0.09}, 0.52^{+0.07}_{-0.05}, 0.37^{+0.05}_{-0.03}]$$

for the SFG + Blazar-1 and SFG + Blazar-2 models, respectively, with the *ultracleanveto* PSF3 instrument response function. With *Alternate prior 1*, we see larger uncertainties, with the PS template capable of absorbing more flux in particular. With *Alternate prior 2*, on the other hand, we find more noticeable differences in the medians as well as in the uncertainties. In particular, for the SFG + Blazar-1 and SFG + Blazar-2 models, we find

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [1.01^{+0.12}_{-0.10}, 1.27^{+0.16}_{-0.31}, 1.25^{+0.12}_{-0.15}, 0.73^{+0.21}_{-0.12}]$$




Figure 35. Best-fit source-count distribution using the Pass 8 *source* PSF0-3 data set and `p8r2` foreground model, but with $|b| > 30^\circ$. The median source-count distribution for the benchmark analysis is shown in blue. (Formatted as in Figure 8.)

and

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.74^{+0.19}_{-0.06}, 0.94^{+0.20}_{-0.19}, 0.61^{+0.17}_{-0.10}, 0.41^{+0.09}_{-0.05}],$$

respectively. In the Blazar-1 model case, it is important to notice that at intermediate energies the NPTF tends to over-predict $I_{\text{blazar-sim}}$ at the $\sim 20\%$ level. With *Alternate prior 3*, the results are

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [1.06^{+0.15}_{-0.09}, 1.10^{+0.14}_{-0.09}, 1.00^{+0.14}_{-0.10}, 0.85^{+0.15}_{-0.11}]$$

and

$$\frac{I_{\text{iso}}^{\text{PS}}}{I_{\text{blazar-sim}}} = [0.92^{+0.16}_{-0.09}, 0.77^{+0.39}_{-0.14}, 0.69^{+0.12}_{-0.08}, 0.53^{+0.10}_{-0.06}],$$

for the Blazar-1 and Blazar-2 models. The *Alternate prior 3* results are consistently closer to unity than the first two alternate prior results.

As may be seen in Figure 7, the median values for the PS intensities recovered from the NPTF analyses with alternate priors are generally consistent with those found in the baseline study. The *Alternate prior 3* PS intensities are slightly enhanced in all energy bins compared to the baseline results

—following our expectations from the simulation results presented above—though the two results are consistent within statistical uncertainties. The recovered source-count distributions, shown in Figure 31, illustrate that the *Alternate prior 3* results are consistent with our baseline results at fluxes above the ~ 1 photon threshold. At lower fluxes, the source-count distributions are slightly softer than in our baseline analysis, as is almost guaranteed by the fact that $n_4 > 1$ with *Alternate prior 3* while $n_4 > -2$ in our baseline analysis.

The *Alternate prior 1* and *Alternate prior 2* results have mean PS intensities similar to those in the baseline analysis, though in the second and third energy bin the upper limits of the credible intervals extend to higher values. This is a reflection of the fact that the recovered source-count distributions, shown in Figures 29 and 30, are softer at low fluxes compared to those in the baseline analysis. This is perhaps due to the fact that the lowest break tends to be at higher flux, and thus the index n_4 is influenced by the data in the vicinity of the break.

APPENDIX C SUPPLEMENTARY RESULTS: HIGH-ENERGY ANALYSIS

This section includes supplementary information pertaining to the high-energy analysis presented in Section 5. In particular, Tables 7 and 8 present the best-fit intensities and source-count parameters for the NPTF analysis of all quartiles of *ultra-clean veto* data. Figures 32–35 show the best-fit source-count distributions for the various systematic studies described in the text.

REFERENCES

- Aartsen, M. G., Abbasi, R., Abdou, Y., et al. 2013a, *Sci*, **342**, 1242856
Aartsen, M. G., Abbasi, R., Abdou, Y., et al. 2013b, *PhRvL*, **111**, 021103
Aartsen, M. G., Abraham, K., Ackermann, M., et al. 2015a, *ApJ*, **809**, 98
Aartsen, M. G., Abraham, K., Ackermann, M., et al. 2015b, *PhRvL*, **115**, 081102
Abazajian, K. N., Blanchet, S., & Harding, J. P. 2011, *PhRvD*, **84**, 103007
Abdo, A. A. 2010, *ApJ*, **715**, 429
Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJ*, **720**, 435
Abramowski, A., Acero, F., Aharonian, F., et al. 2013, *A&A*, **550**, A4
Acero, F., Ackermann, M., Ajello, M., et al. 2015, *ApJS*, **218**, 23
Acero, F., Ackermann, M., Ajello, M., et al. 2016, *ApJS*, **223**, 26
Ackermann, M., Ajello, M., Albert, A., et al. 2015a, *JCAP*, **09**, 008
Ackermann, M., Ajello, M., Albert, A., et al. 2015b, *ApJ*, **799**, 86
Ackermann, M., Ajello, M., Albert, A., et al. 2016b, *PhRvL*, **116**, 151105
Ackermann, M., Ajello, M., Allafort, A., et al. 2012b, *ApJ*, **755**, 164
Ackermann, M., Ajello, M., Allafort, A., et al. 2012c, *Sci*, **338**, 1190
Ackermann, M., Ajello, M., Atwood, W. B., et al. 2012a, *ApJ*, **750**, 3
Ackermann, M., Ajello, M., Atwood, W. B., et al. 2016a, *ApJS*, **222**, 5
Ahlers, M., & Salvado, J. 2011, *PhRvD*, **84**, 085019
Ajello, M., Costamante, L., Sambruna, R. M., et al. 2009, *ApJ*, **699**, 603
Ajello, M., Gasparrini, D., Sánchez-Conde, M., et al. 2015, *ApJL*, **800**, L27
Ajello, M., Romani, R. W., Gasparrini, D., et al. 2014, *ApJ*, **780**, 73
Ajello, M., Shaw, M. S., Romani, R. W., et al. 2012, *ApJ*, **751**, 108
Ando, S., & Pavlidou, V. 2009, *MNRAS*, **400**, 2122
Ando, S., Tamborra, I., & Zandanel, F. 2015, *PhRvL*, **115**, 221101
Angel, J., & Stockman, H. 1980, *ARA&A*, **18**, 321
Bechtol, K., Ahlers, M., Di Mauro, M., Ajello, M., & Vandenbroucke, J. 2015, *ApJ*, submitted (arXiv:1511.00688)
Bengtsson, H. U., Salati, P., & Silk, J. 1990, *NuPhB*, **346**, 129
Bergstrom, L., Edsjo, J., & Ullio, P. 2001, *PhRvL*, **87**, 251301
Bertone, G., Hooper, D., & Silk, J. 2005, *PhR*, **405**, 279
Bhattacharya, D., Sreekumar, P., & Mukherjee, R. 2009, *RAA*, **9**, 1205
Boettcher, M., Reimer, A., Sweeney, K., & Prakash, A. 2013, *ApJ*, **768**, 54
Bottino, A., Donato, F., Fornengo, N., & Scopel, S. 2004, *PhRvD*, **70**, 015005
Bringmann, T., & Weniger, C. 2012, *PDU*, **1**, 194
Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
Ceruti, M., Zech, A., Boisson, C., & Inoue, S. 2015, *MNRAS*, **448**, 910
Chakraborty, N., & Fields, B. D. 2013, *ApJ*, **773**, 104
Clark, G. W., Garmire, G. P., & Kraushaar, W. L. 1968, *ApJL*, **153**, L203
CTA Consortium 2011, arXiv:1111.2183
Dermer, C. D. 2007, in AIP Conf. Ser. 921, The First GLAST Symposium, ed. S. Ritz, P. Michelson, & C. A. Meegan (Melville, NY: AIP), 122
Di Mauro, M., Calore, F., Donato, F., Ajello, M., & Latronico, L. 2014a, *ApJ*, **780**, 161
Di Mauro, M., Cuoco, A., Donato, F., & Siegal-Gaskins, J. M. 2014b, *JCAP*, **1411**, 021
Di Mauro, M., & Donato, F. 2015, *PhRvD*, **91**, 123001
Di Mauro, M., Donato, F., Lamanna, G., Sanchez, D. A., & Serpico, P. D. 2014c, *ApJ*, **786**, 129
Domínguez, A., Finke, J. D., Prada, F., et al. 2013, *ApJ*, **770**, 77
Dubus, G., Contreras, J. L., Funk, S., et al. 2013, *A&A*, **43**, 317
Eddington, A. S. 1913, *MNRAS*, **73**, 359
Fazio, G. G., & Stecker, F. W. 1970, *Natur*, **226**, 135
Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
Fichtel, C. E., Hartman, R. C., Kniffen, D. A., et al. 1975, *ApJ*, **198**, 163
Fichtel, C. E., Simpson, G. A., & Thompson, D. J. 1978, *ApJ*, **222**, 833
Fields, B. D., Pavlidou, V., & Prodanovic, T. 2010, *ApJL*, **722**, L199
Finke, J. D., Razzaque, S., & Dermer, C. D. 2010, *ApJ*, **712**, 238
Fornasa, M., & Sánchez-Conde, M. A. 2015, *PhR*, **598**, 1
Fossati, G., Maraschi, L., Celotti, A., Comastri, A., & Ghisellini, G. 1998, *MNRAS*, **299**, 433
Franceschini, A., Rodighiero, G., & Vaccari, M. 2008, *A&A*, **487**, 837
Ghisellini, G., Celotti, A., Fossati, G., Maraschi, L., & Comastri, A. 1998, *MNRAS*, **301**, 451
Ghisellini, G., Tavecchio, F., Foschini, L., et al. 2010, *MNRAS*, **402**, 497
Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
Gould, R., & Schröder, G. 1966, *PhRvL*, **16**, 252
Hooper, D. 2016, *JCAP*, **09**, 002
Hooper, D., Linden, T., & Lopez, A. 2016, *JCAP*, **08**, 019
Inoue, Y. 2011, *ApJ*, **733**, 66
Kalashev, O. E., Semikoz, D. V., & Sigl, G. 2009, *PhRvD*, **79**, 063005
Kraushaar, W. L., Clark, G. W., Garmire, G. P., et al. 1972, *ApJ*, **177**, 341
Lacki, B. C., Horiuchi, S., & Beacom, J. F. 2014, *ApJ*, **786**, 40
Lee, S. K., Lisanti, M., & Safdi, B. R. 2015, *JCAP*, **1505**, 056
Lee, S. K., Lisanti, M., Safdi, B. R., Slatyer, T. R., & Xue, W. 2016, *PhRvL*, **116**, 051103
Linden, T., Rodd, N. L., Safdi, B. R., & Slatyer, T. R. 2016, arXiv:1604.01026
Loeb, A., & Waxman, E. 2000, *Natur*, **405**, 156
Makiya, R., Totani, T., & Kobayashi, M. A. R. 2011, *ApJ*, **728**, 158
Malyshev, D., & Hogg, D. W. 2011, *ApJ*, **738**, 181
Massaro, F., & Ajello, M. 2011, *ApJL*, **729**, L12
Muecke, A., & Pohl, M. 1999, in ASP Conf. Ser. 159, BL Lac Phenomenon, ed. L. O. Takalo & A. Sillanpää (San Francisco, CA: ASP), 217
Murase, K., Ahlers, M., & Lacki, B. C. 2013, *PhRvD*, **88**, 121301
Murase, K., Beacom, J. F., & Takami, H. 2012, *JCAP*, **1208**, 030
Murase, K., Guetta, D., & Ahlers, M. 2016, *PhRvL*, **116**, 071101
Murase, K., Inoue, S., & Nagataki, S. 2008, *ApJL*, **689**, L105
Narumoto, T., & Totani, T. 2006, *ApJ*, **643**, 81
Pavlidou, V., & Fields, B. D. 2002, *ApJL*, **575**, L5
Pavlidou, V., & Venters, T. M. 2008, *ApJ*, **673**, 114
Sandage, A. R., Kron, R. G., & Longair, M. S. 1995, The Deep Universe (Berlin: Springer)
Singal, J., Petrosian, V., & Ajello, M. 2012, *ApJ*, **753**, 45
Soltan, A. M., & Juchniewicz, J. 1999, *ApL*, **39**, 197
Sreekumar, P., Bertsch, D. L., Dingus, B. L., et al. 1998, *ApJ*, **494**, 523
Stawarz, L., Kneiske, T. M., & Kataoka, J. 2006, *ApJ*, **637**, 693
Stecker, F. W., de Jager, O. C., & Salamon, M. H. 1992, *ApJL*, **390**, L49
Stecker, F. W., & Salamon, M. H. 1996, *ApJ*, **464**, 600
Stecker, F. W., Salamon, M. H., & Malkan, M. A. 1993, *ApJL*, **410**, L71
Stecker, F. W., & Venters, T. M. 2011, *ApJ*, **736**, 40
Strong, A. W., Moskalenko, I. V., & Ptuskin, V. S. 2007, *ARNPS*, **57**, 285
Strong, A. W., Moskalenko, I. V., & Reimer, O. 2004, *ApJ*, **613**, 956
Su, M., Slatyer, T. R., & Finkbeiner, D. P. 2010, *ApJ*, **724**, 1044
Tamborra, I., Ando, S., & Murase, K. 2014, *JCAP*, **1409**, 043
Tavecchio, F. 2014, *MNRAS*, **438**, 3255
Taylor, A. M., Ahlers, M., & Hooper, D. 2015, *PhRvD*, **92**, 063011
Ullio, P., Bergstrom, L., Edsjo, J., & Lacey, C. G. 2002, *PhRvD*, **66**, 123502
Urry, C. M., & Padovani, P. 1995, *PASP*, **107**, 803
Zandanel, F., Tamborra, I., Gabici, S., & Ando, S. 2015, *A&A*, **578**, A32
Zdziarski, A. A., & Böttcher, M. 2015, *MNRAS*, **450**, L21
Zechlin, H.-S., Cuoco, A., Donato, F., Fornengo, N., & Regis, M. 2016a, *ApJL*, **826**, L31
Zechlin, H.-S., Cuoco, A., Donato, F., Fornengo, N., & Vittino, A. 2016b, *ApJS*, **225**, 18