

Luminosity functions consistent with a pulsar-dominated Galactic Center Excess

by

Jack T. Dinsmore

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

BACHELOR OF SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

© 2022 Massachusetts Institute of Technology
All Rights Reserved

The author hereby grants to MIT permission to reproduce and distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Author
Department of Physics
April 30, 2022

Certified by.....
Professor Tracy Slatyer
Thesis Supervisor, Department of Physics

Accepted by.....
Professor Depto Chakrabarty
Physics Associate Head, Department of Physics

Luminosity functions consistent with a pulsar-dominated Galactic Center Excess

by

Jack T. Dinsmore

Submitted to the Department of Physics
on April 30, 2022, in partial fulfillment of the
requirements for the degree of Bachelor of Science.

Abstract

A new population of millisecond pulsars is a long-standing proposed explanation for the excess of GeV-scale gamma rays emanating from the region surrounding the center of the Milky Way (the “Galactic Center excess”). We examine several simple parameterizations of possible luminosity functions for this population, as well as several benchmark luminosity functions proposed in the literature, and compare the predicted populations of resolved point sources to the *Fermi* 4FGL-DR2 point source catalog and a sub-population recently identified using wavelet-based methods. We provide general results that can be used to translate upper limits on the number of resolved point sources associated with the excess, and the fraction of the flux in the excess that can be attributed to resolved sources, into limits on the luminosity function parameter space. We discuss a number of important systematic uncertainties, including in the detection threshold model and the total flux attributed to the excess. We delineate regions of parameter space (containing existing benchmark models) where there is no apparent tension with current data, and the number of total pulsars needed to explain the excess is in the range of $\mathcal{O}(10^{4-5})$. In the future, lowered point source detection thresholds could be achieved either by new analysis methods or new data. An order-of-magnitude reduction in the sensitivity threshold (which may already be achieved by novel analyses probing sub-threshold source populations) could hope to resolve more than 30% of the flux of the excess even in pessimistic scenarios.

Thesis Supervisor: Tracy Slatyer

Title: Associate Professor of Physics

Acknowledgments

I am lucky to have had the support of many on my journey to this thesis. First, I would like to thank Professor Tracy Slatyer. Her constant and generous support as my academic advisor since 2019, and the example she set as a conscientious scientist while supervising our research together, has been a constant source of inspiration throughout my undergraduate years. In graduate school and beyond, I aspire to be as dedicated a mentor to my students as she was to me. I would also like Professor Julien de Wit, who gave me the resources and encouragement I needed to turn a class project into a full-fledged paper, far beyond what I could have hoped for. Since the beginning of my time at MIT, Professor Phil Harris has helped me personally as a recitation instructor, research supervisor, and lecturer, and I am lucky to have had his support. I would also like to thank Professor Josh Pepper of Lehigh University, who gave me my first glimpse of starting a research project from its roots. Finally, I would like to thank Professor Jennie Traschen of UMass Amherst, who spent an hour per week over the summer with a high school student who did not even know calculus, in an internship which started me on my path to astrophysics.

I would also like to thank my friends at MIT, in the physics department and elsewhere, who have remained a source of joy to me for all four years. Finally, I would like to thank my family, who have remained my most steadfast supporters and most treasured role models throughout my life. I could never have gotten to where I am without their love, patience, generosity, and guidance.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 17 |
| 1.1 | The Galactic Center Excess | 17 |
| 1.2 | Millisecond pulsars | 19 |
| 1.3 | MSP luminosity functions | 21 |
| 1.4 | This work | 22 |
| 2 | Models | 25 |
| 2.1 | Modeling the GCE | 25 |
| 2.1.1 | GCE spatial distribution | 25 |
| 2.1.2 | Total GCE flux | 26 |
| 2.2 | Modeling detected point sources | 31 |
| 2.2.1 | Observables | 31 |
| 2.2.2 | Sensitivity models | 38 |
| 2.3 | Luminosity functions | 41 |
| 2.3.1 | General parameterizations | 41 |
| 2.3.2 | Benchmark configurations | 42 |
| 3 | Results | 47 |
| 3.1 | Allowed luminosity function configurations | 47 |
| 3.2 | Number of MSPs in the GCE | 56 |
| 3.3 | Flux distribution | 56 |
| 4 | Discussion | 61 |

| | | |
|----------|---|-----------|
| 4.1 | Systematic uncertainties | 61 |
| 4.1.1 | Dependence on assumed GCE total flux | 61 |
| 4.1.2 | Dependence on sensitivity threshold | 63 |
| 4.1.3 | N_r , R_r , and N_{GCE} with cut around Galactic Center | 64 |
| 4.2 | Future sensitivity | 65 |
| 4.2.1 | Resolved PS flux distributions at higher sensitivity | 66 |
| 4.2.2 | Shape differences between resolved PS flux distributions at higher sensitivity | 69 |
| 5 | Conclusions | 73 |
| A | Scaling of ROIs and spectral ranges | 77 |
| B | Broken power-law fits to GCE spectra | 79 |
| C | Fits to GCE and AIC luminosity functions | 83 |
| D | Conversion between GCE luminosity and flux | 85 |
| E | Calculation of photon energy | 87 |
| F | Analysis of L_{th} for the simplified sensitivity model | 89 |
| G | Flux histograms for the smoothed sensitivity model | 93 |

List of Figures

| | | |
|-----|---|----|
| 2-1 | Previously-derived spectra of photon flux from the GCE in $F_\gamma = E^2 dN_\gamma/dE$, integrated over the ROI with $ l < 20^\circ$ and $2^\circ < b < 20^\circ$, selected from nine analyses of the GCE [60, 16, 25, 2, 32, 7, 6] (note some of these references include multiple analyses). 1σ error bars are reproduced from the same references. Arrows on error bars denote upper limits (i.e. because the 1σ error bars overlap zero). | 28 |
| 2-2 | Spectrum produced by Ref. [25], with a broken power-law fitted with all parameters free (green) and another broken power-law using parameters given by Ref. [16], allowing only the normalization to vary. Arrows on error bars denote upper limits (i.e. because the 1σ error bars overlap zero). | 29 |
| 2-3 | Total flux of the GCE in our $40^\circ \times 40^\circ$ ROI with a $ b \leq 2^\circ$ cut around the Galactic Disk, as determined by the three integration methods discussed in the text, with spectrum range 0.1 GeV – 100 GeV. Also shown (solid black lines) are results from a fit and integral over the 0.1 GeV – 10 GeV domain, to demonstrate that the high energy tail does not dominate the GCE’s flux. | 30 |
| 2-4 | Positions of the 41 wavelet-selected MSP candidates within 0.3° of a 4FGL-DR1 PS (blue circles). Also shown are the 6 PSs added when the maximum separation is increased to 0.55° (orange circles). The PSs’ radii are drawn proportional to their flux listed in the 4FGL-DR1 catalog for the 0.1-100 GeV energy bin. The gray band represents the $ b \leq 2^\circ$ cut masking the Galactic disk. | 34 |

| | | |
|-----|--|----|
| 2-5 | Comparison between the gamma-ray flux distributions of sources in the 4FGL-DR2 (blue) and 4FGL-DR1 (orange) catalogs, and the wavelet-selected 4FGL-DR1 sources of Ref. [60], within our ROI and after the cuts are applied (see text for details). Fluxes are obtained from the relevant catalogs and are evaluated for the 0.1-100 GeV energy band. We show results including only flagged catalog sources (solid blue/orange lines), flagged+unflagged catalog sources (shaded blue/orange regions), and wavelet-selected sources (solid black lines). The vertical line denoting F_{th} represents the weighted average flux threshold of the 4FGL-DR2 catalog within our ROI, as computed in appendix F. | 37 |
| 2-6 | Flux threshold $F_{\text{th}}(\ell, b)$ to include a PS within our ROI in the 4FGL-DR2 catalog (left panel) [4, 9], and the ratio of the 4FGL DR-2 and 4FGL DR-1 flux thresholds (right panel). Fluxes are computed in the 0.1-100 GeV energy band. | 39 |
| 2-7 | <i>Left:</i> Seven benchmark luminosity functions for GCE MSP luminosity functions, normalized to give the luminosity probability density as in eqs. 2.6 to 2.8, for (left panel) dN/dL , (right panel) LdN/dL , and (bottom panel) L^2dN/dL . See table 2.1 for a summary of these luminosity functions. | 46 |

- 3-1 N_r , R_r , and N_{GCE} values for power-law (top two rows) or log-normal (bottom row) luminosity functions, normalized to explain the full GCE flux as determined by Ref. [25]. The benchmark configurations are marked (see text for details). The N_r values (solid) are marked as fractions of the 4FGL-DR2 PSs passing our cuts (see text); the R_r values (dashed) are marked as fractions of the total GCE flux. These fractions are labeled on the contours. The regions where N_r (R_r) is less than 5% of the number of observed resolved sources (GCE flux) are marked with | (-), while regions where both numbers are below this threshold are marked with +. The gray-scale contours indicate the total number of sources required to explain the GCE. The top row (green) fixes $L_{\text{min}} = 10^{29}$ erg/s while the second row (blue) fixes $\alpha = 1.94$. 49
- 3-2 Luminosity functions with fixed $L_{\text{min}} = 10^{31}$ erg s⁻¹ and $\alpha = 1.94$ for the power-law case, (*left*), and fixed $\sigma = 0.5$ for the log-normal case (*right*). L_{max} and L_0 respectively are allowed to vary, with large values at the top. Luminosities greater than the threshold luminosity (dotted line) determined in appendix F are shaded, and the shaded area is proportional to N_r . The top and bottom panels of each figure show the highest and lowest luminosity functions explicitly, so that the vertical axis is labeled. 50
- 3-3 Distribution of resolved PSs predicted by each luminosity function with the standard sensitivity model, compared to the observed distribution of sources passing our cuts in the 4FGL-DR2 catalog (left panel, black dashed line), in the 4FGL-DR2 catalog with flagged sources removed (left panel, black dotted line), and in the wavelet-selected subsample (both panels) with a proximity cut of 0.3° (gray histogram) and 0.55° (hollow histogram). The predicted resolved PS distribution is also shown for the Wavelet 1 benchmark using the simplified sensitivity model (light blue points). Error bars are set to \sqrt{N} , and may therefore be inaccurate for low-count bins. 57

4-1 The parameter space that does not overproduce the wavelet-selected subsample of 4FGL-DR2 PSs ($N_r \leq 47$ and $R_r \leq 0.17$) for the power-law luminosity function (top two plots) and log-normal luminosity function (bottom plot). The total GCE flux has been increased by factors of 2–5 and decreased by factors of 2 and 3; contours of fixed N_r (solid) and fixed flux in resolved PSs (dashed) are labeled with the factor by which F_{GCE} has been multiplied. The colored lines exactly match the wavelet-selected subsample for the nominal GCE flux value (extracted from Ref. [25]). The left power-law plot fixes $\alpha = 1.94$ while the right fixes $L_{\text{min}} = 10^{29} \text{ erg s}^{-1}$ 62

4-2 Histograms of the predicted number of resolvable pulsars required by each luminosity function benchmark studied here to reproduce the GCE, for different models of *Fermi* sensitivity. Darker lines indicate the current sensitivity; lighter lines represent increases in sensitivity by factors of two, five, ten, and twenty. The vertical dotted line represents the energy flux corresponding to a source with an expected photon count of one in the part of the *Fermi* data set used by Ref. [45]. The bins are spaced evenly in log flux space, with width $\Delta F/F = 0.069$. . . 67

4-3 Map of the expected value of λ resulting from assuming one luminosity function (horizontal axis) as the hypothesis while the other (vertical axis) is the true luminosity function. Each luminosity function is normalized to produce the same number of resolved PSs as the assumed-true luminosity function. Each plot represents a different multiplicative increase in the *Fermi* sensitivity level. Diagonals indicate the fit of a luminosity function to itself. The standard sensitivity model is used. . . 71

| | | |
|-----|---|----|
| B-1 | Broken power-law fits to the nine spectra studied in section 2.1.2 to extract the total GCE luminosity. Arrows on error bars denote bars that would descend past the lower edge of the plot. See text for which functions are being fitted. Fits are performed approximating the uncertainties at different energies as uncorrelated, which is why the results of the two fits differ slightly for the “Calore 2015” spectrum taken from Ref. [16]; Ref. [16] took the full covariance matrix into account. . . . | 80 |
| C-1 | Log normal fits to two numerical luminosity functions: the boxy bulge luminosity function from Ref. [53] and the luminosity function from [29]. Uncertainties are assumed to be uncorrelated and represent one standard deviation from the mean. The best fit for a broken power-law with an exponential cutoff is also shown. | 84 |
| F-1 | As figure 3-1, but using the simplified sensitivity model with the average flux threshold of $L_{\text{th}} = 3.4 \times 10^{34} \text{ erg s}^{-1}$ | 91 |
| G-1 | Histogram of the predicted number of resolvable pulsars required by each luminosity function benchmark to reproduce the GCE flux, for different models of <i>Fermi</i> sensitivity under the smoothed sensitivity model. Darker lines indicate the current sensitivity; lighter lines represent increases in sensitivity by factors of two, five, ten, and twenty. See text for details. The bins are spaced evenly in log flux space, with width $\Delta F/F = 0.069$ | 94 |
| G-2 | Map of the expected value of λ resulting from assuming one luminosity function (horizontal axis) as the hypothesis while the other (vertical axis) is the true luminosity function. Each luminosity function is normalized to produce the same number of resolved PSs as the assumed-true luminosity function. Each plot represents a different multiplicative increase in the <i>Fermi</i> sensitivity level. Diagonals indicate the fit of a luminosity function to itself. The smoothed sensitivity model is used. | 95 |

List of Tables

| | | |
|-----|--|----|
| 2.1 | Names, references, and parameters for all seven luminosity functions studied in this paper. All luminosities are given in units of erg s^{-1} . The “Fit?” column contains “Y” when we obtained the given parameters via a least-squares fit to numerical data provided in the reference, and “Pulsars modeled?” contains “Y” when the luminosity function explicitly relies on a pulsar model in their analysis. | 45 |
| 3.1 | Number of resolved PSs, ratio of resolved flux to total flux, and total number of PSs predicted to make up the GCE based on seven luminosity function benchmarks (see section 2.3) and the requirement that the PSs reproduce the entire flux of the GCE, for three different sensitivity models (see section 2.2.2). | 54 |
| 3.2 | p -values comparing the predicted resolved flux distributions for each benchmark luminosity function to the observed distribution, where the overall normalization of the predicted distributions is chosen to be equal to the number of observed resolved PSs in the wavelet-selected subsample of 4FGL-DR2 sources. | 59 |
| 4.1 | Number of resolved MSPs N_r and ratio of resolved flux to total flux R_r for seven luminosity function benchmarks. The flux threshold of the telescope, $F_{\text{th}}(\ell, b)$, has been decreased by a factor of one (i.e., sensitivity is at its current value), two, five, ten, and twenty, with the factor of increase given as superscripts in the header. Apart from the change in threshold, the standard sensitivity model is used. | 66 |

| | | |
|-----|---|----|
| F.1 | Number of resolved PSs, ratio of resolved flux to total flux, and total number of PSs predicted to make up the GCE based on the new $L_{\text{th}} = 3.56 \times 10^{34} \text{ erg s}^{-1}$ simplified sensitivity threshold produced by a flux-weighted average of per-pixel sensitivity provided by Refs. [4, 9]. Compare to table 3.1b, which produces higher values. | 91 |
|-----|---|----|

Chapter 1

Introduction

1.1 The Galactic Center Excess

The *Fermi* Gamma Ray Space Telescope was launched in 2008 to shed light on Galactic and extra-Galactic sources of gamma-rays, which are emitted by some of the most extreme astrophysical events of the Universe, including certain dark matter processes in some models. *Fermi* data, specifically from the Large Area Telescope (LAT), has since led to the discovery of the Fermi Bubbles [56] in 2010, as well as gamma-ray emission from pulsars, supernovae, and many other high-energy phenomena.

In 2009, *Fermi* also detected an unexpected excess of gamma-rays emanating from the region surrounding the Galactic Center (GC) [31, 36, 33]. This signal, known as the Galactic Center Excess (GCE), has a spectral energy distribution that peaks around 1–3 GeV in $E^2 dN/dE$ [25, 16, 60, 32, 7, 6]¹ and early studies found its spacial morphology to be approximately spherically symmetric. These properties suggested a possible dark matter self-annihilation origin for the GCE.

Dark matter is a placeholder term for several connected, poorly understood astrophysical phenomena. The first of these phenomena to be discovered was found by Vera Rubin, Kent Ford, and (separately) Kenneth Freeman in 1970: a tendency for stars near the edges of their host galaxies to orbit more quickly than Newtonian

¹ dN/dE is the number spectral density of photons, meaning the number of photons in an infinitesimal energy bin of width dE , divided by the bin width. Thus, EdN/dE is the energy density. $E^2 dN/dE = EdN/d \ln E$ is the energy density normalized by the bin width in logarithmic space.

gravity predicted [54, 28]. Further inconsistencies with Newtonian gravity were later detected in anisotropies of the Cosmic Microwave Background, gravitational lensing events, galaxy cluster collisions, and more. All of these phenomena can be explained by the presence of a particle or particles which interact so weakly with visible matter that they cannot be detected with telescopes nor experiments in the lab. Only their gravitational interactions reveal their presence.

Many theories have been proposed which produce dark matter particles consistent with observations, and many further predict hitherto undiscovered properties which may give rise to detectable astrophysical signals. One such example is dark matter annihilation, which is a process in which two dark matter particles collide, annihilate, and produce light. This process would have to be extremely rare because it has not yet been detected. But where dark matter gathers in high densities (e.g., at the center of galaxies), the annihilation photons may be just bright enough to observe.

The GCE is a candidate for dark matter annihilation for the following reasons [31, 33, 24, 15, 2]. The Milky Way’s dark matter halo is expected to be roughly spherically symmetric in the GC, like the GCE, with density-squared consistent with the observed flux density of the GCE.² Aside from the halo density, the predicted flux density from dark matter annihilation is also a function of the gamma ray spectral energy density per annihilation event, the annihilation cross section, the dark matter decay channel, and particle mass. Ref. [31] fit these parameters to the flux morphology and spectrum and found values consistent with dark matter models.

Ref. [31] mentions the caveat that background point sources (PSs) provide background diffuse emission which is difficult to subtract from annihilation emission. Indeed, the GCE energy spectrum and morphology also appear to be consistent with a population of largely unresolved Millisecond Pulsars (MSPs), within systematic uncertainties (e.g. [1, 33, 59, 16, 52]). A number of studies using alternate approaches to model the Galactic diffuse emission background have found that the GCE correlates better with the mass distribution of the Galactic bulge rather than the spherical dis-

²The square of dark matter density is compared to the GCE density because dark matter annihilation is a two-particle interaction, so that its probability is proportional to the square of the particle density.

tribution expected from dark matter annihilation [47, 12, 48, 3, 22], although Ref. [25] arrives at the opposite conclusion. Studies of gamma-ray hot spots within the inner Galaxy, and/or the photon statistics of the GCE, have claimed detection of PSs associated with the GCE [11, 43, 44, 60, 13], further supporting the MSP interpretation. However, recent work has shown that some strong earlier claims of PS detection were affected by systematic biases favoring the PS interpretation [40, 42, 41, 23], and the interpretation of Ref. [11] as evidence for PSs that are part of the GCE has been questioned in Ref. [60]. Machine-learning-based analyses have claimed evidence in favor of a PS component [17, 46, 45, 49], but with modest statistical significance (and the potential for as-yet-unaccounted-for systematic errors): Ref. [49] attributes $38^{+9}_{-19}\%$ of the GCE to PSs (i.e. a roughly 2σ detection), and Ref. [45] excludes a smooth (non-PS) fraction for the GCE exceeding 66% at 95% confidence. Alternate photon-statistics methods have found evidence for faint PSs in the inner Galaxy but cannot yet discern whether those PSs are associated with the GCE [14].

1.2 Millisecond pulsars

This thesis assesses the MSP explanation for the GCE without assuming a particular MSP model (unlike Refs. [53], [29], and [37], which do). We briefly review the general properties of neutron stars, pulsars, and MSPs before turning to the MSP models used in this thesis.

A neutron star is one of three possible endpoints of stellar evolution, the other two being white dwarfs (lighter) and black holes (heavier). Their existence was first hinted at with Subrahmanyan Chandrasekhar’s discovery in 1931 that a white dwarf, which is internally buttressed against gravitational collapse by electron degeneracy pressure³, must collapse at masses greater than 1.4 Solar masses due to general relativity (the Chandrasekhar limit) [18]. Two years after the subsequent discovery of the neutron

³Electron degeneracy pressure is a pressure driving very close atoms apart due to the Pauli Exclusion of their electron orbitals. Unlike the radiation pressure which supports living stars against gravity, electron degeneracy pressure does not vanish at low temperature, hence its importance for a dead (cold) star.

in 1932, Fritz Zwicky and Walter Baade proposed that such a massive star might collapse through the white dwarf phase into a dense star of 1.4–2.1 or so Solar masses, which are supported not by electron degeneracy pressure but by nuclear forces [8]. The nature of these forces is sufficiently complicated that predicting neutron star properties is difficult. However, enough is known to estimate that most neutron stars are $\mathcal{O}(10\text{ km})$ in radius.

A star many thousands of kilometers across collapsing into such a small object as a neutron star must gain enormous rotational velocity due to conservation of angular momentum. Through complicated dynamo-like phenomena, these stellar remnants also form large magnetic fields (billions of Tesla in strength [58]) which may not be aligned with the spin pole of the neutron star. The net effect is a concentrated beam of high energy gamma-rays emanating from both magnetic poles of the neutron star. These sweep across the sky as the neutron star rotates at great speeds, and for some neutron stars, one of the beams periodically happens to cross Earth. Such neutron stars are observed as point sources which periodically peak in brightness at very regular intervals, and are called “pulsars.”

Pulsar periods vary greatly, but one class of pulsars are those with period less than about 10 milliseconds. These are millisecond pulsars, which are commonly found in globular clusters. Even though MSPs’ flux oscillates rapidly, the objects are often sufficiently distant that less than one gamma-ray photon is detected per rotation period on average. The MSP will therefore be visible as a point source, but not with the characteristic millisecond period of an MSP. The classification of the pulsar as an MSP must be made through other means.

Ref. [57] predicted a population of these MSPs to exist in the Galactic center with a spectrum peaking at a few GeV. The connection between this population and the GCE (which also peaks at a few GeV) was made by Ref. [1] and is the focus of this thesis.

1.3 MSP luminosity functions

Studies that aim to test the MSP explanation inherently rely on the luminosity function dN/dL of the putative MSP population, i.e., the number density of MSPs as a function of their luminosity. There are two long-standing questions regarding the predictions made by plausible models for the MSP gamma-ray luminosity function: (1) the number of GCE PSs that should be detected individually, or in analyses that probe populations of PSs just below the sensitivity threshold of the relevant telescopes, and (2) the overall number of MSPs required to explain the excess. The first prediction depends primarily on the bright end of the luminosity function, and can be confronted with the number of observed sources or the fraction of the GCE that appears to be due to near-threshold sources. The second prediction is often controlled by the properties of low-luminosity MSPs whose emission can only be observed in aggregate (if at all), and can be compared with theoretical models or empirical inferences for the total number of MSPs ever produced in the Milky Way.

Studies using different models and parameterizations for the luminosity function have given widely varying answers to these questions, leading to differing conclusions on the viability of the MSP hypothesis. Ref. [37], following earlier work in Refs. [20, 38], analyzed the luminosity function for MSPs detected in globular clusters (GLCs) and found that if GCE MSPs had the same luminosity function, they could make up only a few percent of the excess. Refs. [21, 34] calibrated the anticipated number of bright MSPs to observed low-mass X-ray binaries (LMXBs), based on a scenario where MSPs form from LMXB progenitors, and found that such MSPs can contribute only $< 23\%$ of the excess. More recently, however, Ref. [53] has argued that a physical model for the luminosity function of MSPs in the Galactic Bulge, based on observed MSPs in the Galactic disk and correlating other MSP characteristics with their luminosity [39], can fit the excess without overproducing bright sources. This work built on a previous study of MSPs in the Galactic disk [10], which constrained several simple parameterizations of the luminosity function using observational data, and found a luminosity function distinctly different from that inferred by Ref. [37].

Ref. [29] argues that the Galactic Bulge could plausibly host $\mathcal{O}(10^5)$ MSPs formed by accretion-induced collapse (AIC), which would bypass the limit from non-observation of LMXBs, and also predicts a total flux broadly consistent with the GCE, using a luminosity function model based on the results of Ref. [53].

Other studies have simply fitted a parameterized luminosity function to gamma-ray data from the inner Galaxy region. Studies using the Non-Poissonian Template Fitting (NPTF) method [44, 50] have generally used broken power law source count functions to describe a sub-threshold PS population associated with the GCE, and have inferred quite steeply peaked source count functions with most power in sources just below *Fermi*'s PS sensitivity threshold [44, 13]. Ref. [44] consequently requires only a small number of MSPs, $\mathcal{O}(400)$, to explain the entire GCE. Ref. [11] assumed a power-law luminosity function, $dN/dL \propto L^{-\alpha}$ with $\alpha = 1.5$, with cutoffs at minimum and maximum luminosities L_{\min}, L_{\max} , and constrained L_{\max} from the data. The cutoffs are required to prevent the total number of MSPs from diverging. Ref. [60] took a similar approach, but allowed the power-law slope to vary as well as L_{\max} ; their preferred parameters imply $\mathcal{O}(3 \times 10^6)$ MSPs are needed to explain the GCE, primarily because they consider a luminosity function with a steep power-law slope ($\alpha \gtrsim 1.8 - 1.9$) and extrapolate to a relatively small value of L_{\min} . Non-parametric fits to the source count function have also been employed, e.g. in Ref. [45].

1.4 This work

This thesis is based on the work I conducted with Tracy Slatyer in Ref. [26] and has been submitted for publication in *JCAP*. We aim to clarify the differences between these studies; study simple, commonly-used parameterizations of the MSP luminosity function to understand which regions of parameter space remain viable; and explore the resulting range of predictions for ongoing and future PS searches with increased sensitivity. A similar approach was taken by Ref. [52], albeit considering only power-law luminosity functions with sharp cutoffs at minimum and maximum luminosity values. We update and expand on this earlier work by considering a wider range of

luminosity function parameterizations; using up-to-date PS sensitivity information given the most recent PS catalog produced by the *Fermi*-LAT Collaboration; and studying specific models and fits discussed in the literature in recent years.

We begin in section 2.1 by describing our modeling of the GCE. In particular, in section 2.1.2 we discuss the gamma-ray flux we attribute to the GCE, in order to define what we mean by “explaining the GCE”. There are large systematic uncertainties on this total flux, at the factor-of-two level, and this may in itself be responsible for some differences in the literature; identical populations of PSs may explain 10% of the GCE in one study and 20% in another, due to different inferred total fluxes. We consider a range of total fluxes for the GCE drawn from Refs. [25, 16, 60, 32, 7, 6].

In section 2.2, we discuss the population of visible point sources we employ in this analysis and describe how we relate the properties of this population to the MSP luminosity functions (section 2.2.1). Our constraints and forecasts require understanding the sensitivity threshold for point source detection; we discuss several alternative models for this sensitivity threshold in section 2.2.2. In section 2.3 we discuss the luminosity function models we test in this work.

We present our main results in section 3, and discuss some important systematic uncertainties affecting those results in section 4.1. In section 4.2 we discuss the degree to which future analyses and/or observations with increased point source sensitivity would increase the fraction of resolved flux, under our various luminosity function models, and allow us to distinguish between different luminosity functions. We present our conclusions in section 5. Our appendices add detail on several of the intermediate steps needed for our main results, and provide supplementary results and calculations.

Chapter 2

Models

In order to begin making general statements about the MSP properties that could produce a signal like the GCE, we need to develop general models for the GCE, *Fermi*'s sensitivity, and MSP luminosity functions. We will use the GCE's spatial distribution and flux to model the number and locations of the MSPs, the telescope sensitivity to compare observations to our predictions, and the luminosity functions to predict the MSP brightnesses. These three models are discussed in the following three sections.

2.1 Modeling the GCE

2.1.1 GCE spatial distribution

We model the number density distribution of MSPs in the GC as the square of a generalized Navarro-Frenk-White (gNFW) profile [51], based on fits to the observed distribution of GCE flux. The gNFW profile is spherically symmetric, with radial distribution

$$\sqrt{\rho_{\text{GCE}}(r)} \propto \left(\frac{r}{r_s}\right)^{-\gamma} \left(1 + \frac{r}{r_s}\right)^{-3+\gamma}. \quad (2.1)$$

Following [16, 32, 25], we choose $\gamma \approx 1.2$ and $r_s = 20$ kpc, as these parameters match the empirical data reasonably well (although r_s is not strongly constrained by the GCE since the signal is only observed for $r \ll r_s$). Some references use $\gamma = 1$, in

which case the profile is called an “NFW profile” (not “generalized”), e.g. [60]. As discussed above, a number of studies have also found that the GCE is better described by a bulge-like density distribution than the gNFW² profile; we expect the effect of choosing a bulge-like profile instead of gNFW² to be rather small in our analysis, since the two profiles are quite similar where the GCE is bright. Furthermore, the only places we use the detailed spatial distribution of the GCE are (1) in computing the sensitivity to point sources, (2) in translating between flux and luminosity of individual sources, and (3) in computing the ratio of flux emitted between Regions of Interest (ROIs) of different size. We expect errors in the first two calculations due to an incorrect spatial distribution to be relatively small because the *Fermi* sensitivity map is fairly smooth, and in both models the GCE is quite concentrated in the inner Galaxy and originates from sources at a roughly constant distance from Earth. The largest effect of changing the assumed GCE spatial distribution may be via the variation in the inferred total flux, as discussed in section 2.1.2 below; however, this quantity has other substantial systematic uncertainties related to the choice of background modeling.

We study this gNFW distribution within an ROI with $|\ell| < 20^\circ$ and $2^\circ < |b| < 20^\circ$, where we have masked the Galactic disk (consistent with e.g. Refs. [60, 15]). We will generally report results over the energy range $0.1 \text{ GeV} < E_\gamma < 100 \text{ GeV}$, for ease of comparison to point source properties reported in the 4FGL point source catalog [4].

2.1.2 Total GCE flux

In order to discuss PS populations that could potentially generate the gamma-ray flux of the GCE, we need to describe the overall brightness of the GCE. We extract the GCE flux from several previous analyses of GCE energy spectra [60, 16, 25, 2, 32, 7, 6]. Although these studies draw their data from the same source (*Fermi* public data), the inferred GCE spectra differ due to choices in the fitting approach, ROI, and signal and background modeling. For example, to model the spatial distribution of the excess, Refs. [16, 32, 7] fix $\gamma = 1.2$, while Ref. [60] performs the analysis for both $\gamma = 1.0$ and $\gamma = 1.2$, and Refs. [25, 6, 2] allow γ to float in the fit. For the cases

in which γ is fitted, typical values lie in the range 1.0–1.3. The studies all use ROIs centered on the GC, ranging from a $40^\circ \times 40^\circ$ region without the Galactic Disk mask used in this paper, to a $7^\circ \times 7^\circ$ region. All ROIs are centered on $(\ell, b) = (0, 0)$. All studies fix $r_s = 20$ kpc, except for Refs. [2, 32], which use $r_s = 23.1$ kpc; in any case, the GCE spectrum is rather insensitive to r_s [25]. In order to compare studies with different ROIs, we re-scale the inferred flux by the method described in appendix A. The effect of varying γ on the inferred total flux within our ROI is non-negligible — for example, the total flux from the spectra attained by Ref. [60] assuming $\gamma = 1.0$ is $\sim 40\%$ larger than the flux inferred assuming $\gamma = 1.2$ — but as we will see, there are other systematic uncertainties of comparable magnitude.

The manner in which uncertainties in the energy spectrum are reported also varies; some studies report only statistical uncertainties, and some report both statistical and systematic. Refs. [16, 32, 25] report both separately, and for our purposes, we add these in quadrature (this approach may lead to an overestimate of uncertainties in some cases since it neglects correlations between systematic uncertainties).

Figure 2-1 displays all the spectra mentioned above, with ROI rescaling included. Many studies reported flux values in units of flux per steradian; we have multiplied those fluxes by the area of their respective ROIs and then rescaled the flux as described in appendix A to attain an absolute flux from the GCE in our ROI differential in energy. We report our results in terms of $F_\gamma = E^2 dN_\gamma/dE$, where N_γ is the number of incident photons from the ROI per unit exposure (measured in $\text{cm}^2 \text{ s}$).

We compare three methods of extracting the total GCE flux, integrated over energy, from these spectrum analyses. The first method is direct numerical integration of the binned spectrum. This method is most responsive to the data measured by *Fermi* and does not attempt to abstract over it with a smooth function, but it is potentially somewhat noisy and the energy range of the data varies between different studies. Therefore, we also test the effect of fitting a singly broken power-law to the data, and then analytically integrating this function to infer the total integrated GCE




Figure 2-1: Previously-derived spectra of photon flux from the GCE in $F_\gamma = E^2 dN_\gamma/dE$, integrated over the ROI with $|l| < 20^\circ$ and $2^\circ < |b| < 20^\circ$, selected from nine analyses of the GCE [60, 16, 25, 2, 32, 7, 6] (note some of these references include multiple analyses). 1σ error bars are reproduced from the same references. Arrows on error bars denote upper limits (i.e. because the 1σ error bars overlap zero).

flux in the energy range [0.1 GeV, 100 GeV]:

$$F_\gamma = F_0 \begin{cases} (E/E_b)^{-n_1} & E < E_b \\ (E/E_b)^{-n_2} & E > E_b \end{cases}, \quad (2.2)$$

We perform this fit in two ways: (1) where all four parameters of the broken power-law are allowed to float (the normalization constant F_0 , the turnover energy E_b , and the slopes above and below the turnover energy n_2 and n_1), and (2) where all parameters are fixed except F_0 , which is allowed to float. In the latter case we use the parameters determined by Ref. [16]: $E_b = 2.06^{+0.23}_{-0.17}$ GeV, $n_1 = 1.42^{+0.22}_{-0.31}$, $n_2 = 2.63^{+0.13}_{-0.095}$.

An example of the two fits applied to the spectrum of Ref. [25] is shown in figure 2-2. The best fits for all spectra are displayed in appendix B. Figure 2-3 displays the results for the integrated flux F_{GCE} via the three integration methods, from each GCE spectrum studied. In general we find that there is a substantial variation in




Figure 2-2: Spectrum produced by Ref. [25], with a broken power-law fitted with all parameters free (green) and another broken power-law using parameters given by Ref. [16], allowing only the normalization to vary. Arrows on error bars denote upper limits (i.e. because the 1σ error bars overlap zero).

F_{GCE} between different spectra presented in the literature; the variation associated with using different methods to model a given spectrum is comparatively small, and consistent within the nominal uncertainties. Figure 2-3 also shows the results of integrating the flux only up to an energy of 10 GeV, to test whether the variations between different analyses might be due to differences in the high-energy tail. We observe generally that the 10-100 GeV band provides only a subdominant component of the flux, and while this contribution is quite uncertain (being negligible in some analyses and quite substantial in others), it does not appear to be the main source of differences between analyses — there are large variations in the flux in the 0.1-10 GeV band as well.

The inferred flux differs by roughly a factor of two between the highest-flux and lowest-flux scenarios. It seems likely that this can be largely attributed to differences in the modeling of the background components in the fit, between the various analyses, which are known to induce substantial systematic differences in the extracted GCE even when all other analysis choices are held fixed (e.g. [16, 13, 46]). For ease of display, we will show baseline results assuming that “explaining the GCE” means




Figure 2-3: Total flux of the GCE in our $40^\circ \times 40^\circ$ ROI with a $|b| \leq 2^\circ$ cut around the Galactic Disk, as determined by the three integration methods discussed in the text, with spectrum range 0.1 GeV – 100 GeV. Also shown (solid black lines) are results from a fit and integral over the 0.1 GeV – 10 GeV domain, to demonstrate that the high energy tail does not dominate the GCE’s flux.

reproducing the full flux obtained from a broken power-law fit (with all parameters floated) to the data from Ref. [25], $F_{\text{GCE}} = 1.8 \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}$. The corresponding spectral fit parameters are $E_b = 1.35 \text{ GeV}$, $n_1 = 1.11$, $n_2 = 2.58$. This is a recent analysis that agrees well with earlier results from Refs. [16, 6]; it corresponds to a relatively high-flux scenario. In general this means that requiring a source population to generate the GCE predicts more numerous and observable sources, compared to a lower-flux scenario, leading to stronger constraints; a luminosity function that does not overproduce known point sources using this flux choice should also be allowed in a lower-flux scenario. To account for the large uncertainty in the total flux, we will show how our results change for different assumed total fluxes in section 4.1.

2.2 Modeling detected point sources

2.2.1 Observables

To constrain the luminosity function of a hypothetical MSP population responsible for the entire GCE, we track predictions of two observable properties of the population: N_r , the number of resolved PSs, and R_r , the ratio of the flux emitted by those PSs to the total flux of the GCE. Collectively, we call these two values “the observables.” Later, we will also discuss the distribution in flux of the resolved PSs. Predictions for the theoretical values of the observables, given a luminosity function, are discussed in section 2.2.2, and in this section we discuss observational constraints.

4FGL catalog

The 4FGL Point Source Catalog tracks the positions and fluxes of resolved gamma-ray PSs in the sky, in addition to many other PS properties such as their potential origin or any associated sources in other catalogs [4]. Two releases of the catalog currently exist: an 8-year release which we call Data Release 1 (4FGL-DR1)¹, and a 10-year release called 4FGL-DR2.² Earlier versions of the catalog also exist — a 4-year version labeled 3FGL and a 2-year version labeled 2FGL — and have been employed in previous analyses of the GCE.

There are a number of approaches one might take to constraining N_r from these catalogs. The most conservative (i.e. yielding the weakest limits) would be to simply include all sources within the ROI, but it is very unlikely that all such sources are associated with the GCE. Alternatively, one could include only sources with a spectrum sufficiently consistent with the GCE (although it is possible that the spectra of individual GCE sources differ somewhat from the aggregate spectrum), only sources known to be pulsars, or some other subsample. In this work, we use for our baseline sample the full set of 4FGL-DR2 sources in the ROI, minus those which are known to be either outside the GCE region or associated with non-pulsar source classes

¹https://fermi.gsfc.nasa.gov/ssc/data/access/lat/8yr_catalog/

²https://fermi.gsfc.nasa.gov/ssc/data/access/lat/10yr_catalog/

(which are dominantly extragalactic), on the basis that this should be a true upper limit on the number of possible GCE sources detected with the sensitivity relevant to 4FGL-DR2. We then show how the constraints would change if only a fraction of these sources belong to the GCE, in the hope that this will allow for easy translation of our constraints to future analyses that eliminate a larger fraction of 4FGL-DR2 sources as possible members of a GCE source population (or identify members of that population).

Specifically, to obtain our conservative upper limit for N_r , we remove all PSs in the 4FGL-DR2 catalog that are outside the ROI or have a listed association with a non-pulsar source. For the PSs associated with known pulsars, we remove those known to be farther than 2 kpc from the GC based on the ATNF pulsar catalog [35]. This 2 kpc distance corresponds to a maximum angular distance of 13° , but following Ref. [60], we do not exclude sources outside 13° from the GC by default, only those associated with a pulsar whose position in the ATNF catalog places it more than 2 kpc from the GCE.

There are $N_r = 265$ 4FGL-DR2 PSs passing these cuts, seven of which are associated with pulsars. Together, these PSs contribute $1.6 \times 10^{-9} \text{ erg s}^{-1} \text{ cm}^{-2}$ of flux, or $R_r = 91\%$ of the total GCE flux. This includes 4FGL-DR2 sources which have one or more analysis flags and thus should be treated with caution. Cutting the flagged PSs as well, as these are known to be affected by systematic errors and the instructions for use of the catalog indicate they should be used with great care, we have instead $N_r = 109$ and $R_r = 35\%$. In both cases, most of these sources are PSs with unknown origin and unknown distance from the GC, so only a fraction of them are likely to be GCE MSPs. Therefore, these estimates of N_r and R_r should be regarded as upper bounds on the true resolved source populations associated with the GCE.

Wavelet-selected subsample

Ref. [60] recently examined public *Fermi* Pass 8 data (version P8R3, recorded from 4 August 2008 to 20 February 2019), using wavelet-based methods first employed by Ref. [11] to identify a population of isolated spatial peaks. Ref. [11] had identified

these peaks with a sub-detection-threshold source population that could be contributing to the GCE. Ref. [60] found 115 significant peaks (test statistic $\mathcal{S} \geq 4$) within their ROI (which matches the one used in this analysis), and then compared the locations of these peaks to the 4FGL-DR1 catalog. 107 of these peaks were within 0.3° of 103 PSs in the 4FGL-DR1 catalog³, and the authors therefore concluded that the wavelet peaks could be well-approximated as a subset of the 4FGL-DR1 catalog.

The authors of Ref. [60] provided us with the locations and associations of their wavelet peaks, allowing us to reproduce their cuts. We found that of the 103 PSs associated with wavelets, 46 were excluded due to associations with non-pulsar sources (such as active galactic nuclei). A further 15 sources were associated with pulsars and excluded because of distance measures placing them outside 2 kpc of the GC. (Pulsar J1823-3021A was the only pulsar in the list of wavelet-selected sources known to be within this radius, but because it is also a globular cluster member, it was also excluded.) This left $N_r = 41$ resolved wavelet-selected Galactic MSP candidates in the 4FGL-DR1 catalog, contributing $R_r = 14\%$ of our baseline GCE flux. When the proximity cut between the flux peak and the 4FGL source was extended from 0.3° to 0.55° , six other sources were added, yielding $N_r = 47$ PSs, contributing $R_r = 17\%$ of total flux. Both R_r values are slightly smaller than those quoted in Ref. [60] because our baseline GCE flux (derived from Ref. [25]) is larger, as discussed in section 2.1.2. These 41+6 PSs are shown in figure 2-4; we will use them as an example subsample of 4FGL-DR2 that could be attributed to the GCE, and compare our results in this context with those of Ref. [60] (noting that the approach of using a wavelet-selected sample as a proxy for GCE sources was first advanced in Ref. [11]).

Constraints/benchmarks for observables

It is tempting to go further and use the wavelet-selected subsample of 47 PSs to define an upper limit on resolved GCE sources (or even a target number of resolved sources), as was done in Ref. [60]. However, without reproducing the pipeline of Ref. [60], we

³We further found that one of the 8 wavelet peaks not associated with a source in 4FGL-DR1 was coincident with a source in 4FGL-DR2.




Figure 2-4: Positions of the 41 wavelet-selected MSP candidates within 0.3° of a 4FGL-DR1 PS (blue circles). Also shown are the 6 PSs added when the maximum separation is increased to 0.55° (orange circles). The PSs' radii are drawn proportional to their flux listed in the 4FGL-DR1 catalog for the 0.1-100 GeV energy bin. The gray band represents the $|b| \leq 2^\circ$ cut masking the Galactic disk.

do not have a good model for the sensitivity of the wavelet-based selection, which makes it difficult to predict the expected number of sources for a given luminosity function. Ref. [60] estimates the sensitivity using a simple luminosity threshold where sources with luminosity greater than 10^{34} erg/s are all resolved. As a check on this approximation, we show in figure 2-5 the flux distributions of sources in the wavelet-selected subsample and the 4FGL-DR2 and 4FGL-DR1 catalogs (in the latter case, we show separately the results where flagged sources are included or excluded).

We see that within this ROI, the wavelet-selected sources are a subsample of the 4FGL catalogs at all flux levels; i.e. the difference in the total number of sources is not a matter of the wavelet selection having a higher sensitivity threshold relative to the 4FGL catalogs. In fact, the distribution of PSs with respect to flux is rather consistent across different cuts (note that some of the difference between e.g. 4FGL-DR1 and 4FGL-DR2 corresponds to sources moving between flux bins, not sources appearing or disappearing from the catalog), with the main difference being the total number of PSs. As a fraction f of our baseline 4FGL-DR2 sample (including flagged sources), the numbers of sources detected in 4FGL-DR1, 4FGL-DR2 with unflagged sources only, 4FGL-DR1 with unflagged sources only, and the wavelet method of Ref. [60] (with the 0.55° radius cut), correspond respectively to $f = 80\%$, 40% , 35% , and 18% . As noted above, the wavelet-selected PSs contribute roughly 17% of the GCE flux while the baseline 4FGL-DR2 sample (including flagged sources) contributes 91% of the GCE flux. Thus the wavelet-selected PSs contribute 19% of the 4FGL-DR2 sample in flux and 18% in number, consistent with their flux distribution being rather similar.

Consequently, it appears that a GCE PS might well be above both the sensitivity threshold for 4FGL-DR2, and the nominal sensitivity threshold suggested in Ref. [60], and still fail to be detected by the wavelet method (in the sense that there are many 4FGL-DR2 and 4FGL-DR1 sources that pass all the cuts in Ref. [60] and were not associated with a significant wavelet peak). Thus we will use the wavelet-selected subsample as an *example* of a resolved PS population that could be associated with the GCE, but not as a formal upper limit; to use it as a true upper limit would require

an in-depth study of the sensitivity and completeness properties of the wavelet method for identifying resolved PSs.

More generally, we will consider benchmarks for $N_r = 106, 53, 26,$ and $13,$ corresponding to populations generating 40%, 20%, 10%, and 5% of all the 4FGL-DR2 sources passing our cuts. We will separately consider benchmarks of $R_r = 40\%, 20\%, 10\%,$ and $5\%.$ In cases like the wavelet-selected population where the flux distribution of the resolved sources is similar to that of our 4FGL-DR2 sample, the N_r and R_r benchmarks will approximately coincide (because the 4FGL-DR2 cut PSs happen to have a similar total flux to the GCE). In particular, the $N_r = 106, R_r = 40\%$ benchmark is similar to a scenario where the resolved GCE sources coincide with all unflagged 4FGL-DR2 sources passing the cuts, and the $N_r = 53, R_r = 20\%$ benchmark is similar to the case where the resolved GCE sources coincide with the wavelet-selected sources.

There is an independent set of constraints on R_r from the observation that masking all known PSs does not seem to appreciably reduce the flux of the GCE. In particular, Ref. [60] tested the effects of masking all 4FGL-DR1 sources and found that the effect on the inferred GCE spectrum was negligible (compared to masking only 2FGL sources), reducing the GCE flux by less than 10% at all energies. The recent study in Ref. [25], which masks all 4FGL-DR2 sources, infers a very comparable GCE flux to earlier similar studies that masked only the 3FGL [6] or 2FGL [16], with estimated error bars on the total flux at the 10-20% level (although, as discussed above, the scatter between the full range of analyses is larger). The substantial systematic uncertainties in the determination of the GCE spectrum make it difficult to claim a statistically precise quantitative exclusion, but a contribution to the total GCE flux of more than $R_r \sim 20\%$ from resolved sources would appear to be in tension with the results of Ref. [60], and possible tension with the large GCE flux inferred by Ref. [25] in an analysis where all 4FGL-DR2 sources were masked. Thus our $R_r = 18\%$ benchmark can also be viewed as an approximate upper bound on the fraction of flux in resolved sources in order to avoid tension with observations, and the wavelet-selected source population as an example of what saturating that bound




Figure 2-5: Comparison between the gamma-ray flux distributions of sources in the 4FGL-DR2 (blue) and 4FGL-DR1 (orange) catalogs, and the wavelet-selected 4FGL-DR1 sources of Ref. [60], within our ROI and after the cuts are applied (see text for details). Fluxes are obtained from the relevant catalogs and are evaluated for the 0.1-100 GeV energy band. We show results including only flagged catalog sources (solid blue/orange lines), flagged+unflagged catalog sources (shaded blue/orange regions), and wavelet-selected sources (solid black lines). The vertical line denoting F_{th} represents the weighted average flux threshold of the 4FGL-DR2 catalog within our ROI, as computed in appendix F.

might look like.

Total number of MSPs

We will also discuss another feature of a potential MSP population in the GC: the total number of MSPs, resolved or unresolved, denoted N_{GCE} . Given a prediction for N_{GCE} , we can ask whether N_{GCE} can reasonably be achieved in our Galaxy by existing models for the formation and evolution of MSPs. For example, it might be possible to exclude certain GCE luminosity functions on the grounds that they predict unphysically large numbers of very faint MSPs, even if those MSPs are not plausibly resolvable. As an example, the power-law luminosity function proposed in Ref. [60] produces more than three million MSPs, which is very large compared to earlier predictions from population synthesis studies. For example, Ref. [27] predicted a total number of radio-loud pulsars (of all periods) in the Galaxy of around 1.2 million, with only $\mathcal{O}(10\%)$ of those beamed toward us; Ref. [55] predicted a birth rate of MSPs in the Galactic disk of a few $\times 10^{-4}$ per century, translating to a few

$\times 10^4$ MSPs produced over the age of the Galaxy; Ref. [30] proposed that globular cluster disruption could source a population of 5000-16,000 MSPs in the inner Galaxy; and Ref. [29] proposes $\mathcal{O}(10^5)$ MSPs could be produced in the Galactic bulge via accretion-induced collapse. Ref. [60] uses this discrepancy to argue that the MSP hypothesis may be under stress, although this argument depends strongly on the assumed luminosity function. However, if non-detection of GCE PSs continues with increasing sensitivity to PSs, such that all reasonable luminosity functions would predict a number of MSPs exceeding the yield of all possible production mechanisms, then in future it might in principle be possible to exclude the MSP hypothesis.

2.2.2 Sensitivity models

The *Fermi* telescope does not detect every pulsar (or PS) in the GC; position-dependent background emission obscures dimmer PSs, and the faintest PSs may not produce a statistically significant number of photons at all. Inclusion of a source in the 4FGL catalogs generally requires a minimum value of the test statistic (TS) describing the likelihood improvement from adding the source to the model ($TS > 25$ for the likelihood test with non-curved spectra).

In our main analysis, we account for these factors by using a position-dependent flux threshold $F_{\text{th}}(b, l)$ published by the *Fermi* team for the DR2 catalog [4, 9] to model the catalog’s threshold sensitivity. If a PS emits flux $F > F_{\text{th}}(b, l)$, we model it as resolved, and if $F < F_{\text{th}}(b, l)$, we model it as unresolved. The position dependence of this flux threshold, restricted to our ROI and for the energy range 0.1 – 100 GeV, is shown in figure 2-6. We will refer to this sensitivity model as the standard sensitivity model.

This approach is an approximation — in reality, Poisson fluctuations in the observed number of photons from sources or backgrounds may cause sources to move across the threshold in either direction, and the published sensitivity map also assumes a specific spectrum for the PSs — but we expect the resulting systematic errors to be small except perhaps for PS populations with flux distributions peaked very close to the threshold. We will validate this approach by comparing the flux




Figure 2-6: Flux threshold $F_{\text{th}}(\ell, b)$ to include a PS within our ROI in the 4FGL-DR2 catalog (left panel) [4, 9], and the ratio of the 4FGL DR-2 and 4FGL DR-1 flux thresholds (right panel). Fluxes are computed in the 0.1-100 GeV energy band.

distribution of observed sources to the predicted distribution for plausible luminosity functions.

To calculate the required properties of GC PS populations given a luminosity function $P(L) \propto dN/dL$, we write:

$$\begin{aligned}
 F_{\text{GCE}} &= \int_{\Omega} d\Omega \int_0^{\infty} s^2 ds A \rho_{\text{GCE}}(r) \int_{L_{\text{min}}}^{\infty} dL \frac{L}{4\pi s^2} P(L), \\
 N_{\text{GCE}} &= \int_{\Omega} d\Omega \int_0^{\infty} s^2 ds A \rho_{\text{GCE}}(r), \\
 F_r &= \int_{\Omega} d\Omega \int_0^{\infty} s^2 ds A \rho_{\text{GCE}}(r) \int_{4\pi s^2 F_{\text{th}}(\ell, b)}^{\infty} dL \frac{L}{4\pi s^2} P(L), \\
 N_r &= \int_{\Omega} d\Omega \int_0^{\infty} s^2 ds A \rho_{\text{GCE}}(r) \int_{4\pi s^2 F_{\text{th}}(\ell, b)}^{\infty} dL P(L).
 \end{aligned} \tag{2.3}$$

Here Ω denotes our $20^\circ \times 20^\circ$ ROI with $|b| < 2^\circ$ masked, and A is the coefficient of the RHS of eq. 2.1 — i.e. A governs the number of PSs — and is fixed by forcing F_{GCE} to equal the observed value. In eq. 2.3, r represents the distance to the GC from the point of integration and is determined by the law of cosines: $r^2 = s^2 + r_c^2 - 2r_c s \cos b \cos \ell$, where $r_c = 8.5$ kpc is the approximate distance from the Earth to the GC (as given in e.g. [11]) and s is the distance between Earth and the point of integration.

We contrast the standard sensitivity model with a simplified sensitivity model used by Ref. [60]. The simplified model assumes that all PSs with luminosity $L > L_{\text{th}}$ are resolved, and none with $L < L_{\text{th}}$. In Ref. [60] and here, L_{th} is fixed throughout the sky at $10^{34} \text{ erg s}^{-1}$. The four properties we intend to measure are then given by

$$\begin{aligned} L_{\text{GCE}} &= N_{\text{GCE}} \int_{L_{\text{min}}}^{\infty} LP(L)dL, & L_r &= N_{\text{GCE}} \int_{L_{\text{th}}}^{\infty} LP(L)dL, \\ N_r &= N_{\text{GCE}} \int_{L_{\text{th}}}^{\infty} P(L)dL, \end{aligned} \quad (2.4)$$

where N_{GCE} is a normalization constant, fixed by requiring that the luminosity of the GCE L_{GCE} reproduces the flux F_{GCE} observed. The conversion between L_{GCE} and F_{GCE} required to force this equivalence is outlined in appendix D. We call this sensitivity model the simplified model.

The threshold value L_{th} typically has a strong impact on the values given by eq. 2.4. A more detailed analysis of the validity and effect of the $10^{34} \text{ erg s}^{-1}$ estimate is carried out in appendix F.

The final sensitivity model, which we label the ‘‘smoothed model’’, has been used elsewhere in the literature to model the probability for a MSP to be detected as a pulsar, not simply as a gamma-ray PS. Given a sensitivity model of this type, we could confront our predicted N_r and R_r values with the observed N_r and R_r values from *detected pulsars* in the ROI relevant to the GCE, potentially allowing us to discard non-pulsar background sources and leading to significantly stronger constraints. We adopt the prescription of Ref. [53],

$$P_r(F) = \frac{1}{2} \left(1 + \text{erf} \left(\frac{\log_{10} F - (\log_{10} F_{\text{th}}(\ell, b) + K_{\text{th}})}{\sqrt{2}\sigma_{\text{th}}} \right) \right). \quad (2.5)$$

Here rather than having a step-function threshold, for each source we assign a probability of detection $P_r(F)$ as a function of its flux F and position. A constant parameter $10^{K_{\text{th}}} \text{ erg s}^{-1}$ is added to the baseline 4FGL-DR2 threshold $F_{\text{th}}(\ell, b)$ (as used in our standard sensitivity model) to accommodate a higher threshold for pulsar identification. The width of the threshold — accounting for uncertainties in K_{th} or the

published thresholds, statistical variations of photon counts, or temporal variation in background radiation — is controlled by a dimensionless parameter σ_{th} . Ref. [53] extracted values for K_{th} and σ_{th} from a Markov chain Monte Carlo fit to globular cluster MSPs, finding $K_{\text{th}} = 0.45$ and $\sigma_{\text{th}} = 0.28$; while the inner Galaxy region is quite a different environment from globular clusters, we adopt these values in this work for illustration.

In this framework, the observables F_r and N_r are calculated simply by multiplying the integrand of the luminosity integral in eq. 2.3 by $P_r(L/4\pi s^2)$ for F_r and N_r and removing the lower limit of integration. (Note that there is no need to apply this modification to F_{GCE} and N_{GCE} , because we do not require that the N_{GCE} pulsars making up the total F_{GCE} flux are resolved.)

2.3 Luminosity functions

2.3.1 General parameterizations

Luminosity functions for MSP populations are frequently parameterized either as a power-law with a cutoff or break, or as a log-normal function with a broad peak around some characteristic luminosity. In this work we study each of these cases in general, to understand what regions of parameter space are still permitted to explain the excess and what they predict for resolved and unresolved sources, and compare with specific benchmark points given in the literature.

Our first parameterization is a power-law with an exponential cutoff at high flux and a step-function cutoff at low flux (as used e.g. in Refs. [60, 10]), shown here in its normalized form:

$$\frac{dN}{dL} \propto P_{\text{PL}}(L) = L^{-\alpha} \exp\left(-\frac{L}{L_{\text{max}}}\right) \left[\Gamma\left(1 - \alpha, \frac{L_{\text{min}}}{L_{\text{max}}}\right) L_{\text{max}}^{1-\alpha}\right]^{-1}. \quad (2.6)$$

This luminosity function restricts the range of luminosities to $[L_{\text{min}}, \infty)$, where L_{min} , L_{max} , and α are free parameters. It is in practice quite similar to a power-law with a step-function cutoff or sharp break around L_{max} ; the details of the cutoff at L_{min} are

not generally observable, since they describe the behavior of very faint point sources.

Our second parameterization is a log-normal luminosity function (as used in e.g. Ref. [37]):

$$\frac{dN}{dL} \propto P_{\text{LN}}(L) = \frac{\log_{10} e}{\sigma\sqrt{2\pi}L} \exp\left(-\frac{(\log_{10} L - \log_{10} L_0)^2}{2\sigma^2}\right), \quad (2.7)$$

where L_0 and σ are free parameters. Here, all values of $L > 0$ are allowed.

We also study a broken power-law luminosity function:

$$\frac{dN}{dL} \propto P_{\text{BPL}}(L) = \left(\frac{(1-n_1)(1-n_2)}{L_b(n_1-n_2)}\right) \begin{cases} (L/L_b)^{-n_1} & L < L_b \\ (L/L_b)^{-n_2} & L > L_b \end{cases}, \quad (2.8)$$

with parameters L_b , n_1 , and n_2 . We use this luminosity function to benchmark our results, but do not scan the parameter space as we do for the power-law and log-normal functions.

2.3.2 Benchmark configurations

To better interpret these functional forms, we further study seven benchmark configurations which have been proposed in the literature.

We use two benchmarks for the power-law luminosity function from Refs. [60] and [11]. Both sources perform a wavelet search on GCE data, isolate point sources in the GCE, and test what properties are required for a power-law luminosity function to match these observations.

Ref. [60] found that $L_{\text{min}} = 10^{29} \text{ erg s}^{-1}$, $L_{\text{max}} = 10^{35} \text{ erg s}^{-1}$, and $\alpha = 1.94$ exactly reproduces the observed number of resolved wavelet-selected PSs $N_r = 47$ and the corresponding flux fraction R_r , while producing the full GCE flux. (As noted above, this prediction employs the simplified sensitivity model and a lower GCE flux than our baseline analysis.) They find that this model predicts approximately three million MSPs (almost entirely unresolved) in the GCE. Although this luminosity function matches Ref. [60]’s observations of N_r and R_r , it is unusual in that it has a very steep

slope / large spectral index α with a low L_{\min} , corresponding to an average luminosity $1.9 \times 10^{30} \text{ erg s}^{-1}$ that is much lower than the other six benchmarks discussed here. We label this benchmark as “Wavelet 1”.

Ref. [11] included in their analysis 13 resolved PSs in the 3FGL catalog, together with a population of wavelet peaks not associated with 3FGL sources. They assumed a power-law slope of 1.5 and a step-function cutoff at $L = L_{\max}$. Fitting to this population (assuming all wavelet peaks were associated with the GCE) yielded a best-fit value for the cutoff of $L_{\max} \approx 7 \times 10^{34} \text{ erg s}^{-1}$. We adopt these parameters, but with an exponential cutoff at the same L_{\max} . This configuration has average luminosity $7.4 \times 10^{31} \text{ erg s}^{-1}$, and we label this benchmark “Wavelet 2”.

We benchmark the log-normal luminosity function with three configurations found in the literature. Ref. [37] fits a log-normal model to data from globular cluster MSP populations, yielding values $L_0 = 8.8_{-4.1}^{+7.9} \times 10^{33} \text{ erg s}^{-1}$ and $\sigma = 0.62_{-0.16}^{+0.15}$. We label this luminosity function the “GLC” benchmark, since it is based on observations of globular clusters.

Ref. [53] proposes a more intricate luminosity function, derived from a model of the relationship between MSP spectrum and physical properties. They split the MSP population into three components: a Disk component containing most of the Galactic resolved MSPs, and a Boxy Bulge and Nuclear Bulge component which both contained unresolved MSPs and make up the GCE. A luminosity function is generated for each. We use the Boxy Bulge luminosity function, since the Boxy Bulge extends out to kpc scales whereas the Nuclear Bulge is smaller in radial extent and lies mostly within our mask; that said, the luminosity functions are very similar. While this luminosity function is generated numerically, it very closely resembles a log-normal curve with $L_0 = 1.3 \times 10^{32} \text{ erg s}^{-1}$ and $\sigma = 0.70$ (these values were obtained from a simple least-squares fit, detailed in appendix C). We label this luminosity function the “GCE” benchmark, as it is fitted in part based on observations of the Bulge and is intended to fit the GCE.

Another recent study analyzes the possibility that the hypothetical MSP population in the GCE is generated not via the spin-up of old pulsars in low-mass X-ray

binary progenitor systems, but directly via accretion-induced collapse (AIC) [29]. The population synthesis model developed in this work yields a numerical luminosity function for MSPs in the GC, reported as a function of flux. It predicts hundreds of thousands of low-flux MSPs in the GC with very few being bright enough for detection. For ease of comparison to the other luminosity functions discussed here, we convert the reported flux distribution to a luminosity distribution as described in appendix D, and show in appendix C that this distribution closely resembles a log-normal function with $L_0 = (4.3 \pm 0.2) \times 10^{30} \text{ erg s}^{-1}$ and $\sigma = 0.94$. These parameters were obtained via a least squares fit in appendix C. We label this luminosity function the “AIC” benchmark, as it employs an AIC-based population synthesis model.

Our first broken power-law benchmark is derived from Ref. [10], which uses a self-consistent Bayesian analysis to fit a luminosity function, spatial distribution, and other properties to observed MSPs in the Galactic Disk. That work derives parameters of $n_1 = 0.97$, $n_2 = 2.60$, and $L_b = 1.7 \times 10^{33} \text{ erg s}^{-1}$, additionally impose upper and lower limits on the luminosity of MSPs of $10^{37} \text{ erg s}^{-1}$ and $10^{30} \text{ erg s}^{-1}$ respectively. We make use of these limits as well by modifying the normalization of eq. 2.8. We call this benchmark the Disk benchmark.

Our second broken power-law benchmark is derived from Ref. [44]’s search for point sources in the GCE via Non-Poissonian Template Fitting (NPTF). In this paper the flux distribution of sources was modeled as a broken power-law, with parameters fitted to the data assuming a gNFW-squared-distributed population of MSPs (in addition to separate extragalactic and disk PS populations with different luminosity functions), and data restricted to a $\sim 2 - 12 \text{ GeV}$ energy range. We discuss the conversion to our energy range and to luminosity rather than flux in appendices D and E. The resulting best-fit parameters are $n_1 = -0.66$, $n_2 = 18.2$, and $L_b = 2.5 \times 10^{34} \text{ erg s}^{-1}$; due to the size of n_2 , this luminosity function behaves similarly to a power-law with a step-function cutoff at L_b , and predicts that most of the power in the GCE arises from MSPs with flux comparable to *Fermi*’s sensitivity. We call this luminosity function the “NPTF” benchmark.

All the above-mentioned luminosity functions are shown in figure 2-7; we plot

| Name | Ref. | L_{\min} | L_{\max} | α | Fit? | Pulsars modeled? |
|-----------|------|----------------------|--------------------|----------|------|------------------|
| Wavelet 1 | [60] | 10^{29} | 10^{35} | 1.94 | N | N |
| Wavelet 2 | [11] | 10^{29} | 7×10^{34} | 1.5 | N | N |
| | | L_0 | σ | | | |
| GLC | [37] | 8.8×10^{33} | 0.62 | | N | Y |
| GCE | [53] | 1.3×10^{32} | 0.70 | | Y | Y |
| AIC | [29] | 4.3×10^{30} | 0.94 | | Y | Y |
| | | L_b | n_1 | n_2 | | |
| NPTF | [44] | 1.7×10^{33} | 0.97 | 2.60 | N | N |
| Disk | [10] | 2.5×10^{34} | -0.66 | 18.2 | N | N |

Table 2.1: Names, references, and parameters for all seven luminosity functions studied in this paper. All luminosities are given in units of erg s^{-1} . The “Fit?” column contains “Y” when we obtained the given parameters via a least-squares fit to numerical data provided in the reference, and “Pulsars modeled?” contains “Y” when the luminosity function explicitly relies on a pulsar model in their analysis.

dN/dL , $dN/d\ln(L) = LdN/dL$ (which shows how the total number of pulsars is distributed across e-folds in L), and $LdN/d\ln(L) = L^2dN/dL$ (which shows how the total luminosity is distributed across e-folds in L) for clarity. We note that despite the different parameterizations, several of these benchmarks are roughly similar in shape; the main exceptions are the GLC and NPTF luminosity functions, which have much more power near the high-luminosity cutoff than others, and the “Wavelet 1” benchmark, which is flatter than the others. The names, associated references, and parameter values of all these luminosity functions are repeated in table 2.1.

It is important to note that all these luminosity functions are peaked at flux values below the *Fermi* sensitivity threshold around a few times $10^{34} \text{ erg s}^{-1}$ — well below, in most cases. The high luminosity tail therefore dominates the resolvable MSPs and hence our observables. Our analysis is thus primarily sensitive to the luminosity function parameters insofar as they affect this tail, and to the sensitivity model, which dictates how much of the tail is exposed to *Fermi*.




Figure 2-7: *Left:* Seven benchmark luminosity functions for GCE MSP luminosity functions, normalized to give the luminosity probability density as in eqs. 2.6 to 2.8, for (left panel) dN/dL , (right panel) LdN/dL , and (bottom panel) L^2dN/dL . See table 2.1 for a summary of these luminosity functions.

Chapter 3

Results

3.1 Allowed luminosity function configurations

Using each of the luminosity function parameterizations and benchmarks outlined in section 2.3, we extract the total expected number of resolved MSPs N_r , the ratio of the flux received from those resolved MSPs to the total flux of the GCE R_r , and the total number of MSPs N_{GCE} by forcing each function to reproduce the observed flux of the GCE as discussed in section 2.1.2. This is done for the parameter space of the power-law and log-normal luminosity functions (Eqs. 2.6 and 2.7), with N_{GCE} displayed as a contour map in figures 3-1a and 3-1b, using the standard and the simplified sensitivity models respectively.

The first two rows of these figures explore the parameter space for the power-law luminosity functions with a cutoff at L_{max} . In the first row we hold L_{min} fixed at $L_{\text{min}} = 10^{29} \text{ erg s}^{-1}$, and vary L_{max} and α . In the second row we instead hold α fixed at the value corresponding to the Wavelet 1 benchmark ($\alpha = 1.94$), to explore the effects of varying L_{min} . The third row explores the parameter space for log-normal luminosity functions, in terms of L_0 and σ .

The observational constraints of N_r constant and R_r constant each trace out a one-parameter family of luminosity functions that are also displayed for various reference values. For N_r , we display $f = 100\%$, 40% , 20% , 10% , and 5% , where f is the fraction of the DR2 catalog modeled as part of the GCE. For R_r , we display the fraction of

GCE modeled as produced by resolved sources, with reference values $R_r = 40\%$, 20% , 10% , and 5% . The regions with even less N_r and R_r are marked with + symbols to denote that they are still consistent with observations. Luminosity functions slightly above the observational constraints may still be marginally allowed, given both the Poisson error bar in N_r and systematic uncertainties arising e.g. from the choice of GCE flux. The observational curves are intended to indicate the regions of parameter space where tension with observational constraints starts to become a concern.

In figure 3-2, we display a “cross section” of figure 3-1. That is, we display several luminosity functions on the line $L_{\min} = 10^{31} \text{ erg s}^{-1}$, $\alpha = 1.94$ for the power-law case ($\sigma = 0.5$ for the log-normal case.) of figure 3-1 and shade the segment of the luminosity function brighter than the threshold. This shaded area is proportional to N_r . The threshold was chosen to be the average threshold determined by appendix F, so that figure 3-2 is comparable to the N_r curves shown in figure 3-1a (the standard sensitivity model). These figures more clearly illustrate the regions of figure 3-1a which maximize N_r for fixed L_{\min} or fixed σ .

Figure 3-2 more clearly demonstrates why maxima in N_r sometimes occur for the standard sensitivity model when L_{\min} and α (power-law case) or σ (log-normal case) are fixed. These maxima are visible in figure 3-1a. For low L_{\max} or L_0 (bottom of figure 3-2), N_r is low simply because the high-luminosity tail is too small. For large L_{\max} or L_0 (top of figure 3-2), N_r is low because so much flux is produced by the bright luminosity function that it has to be normalized to produce fewer pulsars in order to reproduce the observed GCE flux. In figure 3-2, the renormalization is visible as a decrease in height of the luminosity function peak, and the maximum in N_r is shown by the fact that the area of the colored region is not greatest at the top or bottom of the plots but somewhere in the middle. These effects are most visible in the log-normal case.

Relevant luminosity function benchmarks are also marked in figure 3-1; only the Wavelet 1 benchmark is shown in the second row because the choice of α in the second row is inconsistent with other benchmarks. The NPTF benchmark is not shown in the first row of power-law plots because its best-fit value for α is well off the left-hand

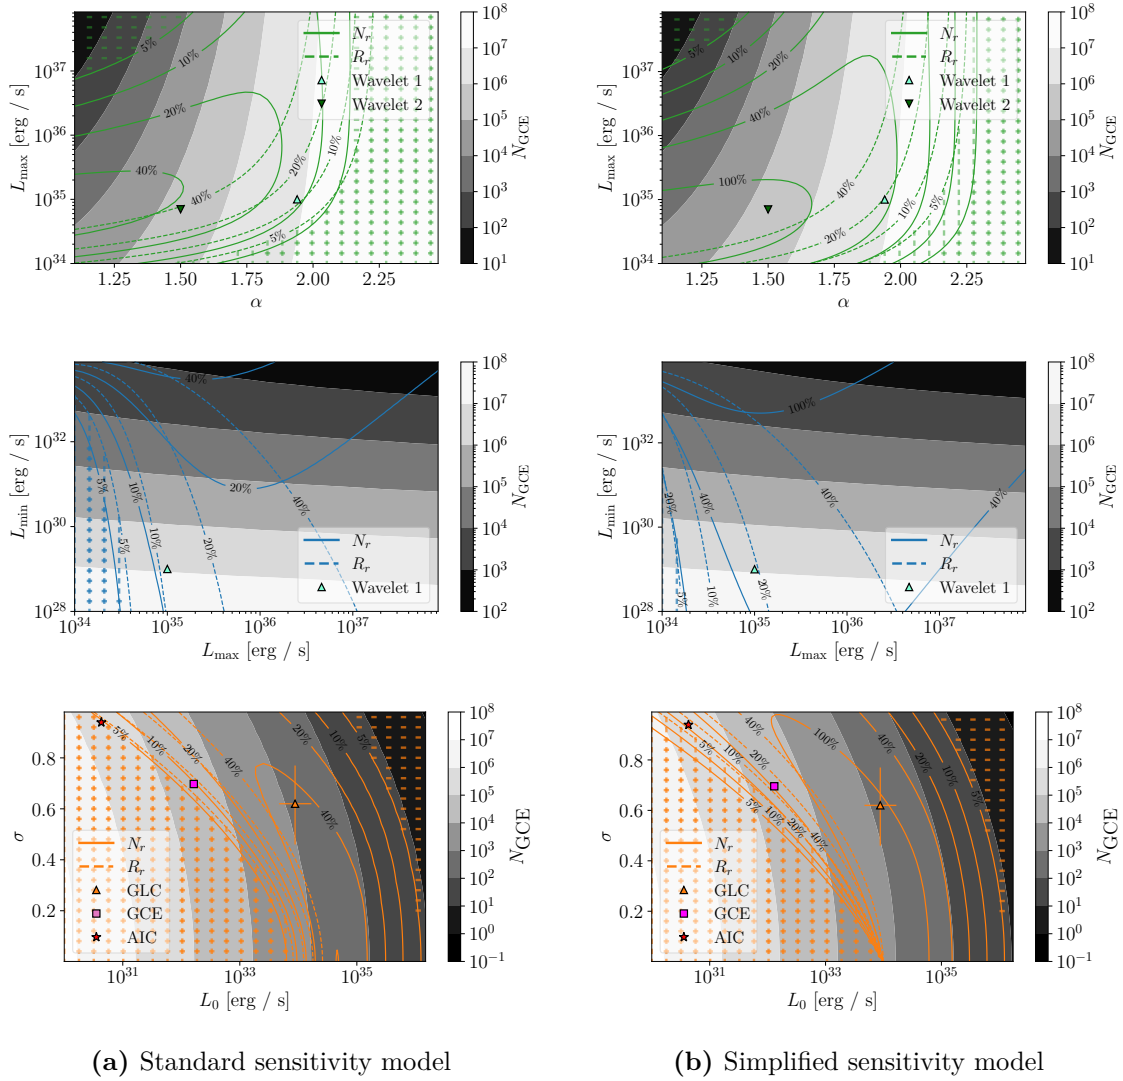


Figure 3-1: N_r , R_r , and N_{GCE} values for power-law (top two rows) or log-normal (bottom row) luminosity functions, normalized to explain the full GCE flux as determined by Ref. [25]. The benchmark configurations are marked (see text for details). The N_r values (solid) are marked as fractions of the 4FGL-DR2 PSs passing our cuts (see text); the R_r values (dashed) are marked as fractions of the total GCE flux. These fractions are labeled on the contours. The regions where N_r (R_r) is less than 5% of the number of observed resolved sources (GCE flux) are marked with | (-), while regions where both numbers are below this threshold are marked with +. The gray-scale contours indicate the total number of sources required to explain the GCE. The top row (green) fixes $L_{\text{min}} = 10^{29}$ erg/s while the second row (blue) fixes $\alpha = 1.94$.




Figure 3-2: Luminosity functions with fixed $L_{\min} = 10^{31} \text{ erg s}^{-1}$ and $\alpha = 1.94$ for the power-law case, (*left*), and fixed $\sigma = 0.5$ for the log-normal case (*right*). L_{\max} and L_0 respectively are allowed to vary, with large values at the top. Luminosities greater than the threshold luminosity (dotted line) determined in appendix F are shaded, and the shaded area is proportional to N_r . The top and bottom panels of each figure show the highest and lowest luminosity functions explicitly, so that the vertical axis is labeled.

side of the plot, near the bottom of the figure. It predicts observations of $N_r = 111$ and $R_r = 0.26$ via the standard sensitivity model (and $N_r = 770$ and $R_r = 0.93$ for the simplified sensitivity model). The expected values of N_r , R_r , and N_{GCE} for all benchmark configurations are given in table 3.1 for the standard sensitivity model and the alternatives.

Focusing on the standard sensitivity model, we observe that in the power-law case, requiring that resolved sources make up less than $\sim 10 - 20\%$ of the GCE would drive us either to large values of $\alpha \gtrsim 2$ or values of $L_{\text{max}} \lesssim 10^{35}$ erg/s. In other words, if L_{max} is substantially larger than the typical sensitivity threshold, a steep power-law slope is required to avoid overproduction of bright sources, and furthermore this implies a very large number of point sources ($\gtrsim 10^7$). This case corresponds to the solution identified in Ref. [60]. If we only restricted the *number* of bright sources and not their flux, there would be another possible solution where a small number of very bright sources explain the whole GCE, but this is observationally excluded (as masking the bright sources does not remove the GCE). Increasing the sensitivity (as shown in the right-hand panels) pushes the threshold down and hence reduces the value of L_{max} compatible with a wide range of choices for the power-law index, where the number of required MSPs is modest. In the second row of figures, we see that when α is held fixed at 1.94, demanding that the resolved sources make up less than $\sim 10 - 20\%$ of the GCE would primarily constrain L_{max} (to a value between $\sim 10^{34} - 2 \times 10^{35}$ erg/s), whereas the total number of sources is primarily controlled by L_{min} ; if a larger flux fraction in resolved sources is permitted, the limit on L_{max} will similarly rise, but without meaningfully constraining L_{min} and the total number of required sources.

Figure 3-1b shows that the Wavelet 1 power-law benchmark used by Ref. [60], which matches the properties of the wavelet-selected subsample exactly when that reference’s GCE luminosity and the simplified sensitivity model are used ($N_r = 47$ and $R_r = 0.2$), fails to match these properties in our analysis when the simplified sensitivity model is used. The reason is the larger GCE flux value we have taken from Ref. [25]. This flux is somewhat more than twice that extracted from Ref. [60],

and therefore more than twice as many MSPs are required to reproduce it; with the simplified sensitivity model, this scenario would thus predict more than double the number of resolved sources that occur in the wavelet-selected subsample (while closely matching the *fraction* of the GCE flux attributed to that subsample), since it was designed to match that subsample with a smaller GCE flux. However, using the standard sensitivity model, we see that this benchmark point is actually quite consistent with the wavelet-associated subsample when our standard sensitivity model is applied (see table 3.1); the simplified sensitivity model used in Ref. [60] overestimates the number of detected sources by a larger factor than the effect of varying the flux.

We can ask where the point is that would correspond to matching both N_r and R_r for the wavelet-selected subsample, as an example of a possible resolved-source population yielding roughly 20% of the total GCE flux (and roughly 20% of the 4FGL-DR2 sources), if we use the standard sensitivity model and the flux from Ref. [25] (while keeping in mind that there are substantial systematic uncertainties that can modify this value). Figure 3-1a indicates that in this case the configuration that matches both N_r and R_r values moves to $\alpha \approx 1.8$ and $L_{\max} \approx 9 \times 10^{34} \text{ erg s}^{-1}$ if L_{\min} is fixed, or to $L_{\max} \approx 1 \times 10^{35} \text{ erg s}^{-1}$ and $L_{\min} \approx 8 \times 10^{31} \text{ erg s}^{-1}$ if α is fixed instead.

These configurations possess average luminosities of $\sim 5 \times 10^{30} \text{ erg s}^{-1}$ and $3 \times 10^{33} \text{ erg s}^{-1}$ respectively, compared with $1.9 \times 10^{30} \text{ erg s}^{-1}$ of the original benchmark. They also predict millions and hundreds of thousands of MSPs respectively, compared to 8.5 million for the original benchmark as computed by our analysis. (Recall that due to different modeling choices, Ref. [60] found a slightly smaller value.) The change in average luminosity and total number of MSPs demonstrates the strong dependence of the power-law configuration satisfying observational constraints, and the associated total number of sources, on the sensitivity model used (as well as the assumed GCE flux); the large number of MSPs inferred by Ref. [60] is not a requirement, even when a power-law luminosity function model is imposed.

As shown in table 3.1, Wavelet 2 and NPTF benchmarks both require high N_r values — over 100 resolved sources — to reproduce the excess when using our standard sensitivity model; in other words, if these scenarios were realized, more than 40% of all

4FGL-DR2 sources (including flagged sources) should be real and associated with the GCE source population. These scenarios also predict that a relatively large fraction of the flux of the excess should originate from resolved sources (38% for Wavelet 2 and 26% for NPTF), which already appears to be in some tension with observations, and this particular tension is not dependent on the chosen GCE flux value. The smoothed sensitivity model (used to estimate the detection threshold for MSPs elsewhere in the sky) predicts tens of detected inner-Galaxy MSPs for both of these scenarios. This may also be in conflict with measurements, since the 4FGL-DR2 contains 26 known MSPs in the ROI and only 2 are known to be within 2 kpc of the GC, with 5 not in the ATNF catalog and therefore having unknown position.

A dedicated analysis of the effects of masking 4FGL-DR2 sources, the sensitivity to pulsars in this population, and/or the spatial distribution of these sources, could potentially sharpen these statements. Note that while both of these benchmarks were derived from examination of older inner Galaxy gamma-ray data, it is still quite reasonable for them to be in tension with the data now that the sensitivity of the catalogs have improved; the NPTF analysis may also have overestimated near-threshold PSs due to systematics from background mismodeling [40, 42, 41].

For the log-normal luminosity functions, we see that there are two broad regions of parameter space consistent with a small value of N_r , but one of these regions — where L_0 is very large and so the number of sources required to explain the GCE is small — is excluded by any plausible limit on R_r (this again corresponds to the scenario where a handful of bright resolved sources explain the GCE; the fact that masking 4FGL sources does not eliminate the GCE excludes this region). In the remaining region of parameter space, the N_r and R_r constraint lines are quite similar, suggesting that log-normal luminosity functions generating a fraction close to R_r of the GCE will also quite generically predict a number of observed sources comparable to N_r (and vice versa). However, the constraint lines spans several orders of magnitude in the total number of allowed sources, from around 10^3 up to 10^5 or more.

The GCE and AIC benchmarks correspond to $N_r = 5.5$ and 3.5 respectively by our smoothed analysis (table 3.1a), compared with the three pulsars originally

| Luminosity function | N_r | R_r | N_{GCE} |
|---------------------|-------|-------|-------------------|
| Wavelet 1 | 31 | 0.11 | 8.5×10^6 |
| Wavelet 2 | 98 | 0.38 | 2.2×10^5 |
| GLC | 124 | 0.72 | 670 |
| GCE | 20 | 0.059 | 3.5×10^4 |
| AIC | 12 | 0.039 | 3.6×10^5 |
| NPTF | 111 | 0.26 | 970 |
| Disk | 30 | 0.13 | 2.6×10^4 |

(a) Standard sensitivity model

| | | | |
|-----------|-----|-------|-------------------|
| Wavelet 1 | 120 | 0.19 | 8.5×10^6 |
| Wavelet 2 | 320 | 0.59 | 2.2×10^5 |
| GLC | 310 | 0.91 | 660 |
| GCE | 120 | 0.13 | 3.4×10^4 |
| AIC | 61 | 0.078 | 3.6×10^5 |
| NPTF | 770 | 0.93 | 960 |
| Disk | 140 | 0.22 | 2.5×10^4 |

(b) Simplified sensitivity model

| | | | |
|-----------|-----|-------|-------------------|
| Wavelet 1 | 9.6 | 0.053 | 8.5×10^6 |
| Wavelet 2 | 32 | 0.18 | 2.2×10^5 |
| GLC | 50 | 0.45 | 670 |
| GCE | 5.5 | 0.024 | 3.5×10^4 |
| AIC | 3.5 | 0.018 | 3.6×10^5 |
| NPTF | 25 | 0.077 | 970 |
| Disk | 10 | 0.078 | 2.6×10^4 |

(c) Smoothed sensitivity model

Table 3.1: Number of resolved PSs, ratio of resolved flux to total flux, and total number of PSs predicted to make up the GCE based on seven luminosity function benchmarks (see section 2.3) and the requirement that the PSs reproduce the entire flux of the GCE, for three different sensitivity models (see section 2.2.2).

identified as possibly bulge-population members for the GCE benchmark in Ref. [53] and the $0.15_{-0.0}^{+0.37}$ members originally predicted to be resolved for the AIC benchmark in Ref. [29]. We use the smoothed analysis for comparison purposes here because the sensitivity models used in Refs. [53, 29] are most similar to our smoothed model in that they model the resolution of PSs *as MSPs* rather than as PSs alone. Our results for N_r are larger for several reasons. Firstly, the original luminosity functions of both the GCE and AIC benchmarks were numerical, and we fitted a log-normal distribution as an approximation and used this fit to extrapolate to higher luminosities. For the AIC case, the original numerical result has a sharp cutoff at some maximum luminosity; if we cut our fitted luminosity function at the same point then our predicted N_r falls to less than one. Furthermore, slight differences between our smoothed sensitivity model and those used by Refs. [53, 29] also affect N_r estimates. In the GCE case, our predicted N_r does not exactly describe the expected number of pulsars that are found to have a high probability of association with the bulge population (which is what Ref. [53] measures); the latter number also depends on the distinguishability of bulge pulsars from other pulsars. Finally, our relatively large estimate for the flux of the GCE may also be responsible for our larger estimates of N_r .

When we consider specific benchmarks studied in the literature, figure 3-1 demonstrates that the GLC log-normal benchmark predicts a GCE dominated by resolved sources ($R_r = 0.72$), which is quite inconsistent with observations; this is consistent with Ref. [37], which proposed this luminosity function and argued it could not explain the GCE. On the other hand, the GCE and AIC benchmarks show no tension with observations, predicting a modest but not dominant contribution to the resolved sources that pass all our cuts (10-20 sources, yielding 3-5% of the GCE flux). Note that the predictions from these benchmarks for resolved sources are rather similar (differing by a factor of about 1.5); the main difference in these luminosity functions occurs at the low-luminosity end and is responsible for a roughly one-order-of-magnitude difference in the required number of sources (3.5×10^4 for the GCE benchmark and 3.6×10^5 for the AIC benchmark).

Because the smoothed sensitivity model has an overall offset in the detection

threshold (controlled by the parameter K_{th}), we expect it to yield lower N_r and R_r values, and identical N_{GCE} , since the sensitivity model does not affect the total number of pulsars. From table 3.1c we see that values of N_r fall by a factor of 2 – 3 in most cases compared to the standard sensitivity model, and R_r falls by a factor of around two, demonstrating that all seven configurations studied are roughly equally responsive to the difference between the smoothed and standard sensitivity models.

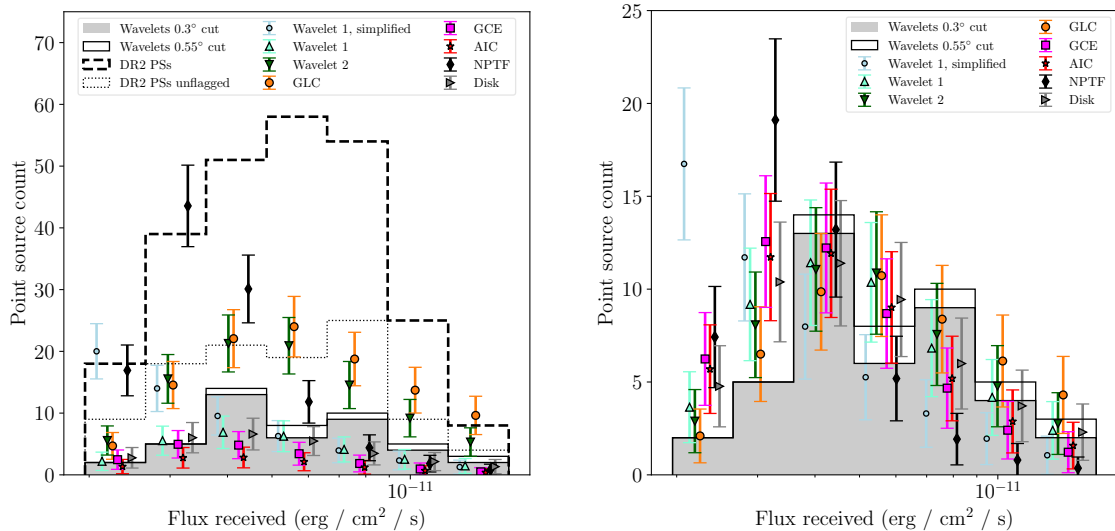
3.2 Number of MSPs in the GCE

For the power-law luminosity function, Ref. [60] determined that about three million MSPs are required to reproduce the GCE using a simplified sensitivity model. Using the total GCE flux inferred from Ref. [25], which is 2–3 times higher, but the same parameters, that number inflates to 8.5 million. As discussed above, power-law configurations that match the wavelet-selected subsample of Ref. [60] can be consistent with only $\mathcal{O}(10^4)$ MSPs if we use the flux value of Ref. [25] and the standard sensitivity model (figure 3-1a), largely because L_{min} is not observationally well-constrained but has a large impact on the total required number of sources.

For log-normal luminosity functions, figure 3-1 demonstrates that N_{GCE} is much more sensitive to L_0 than to σ for $\sigma < 1$. This is intuitively sensible, as to the degree that L_0 determines the average luminosity and hence the average flux of sources, it should simply be inversely related to the required number of sources. We can also read off from the figure that $N_{\text{GCE}} < 1000$ is generally inconsistent with requiring $R_r \lesssim 20\%$, and $N_{\text{GCE}} \ll 10,000$ is only achievable with low σ .

3.3 Flux distribution

Until this point, we have compared predictions for the population of MSPs in the inner Galaxy to observations only via the number and flux of resolved PSs. We may expand this analysis by comparing the predicted flux distribution of resolved MSPs to the observed distribution for resolved PSs in the inner Galaxy and the 0.1-100



(a) Each luminosity function is forced to reproduce the flux of the GCE. (b) Each luminosity function is forced to reproduce 47 PSs, which is the observed number of sources in the wavelet-selected subsample.

Figure 3-3: Distribution of resolved PSs predicted by each luminosity function with the standard sensitivity model, compared to the observed distribution of sources passing our cuts in the 4FGL-DR2 catalog (left panel, black dashed line), in the 4FGL-DR2 catalog with flagged sources removed (left panel, black dotted line), and in the wavelet-selected subsample (both panels) with a proximity cut of 0.3° (gray histogram) and 0.55° (hollow histogram). The predicted resolved PS distribution is also shown for the Wavelet 1 benchmark using the simplified sensitivity model (light blue points). Error bars are set to \sqrt{N} , and may therefore be inaccurate for low-count bins.

GeV energy band (as described in figure 2-5). This will also serve as a cross-check on our sensitivity threshold estimate, as we should see both the observed and predicted sources fall off at low fluxes.

The left panel of figure 3-3 shows the observed flux distributions for 4FGL-DR2 sources (with flagged sources included/excluded) and the wavelet-selected subsample of 47 sources (or 41 with a more stringent proximity requirement), together with the predicted distributions for the seven benchmark luminosity functions, normalized to explain the total GCE flux. The right panel shows the wavelet-selected subsample of observed sources, and the predicted source distributions based on the benchmark luminosity functions, normalized to match the total number of sources (47) in this population. In some cases this latter normalization would require overproducing the full GCE; the purpose is to test the similarity in the distributions of resolved sources

between models, when they are forced to have a common normalization. Both panels are calculated with the standard sensitivity model; for comparison, we also display the results of using the Wavelet 1 luminosity function with the simplified sensitivity model, as in Ref. [60].

From the left panel, we observe that binning in flux clarifies the degree of consistency between various luminosity function models and the data. Consistency of the NPTF model with the data would require essentially all 4FGL-DR2 sources (including flagged sources) with a flux below $4 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$ to be attributed to the GCE population; the tension arises from the lowest-flux bins. In contrast, the GLC model consistently attributes an appreciable fraction of the 4FGL-DR2 sources at all flux levels to the GCE, including slightly overpredicting the total sources in the highest-flux bin; the Wavelet 2 model has very similar behavior to the GLC model, but with a slightly lower normalization at high fluxes. As expected from the total flux and source counts calculated earlier, the Wavelet 1, GCE, AIC and Disk models all predict quite low numbers of counts and will be difficult to exclude observationally with current data. Finally, the simplified sensitivity model clearly does not capture the observed drop-off in the number of sources at low fluxes due to decreased sensitivity, signaling that the threshold is underestimated.

This last point is even more apparent in the right panel, where the observed wavelet-selected population has a decline at fluxes below $\sim 3 - 4 \times 10^{-12} \text{ erg/cm}^2/\text{s}$ that is also observed in all the predictions using the standard sensitivity model; the example with the simplified sensitivity model matches the overall number of sources but overpredicts low-flux sources at the expense of high-flux ones. By eye, the distributions of resolved sources for all the luminosity function benchmarks with the standard sensitivity model are rather similar (once their normalization is fixed): one quantitative measure of their differences is the p -value associated with the fit of each benchmark resolved-PS flux distribution to the observed distribution of sources in the wavelet-selected subsample. Note that, without a guarantee that the wavelet-selected subsample is representative of GCE PSs, poor p -values cannot be taken to exclude a specific model. We use the subsample here only as a plausible example of such a pop-

| Luminosity function | Number fixed p -value |
|-----------------------------|-------------------------|
| Wavelet 1 | 0.64 |
| Wavelet 2 | 0.86 |
| GLC | 0.89 |
| GCE | 0.0080 |
| AIC | 0.046 |
| NPTF | 2.2×10^{-10} |
| Disk | 0.30 |
| Wavelet 1, simplified model | 7.3×10^{-8} |

Table 3.2: p -values comparing the predicted resolved flux distributions for each benchmark luminosity function to the observed distribution, where the overall normalization of the predicted distributions is chosen to be equal to the number of observed resolved PSs in the wavelet-selected subsample of 4FGL-DR2 sources.

ulation. However, low p -values do indicate strong differences between this example distribution and that predicted by the luminosity function benchmarks, suggesting that this luminosity function could potentially be distinguished from others with knowledge of the bright end of the GCE source population (at current sensitivity).

In table 3.2, we show the p -values associated with the fit of each benchmark resolved flux distribution to the observed distribution shown in figure 3-3b. Specifically, we let $\chi^2 = -2 \ln(\mathcal{L}/\mathcal{L}_0)$ where \mathcal{L} gives the probability of achieving the binned data given the predicted resolved flux distribution, defined via a multinomial distribution (since the total sources are fixed). \mathcal{L}_0 is the probability of achieving exactly the model expectation values in each bin based on the analytically extended multinomial-distributed probability distribution function. We then extract a p -value using a χ^2 distribution with degrees of freedom equal to the number of bins minus one.

We find that all of our re-scaled benchmark flux distributions, using the standard sensitivity model, are quite consistent in shape with the observed flux distribution of wavelet-selected PSs, except for the NPTF benchmark and (at much lower significance) the GCE benchmark. Even in the scenario where all these PSs, and only these PSs, were associated with the GCE (where we would have maximal signal to noise), it would therefore be challenging to differentiate luminosity functions similar to these benchmarks based on the observed flux distribution of the resolved sources. The NPTF is the one major exception: such a strongly peaked luminosity function

could potentially be clearly distinguished from other benchmarks. Of the benchmarks that do fit the data, those that fit best generally possess brighter peak luminosities because all the benchmarks we survey generally underproduce bright MSPs compared to the wavelet-selected population, except for the GLC benchmark (see figure 3-3b).

Chapter 4

Discussion

In the following two sections, we discuss the degree to which our results depend on systematic uncertainties such as the GCE flux, the sensitivity threshold, and the cuts we place on known and predicted PS to separate GCE PSs from background. We also discuss the ability for future increases in PS sensitivity to distinguish between luminosity functions.

4.1 Systematic uncertainties

4.1.1 Dependence on assumed GCE total flux

Figure 4-1 shows luminosity function configurations that follow the observational constraints at different GCE fluxes, demonstrating how sensitive the regions corresponding to specific N_r and R_r values are to the total flux. To make the plots easier to follow, we have fixed the baseline N_r and R_r contours at the values corresponding to the wavelet-selected subsample of PSs, $N_r = 47 \approx 18\%$ of the 4FGL-DR2 PSs, and $R_r \approx 0.17$. We then ask the question: suppose we wish to maintain the predicted number of sources and their fluxes, then how does the region of parameter space that does not overproduce those sources/fluxes evolve as F_{GCE} changes? In particular, if F_{GCE} doubles, R_r must halve, to maintain the total flux in resolved sources. The standard sensitivity model is used in all cases.

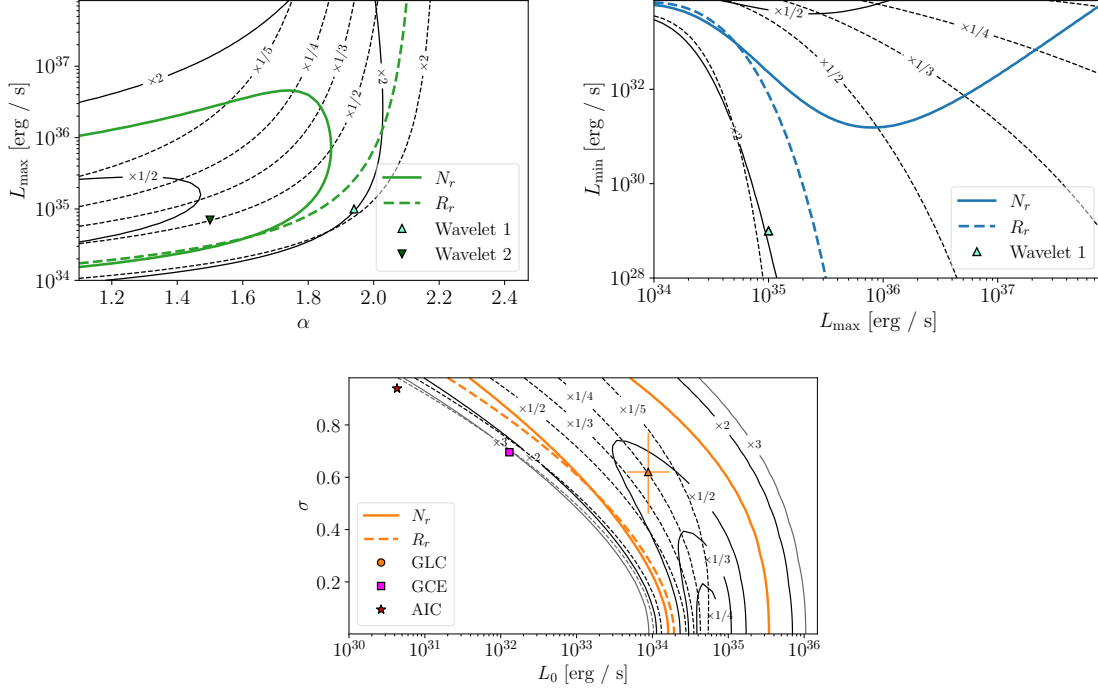


Figure 4-1: The parameter space that does not overproduce the wavelet-selected subsample of 4FGL-DR2 PSs ($N_r \leq 47$ and $R_r \leq 0.17$) for the power-law luminosity function (top two plots) and log-normal luminosity function (bottom plot). The total GCE flux has been increased by factors of 2–5 and decreased by factors of 2 and 3; contours of fixed N_r (solid) and fixed flux in resolved PSs (dashed) are labeled with the factor by which F_{GCE} has been multiplied. The colored lines exactly match the wavelet-selected subsample for the nominal GCE flux value (extracted from Ref. [25]). The left power-law plot fixes $\alpha = 1.94$ while the right fixes $L_{\min} = 10^{29} \text{ erg s}^{-1}$.

In all three parameter spaces investigated, the allowed region expands (contracts) when F_{GCE} is decreased (increased). The change is especially rapid in the power-law case. An increase by a factor of two in F_{GCE} will cause the Wavelet 1 benchmark to overpredict the wavelet-selected subsample. A decrease by a factor of two implies that $L_{\text{min}} \lesssim 10^{33} \text{ erg s}^{-1}$ satisfies $N_r \leq 47$ for a wide range of L_{max} values, so that nearly all luminosity functions produce fewer PSs than observed in the wavelet-selected subsample. The region allowed by the R_r constraint also changes rapidly; if L_{min} and α are fixed at the Wavelet 1 benchmark, then a reduction in F_{GCE} by a factor of two allows L_{max} to increase by a factor of ten and remain in the allowed region. A decrease in F_{GCE} by a factor of 2–3 would also bring the Wavelet 2 benchmark into consistency with the observables associated with the wavelet-selected subsample.

For the log-normal luminosity function, neither constraint is as sensitive to increases in F_{GCE} . The R_r constraint is not very sensitive to decreases either, but a decrease in F_{GCE} by a factor of 2–3 will cause the GLC benchmark configuration to satisfy the $N_r \leq 47$ condition to not overproduce the wavelet-selected subsample. Using another of the GCE spectra analyzed in section 2.1.2 could achieve this amount of decrease. However, the GLC configuration would still be disallowed by the R_r constraint unless F_{GCE} were decreased by a factor of $\gtrsim 5$.

4.1.2 Dependence on sensitivity threshold

Figure 3-1 and table 3.1 both demonstrate that the simplified and standard sensitivity models yield starkly different predictions for the observed PSs. The allowed parameter regions (for given choices of N_r and R_r) differ markedly, especially in the power-law case, between the two models, and predictions for N_r can differ by as much as a factor of ~ 7 for benchmarks. Predictions of R_r differ by as much as a factor of three. This is mostly because the simplified flux threshold $L_{\text{th}} = 10^{34} \text{ erg s}^{-1}$ used in Ref. [60] is low, as demonstrated by the fact that the lowest value in *Fermi*'s sensitivity map in the ROI corresponds to a luminosity of $1.7 \times 10^{34} \text{ erg s}^{-1}$ according to the luminosity-to-flux conversion outlined in appendix D. Ref. [60] also states results with a threshold of $L_{\text{th}} = 3 \times 10^{34} \text{ erg s}^{-1}$, which appears more accurate.

We recompute results for the simplified sensitivity model with a larger threshold in appendix F and show that they mirror results from the standard model much better. But differences remain between the standard and simplified model, caused by the fact that bright pulsars distant from Earth have lower flux and therefore can appear unresolved in the standard sensitivity model, while the simplified sensitivity model always marks these pulsars as resolved. Also, the threshold is correlated with MSP population density due to the larger backgrounds close to the Galactic plane, which is not accounted for in the simplified model.

4.1.3 N_r , R_r , and N_{GCE} with cut around Galactic Center

In obtaining our 4FGL-DR2 baseline sample, we cut all MSPs known to be > 2 kpc from the GC as likely unrelated to the GCE, following Ref. [60]. This leaves many PSs of unknown distance from the GC, which might pass or fail the 2 kpc cut if their positions were known (although those which are more than 14° from the GC must fail the 2 kpc cut by virtue of their angular separation alone). When predicting N_r , we do not cut MSPs more than 2 kpc away from the GC in our model. This failure to cut potentially gives rise to a discrepancy between our signal prediction and the data we use for comparison.

Specifically, if the list of observed PSs contains all pulsars in the ROI, regardless of their distance from the GC, then our predicted N_r values can be compared to the observed N_r values. However, if the observed sources are all within 2 kpc of the GC, then we should cut all predicted MSPs more distant than 2 kpc from the GC. In practice, we know that the observed N_r values lie somewhere between these two cases — we have cut pulsars with known galactocentric distances greater than 2 kpc, but we did not cut PSs with angular separation greater than 2 kpc of the GC, and many of the remaining sources are of unknown distance from the GC. In this subsection, we therefore compute how much predicted N_r and R_r change if distant MSPs are cut, to bracket the associated systematic uncertainties (an alternative approach would be to remove the cut and see how much the number of sources changes).

Our default treatment assumes this cut is quite inefficient at removing resolved

sources more than 2 kpc from the GC, and thus does not modify the predicted number of sources. In the opposite limit where we treat this cut as perfectly efficient at removing resolved sources more than 2 kpc from the GC, R_r and N_r fall by 55-75% for most of the luminosity function models (excluding the NPTF model), and for the NPTF model, both N_r and R_r fall by 80%. The strength of this effect for the NPTF model is due to its rather narrow luminosity function and the proximity of the luminosity function peak to the threshold, meaning that cutting sources that lie closer to Earth may make the difference between the population being detectable and undetectable.

This raises the question of the degree to which our predictions rely on extrapolating the GCE into regions of the sky where it may not have been significantly detected. In the main analysis, these regions are modeled with a gNFW² distribution for the source density, but this distribution may not be accurate at high galactocentric distance. To test the importance of the contributions from these regions, we can look at the effect of truncating the source density distribution at 2 kpc from the GC, excluding both resolved sources *and* GCE flux from more distant points. For a fixed point in parameter space, the result is a decrease in F_{GCE} and N_{GCE} of about 30%; consequently, to match the GCE flux requires a higher normalization for the source population within 2 kpc of the GC. In this case the predicted values for N_r and R_r decrease by 40-65% and 5-45% respectively for non-NPTF benchmarks, and by 75% for the NPTF benchmark.

The effect of either truncating the source density distribution or assuming a perfect efficiency for the 2 kpc cut is thus to reduce the number and flux of predicted sources and move all benchmarks away from the regions where there may be tension with the data, in figure 3-1, with a particularly strong effect for the NPTF benchmark.

4.2 Future sensitivity

In this section, we determine the capability of an improved γ -ray telescope to constrain the luminosity function parameters of an MSP population in the GC. We simulate an

| Luminosity function | $N_r^{\times 1}$ | $R_r^{\times 1}$ | $N_r^{\times 2}$ | $R_r^{\times 2}$ | $N_r^{\times 5}$ | $R_r^{\times 5}$ | $N_r^{\times 10}$ | $R_r^{\times 10}$ | $N_r^{\times 20}$ | $R_r^{\times 20}$ |
|---------------------|------------------|------------------|------------------|------------------|------------------|------------------|-------------------|-------------------|-------------------|-------------------|
| Wavelet 1 | 31 | 0.11 | 77 | 0.17 | 230 | 0.25 | 470 | 0.31 | 955 | 0.37 |
| Wavelet 2 | 98 | 0.38 | 210 | 0.53 | 490 | 0.69 | 830 | 0.78 | 1300 | 0.84 |
| Log normal, GLC | 124 | 0.72 | 220 | 0.85 | 380 | 0.95 | 490 | 0.98 | 570 | 0.99 |
| Log normal, GCE | 20 | 0.059 | 73 | 0.12 | 340 | 0.25 | 930 | 0.39 | 2200 | 0.55 |
| Log normal, AIC | 12 | 0.039 | 41 | 0.071 | 180 | 0.14 | 520 | 0.22 | 1400 | 0.32 |
| NPTF | 111 | 0.26 | 460 | 0.70 | 810 | 0.96 | 907 | 0.99 | 940 | 0.999 |
| Disk | 30 | 0.13 | 89 | 0.20 | 370 | 0.34 | 1032 | 0.50 | 2500 | 0.69 |


Table 4.1: Number of resolved MSPs N_r and ratio of resolved flux to total flux R_r for seven luminosity function benchmarks. The flux threshold of the telescope, $F_{\text{th}}(\ell, b)$, has been decreased by a factor of one (i.e., sensitivity is at its current value), two, five, ten, and twenty, with the factor of increase given as superscripts in the header. Apart from the change in threshold, the standard sensitivity model is used.

increase in sensitivity of GCE measurements by reusing the same $F_{\text{th}}(\ell, b)$ sensitivity map (figure 2-6) used in our standard sensitivity model, but with an overall multiplicative decrease. In particular, we study cases where the sensitivity threshold is decreased by a factor of two, five, ten, and twenty and reproduce some of the analyses described earlier in this paper.


4.2.1 Resolved PS flux distributions at higher sensitivity

Table 4.1 indicates that even a doubling in sensitivity of *Fermi* is expected to greatly increase the number of resolvable MSPs in the GC; N_r increases by a factor of 2-4 across our benchmark luminosity functions, and the smallest change is for the GLC model which is already in severe tension with the data due to its large R_r value. R_r increases by a similar factor for models where most of the GCE flux is not already resolved. Greater sensitivity increases, at the 10 – 20 \times level, would be expected to resolve at least 20 – 30% of the GCE even in the most pessimistic benchmark cases.


In figure 4-2, we show the forecast histograms for the expected number of resolved sources as the PS sensitivity increases. MSPs with flux $F \gtrsim 5 \times 10^{-12} \text{ erg s}^{-1} \text{ cm}^{-2}$ are predicted to already be visible in the *Fermi* data set, so that the MSPs resolvable only by a sensitivity increase are low luminosity. With a five-, ten-, or twenty-fold sensitivity increase, we also see that a large fraction of MSPs (approaching 1) are resolved for the GLC and NPTF benchmarks. This fact is quantitatively visible by



(a) Power law luminosity function



(b) Log normal luminosity functions



(c) NPTF luminosity function

Figure 4-2: Histograms of the predicted number of resolvable pulsars required by each luminosity function benchmark studied here to reproduce the GCE, for different models of *Fermi* sensitivity. Darker lines indicate the current sensitivity; lighter lines represent increases in sensitivity by factors of two, five, ten, and twenty. The vertical dotted line represents the energy flux corresponding to a source with an expected photon count of one in the part of the *Fermi* data set used by Ref. [45]. The bins are spaced evenly in log flux space, with width $\Delta F/F = 0.069$.

comparing table 4.1 to the N_{GCE} entries of table 3.1a, and qualitatively visible in figure 4-2 because the distributions begin to take on the shape of the true luminosity function, with the peak of the GLC log-normal curve and the low-luminosity branch of the NPTF broken power-law visible for high sensitivities.

One might ask whether the highest sensitivities studied here are ever plausibly achievable. While achieving these sensitivities purely through increased integration time with *Fermi* does not seem feasible, novel analyses of the existing *Fermi* data set can shed light on the properties of source populations below the current detection threshold. Ref. [45] claims sensitivity to the source count function (closely related to the luminosity function) down to fluxes corresponding to a single photon per source (or even lower in principle), using a neural-network-based analysis, within the subset of *Fermi* data used in that work (which is restricted to a narrower range of energies, and includes other cuts, relative to the data used to generate the 4FGL-DR2 catalog). We convert the one-photon threshold flux in the analysis of Ref. [45] to our 0.1-100 GeV energy band and find it corresponds to a PS flux of $1.91 \times 10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$, which is roughly twenty times lower than the approximate current *Fermi* threshold we calculate in appendix F. This “one photon” line is marked on figure 4-2. Other analyses may be sensitive to one-photon source populations, such as the NPTF technique [19] and another parametric likelihood method known as the Compound Poisson Generator [23].

Consequently, if neural-network-based analyses can achieve sensitivity to PS populations down to their one-photon threshold, it is intriguing that recent studies of this type appear to detect a PS fraction around the 30–40% level; Ref. [45] claims to exclude a non-PS contribution greater than 66% at 95% confidence, while Ref. [49] finds a best-fit PS contribution of 38% (which differs from zero by 2σ). In our analysis, the benchmarks that predict the smallest R_r at this sensitivity level (Wavelet 1 and AIC benchmarks) also predict $R_r = 30\text{--}40\%$; modestly higher R_r values (as found for the GCE benchmark) are also consistent with the results of [45, 49].

Figure 4-2 also demonstrates that flux histograms predicted by different luminosity function configurations begin to visibly diverge at higher sensitivity. The

log-normal populations in particular are visibly distinct; the GLC luminosity function predicts a flatter distribution of MSP flux with peak near $1 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$ while the GCE and GLC benchmarks predict larger and thinner peaks near $F = 3 \times 10^{-13} \text{ erg s}^{-1} \text{ cm}^{-2}$ at twenty-times-greater sensitivities. We perform an initial quantitative analysis of this point below.

4.2.2 Shape differences between resolved PS flux distributions at higher sensitivity

The flux distribution of a population of resolved sources could potentially be a more powerful consistency test of luminosity function models, compared to simply checking the total flux and number of sources, as discussed in section 3.3. However, as demonstrated in that section, the similarity of the high-flux tails of different benchmark luminosity functions, and their tendency to underpredict currently observed sources, makes this test rather uninformative at the present level of sensitivity. In this section we discuss the degree to which improved sensitivity could allow us to distinguish benchmark luminosity function models, in the idealized case where we can accurately identify a subpopulation of inner Galaxy pulsars.

To determine our ability to differentiate between the flux distribution of resolved PSs predicted by different luminosity functions, for every pair of luminosity functions, we take one to be the hypothesis while the other is the true luminosity function. We then draw a mock data set D of observed source fluxes F_i from the true (predicted) flux distribution, and determine the unbinned likelihood for the mock data set given the hypothesis luminosity function P_{hyp} . We draw a number of data points equal to the expected number of resolved MSPs N_r predicted by the true luminosity function, under the assumption that the source population explains 100% of the GCE (these values are given in table 4.1). Because we are interested in discriminating between the high-flux distributions of different source populations based on their shape, not just their normalization, we normalize the hypothesis luminosity function so that it produces the same expected number of resolved PSs as the assumed-true luminosity

function. (Note that this does mean that some hypotheses could be independently excluded by the fact that they overproduce the total GCE flux; in this section we are just using the benchmark luminosity functions as examples of scenarios with somewhat different forms for the high-flux tail.)

We define the likelihood of the mock data given the hypothesis to be:

$$\mathcal{L}_{\text{hyp}} = \prod_{F_i \in D} P(F_i), \quad P(F_i) \propto P_r(F_i) \int_{\Omega} d\Omega \int s^4 ds \rho_{\text{GCE}}(r) P_{\text{hyp}}(4\pi s^2 F_i), \quad (4.1)$$

where $P(F_i)$ is the probability density function describing the probability that a PS is detected with flux F_i . Here, as in eq. 2.3, Ω is the region of interest, s is the distance from the point of integration to Earth, and r is the distance from the point of integration to the GC, which is given by $r^2 = s^2 + r_c^2 - 2sr_c \cos b \cos \ell$, where $r_c = 8.5 \text{ kpc}$ is the distance between Earth and the GC. The probability of a certain PS with flux F being resolved by *Fermi* given that it exists is denoted by $P_r(F)$, which is the sensitivity model (e.g. eq. 2.5).

To test the similarity of the two flux distributions, we calculate the test statistic,

$$\lambda = -2 \ln(\mathcal{L}_{\text{hyp}}/\mathcal{L}_{\text{true}}), \quad (4.2)$$

where $\mathcal{L}_{\text{true}}$ is obtained by substituting P_{hyp} for P_{true} in eq. 4.1. We expect this number to be low for similar luminosity functions (zero for identical luminosity functions) and higher for different ones.

Figure 4-3 displays the λ values for each pair of luminosity functions, averaged over 10,000 sample data sets drawn from each luminosity function, for each sensitivity level studied. As previously, a rescaled version of the standard sensitivity model is used.

As expected, the flux distributions become more distinct as sensitivity increases, as demonstrated by increasing values of λ . This is largely due to the fact that N_r increases when sensitivity increases, thereby amplifying the likelihood ratio. Note that our procedure has the effect that scenarios where the true model has many

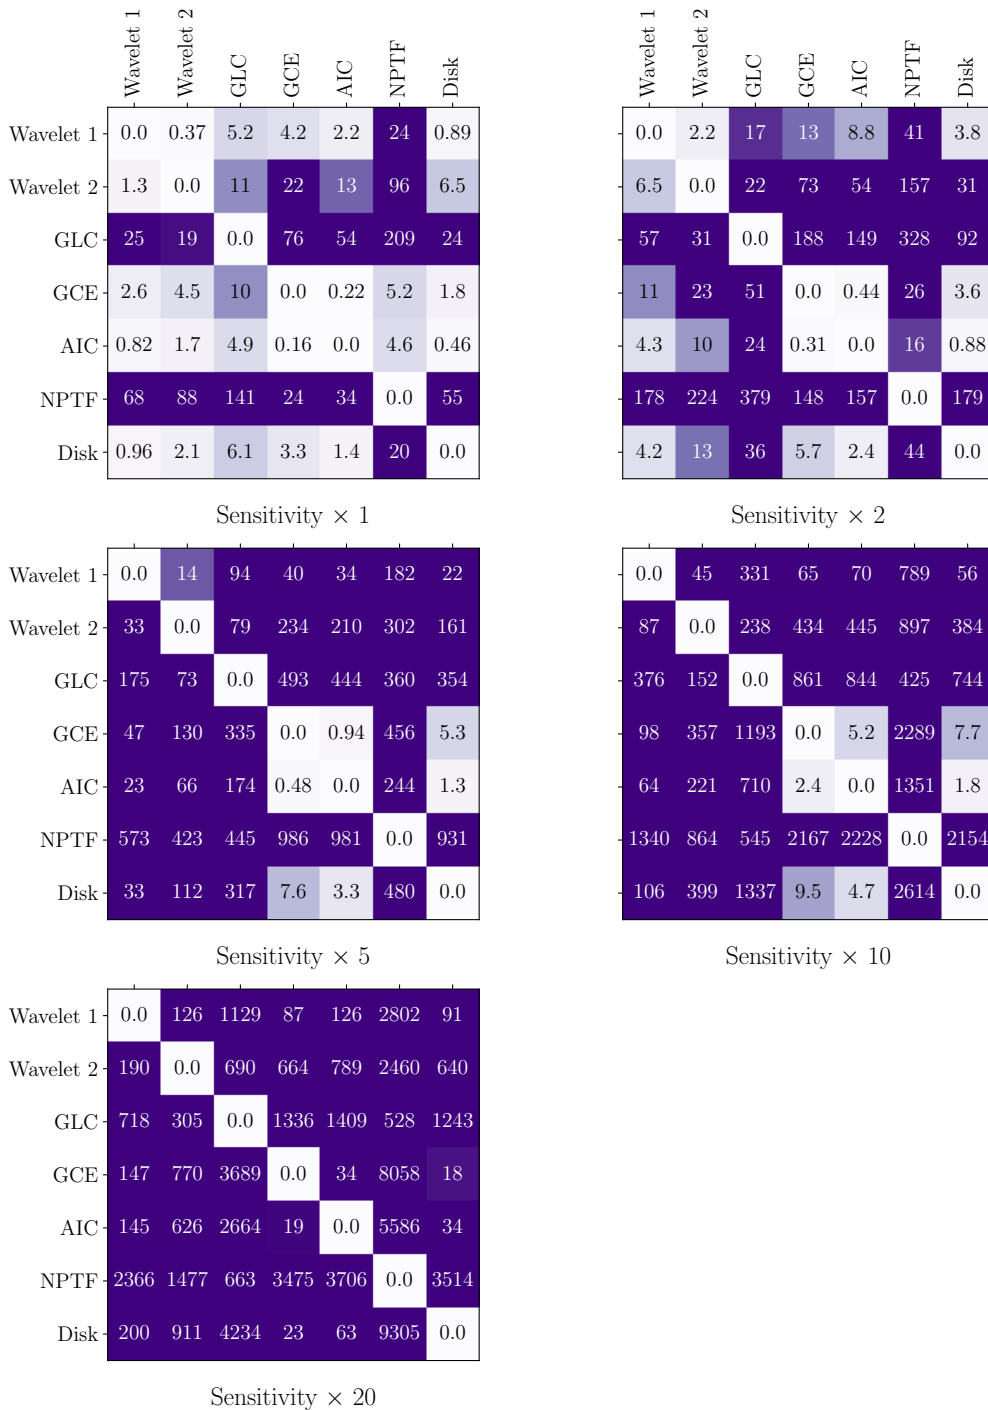


Figure 4-3: Map of the expected value of λ resulting from assuming one luminosity function (horizontal axis) as the hypothesis while the other (vertical axis) is the true luminosity function. Each luminosity function is normalized to produce the same number of resolved PSs as the assumed-true luminosity function. Each plot represents a different multiplicative increase in the *Fermi* sensitivity level. Diagonals indicate the fit of a luminosity function to itself. The standard sensitivity model is used.

bright resolved sources allow models to be more easily distinguished, because the true model's N_r sets the number of source fluxes in D (defined in eq. 4.1), and thereby the sensitivity of λ to small differences between P_{hyp} and P_{true} . For example, the GCE and AIC luminosity functions are hard to distinguish from the Disk luminosity function even at high sensitivity, when the GCE or AIC functions (with relatively few bright sources) are assumed true. However, when the Disk luminosity function is assumed true, the difference with GCE/AIC increases due in part to the large number of resolved sources produced by the Disk luminosity function. Other similar pairs of luminosity functions highlighted in figure 4-3 are the GCE and AIC functions which share the same log-normal functional form, and the power-law functions Wavelet 1 and 2 to a lesser extent. The NPTF luminosity function benchmark is distinctly different from the others, even at low sensitivity, as discussed above.

Until now, we have avoided showing the resolved flux distribution predicted by the smoothed sensitivity model because that model is designed to predict the probability that a PS is resolved as a MSP specifically. It therefore predicts many fewer resolved sources than the standard model and the threshold is shifted to higher flux. But the shape and slope of the smoothed cutoff is largely independent of K_{th} , which merely shifts the MSP detection threshold to higher luminosities (eq. 3.3) and is responsible for much of the reduction in N_r . The above analysis, including figures 4-2 and 4-3, may therefore be informative when recalculated for the smoothed model; it may also be easier in practice to avoid contamination from non-GCE sources when focusing on resolved MSPs rather than just resolved PSs. This exercise is done in appendix G.

Chapter 5

Conclusions

We have explored the total number of MSPs, number of resolved MSPs, and the fraction of flux emitted by resolved MSPs required to produce the observed GCE flux. We extract these properties, for a detection threshold corresponding to the 4FGL-DR2 *Fermi*-LAT source catalog, in a broad scan over the parameter space of the commonly-used power-law and log-normal luminosity function parameterizations. We also benchmark these functions and a broken power-law luminosity function with seven configurations found in the literature. We compared the predicted resolved source populations with sources in the 4FGL-DR2 catalog that pass cuts to reject extragalactic sources, as well as with a subpopulation of such sources identified by the wavelet-based analysis of Ref. [60] (following up on the earlier study of Ref. [11]).

We found that observationally viable luminosity functions can produce between hundreds and millions of MSPs in the GCE without being excluded by overproduction of resolved sources. The high end of this range corresponds to luminosity functions with very low average gamma-ray luminosities ($\lesssim 10^{31}$ erg s $^{-1}$ in the 0.1-100 GeV energy band) compared to physically or observationally motivated benchmarks. The low end of this range, with less than $\mathcal{O}(10^4)$ MSPs, generally requires a strongly peaked luminosity function with relatively high average luminosity. Luminosity functions derived from physical models of MSPs and fitted to GCE data can conform to observations and generally predict tens to hundreds of thousands of MSPs, without any tension with studies finding that masking resolved PSs does not significantly reduce

the GCE [60].

Of the benchmarks we tested, the luminosity functions inferred from early non-Poissonian template fitting analyses [44], from observations of globular clusters [37], and from a search for wavelet peaks in older *Fermi* data [11] appear to produce more resolved sources/flux than are contained in the subpopulation identified via wavelet methods in Ref. [60], while the other benchmarks are consistent with this subpopulation (underproducing it). However, it is not clear whether this subpopulation can truly be used as an upper bound on resolved GCE sources, as there are a large number of sources in 4FGL-DR2 which satisfy the same cuts as the wavelet-selected population (imposed to reject extragalactic sources) except that they were not significantly detected by the wavelet analysis. This is not an issue of the wavelet analysis simply having a higher sensitivity threshold than the 4FGL-DR2 analysis; many of the 4FGL-DR2 sources not detected by the wavelet method are quite bright. To strengthen the constraints on the GCE PS population, it would be helpful to better understand the completeness properties of the two source lists above the 4FGL-DR2 sensitivity threshold, and to further study the properties of 4FGL-DR2 sources to separate possible GCE candidate sources from others. We have provided estimates of how the constraints on PS populations would behave as a function of the fraction of resolved PSs attributed to the GCE, or the fraction of the total GCE flux attributed to resolved sources, both with current data and with an improved PS sensitivity threshold.

We also tested the ability of current and future analyses to distinguish between the benchmark luminosity functions, finding that the high-flux tails of these luminosity functions are rather similar and would currently be difficult to distinguish even if we could correctly identify all the GCE PSs above threshold. However, improvements in sensitivity — either from future data or from improved analyses — could change this conclusion. We explored the fraction of the GCE that would be resolved by analyses with point-source sensitivity down to the 1-photon threshold of Ref. [45], and found that the benchmark scenarios generally predicted that fractions exceeding 30% of the GCE flux would be resolved in this case. This estimate is intriguingly similar to the

claimed fraction of the GCE tentatively detected as PSs in Ref. [49], and consistent with the limit in Ref. [45].

We found there are a number of significant systematic uncertainties which must be taken into account when claiming to exclude or match specific luminosity function models. For example, the total flux attributed to the GCE varies by up to a factor of two between different studies; we choose a benchmark value based on a recent analysis [25] and then show the effects of modifying this choice. Varying the total GCE flux modifies the required properties of a PS population explaining some fraction of the GCE; for a fixed luminosity function, a higher total flux implies a higher number of resolved sources (although they will yield the same total fraction of the GCE flux). Furthermore, especially for steeply peaked luminosity functions where much of the power is in sources that are barely resolved or barely unresolved, the distribution of the GCE at large distances from the GC is potentially very important — while most of the total flux of the GCE may originate from the region around the GC, the resolved sources can dominantly originate from regions closer to Earth and away from the line of sight to the GC, where the sensitivity is improved. Finally, the sensitivity modeling is important; we have demonstrated that a simple sensitivity estimate employed in Ref. [60] predicts quite a different distribution for the fluxes of resolved sources compared to observations, and suggest an improved prescription.

Appendix A

Scaling of ROIs and spectral ranges

To establish a value for the total flux of the GCE, we draw on several analyses of the GCE spectrum in section 2.1.2. However, not all the analyses we study use the same ROI as ours. To convert between our ROI and others', we assume a gNFW squared spatial distribution of MSPs in the GCE as discussed in section 2.1.1. Then we calculate the ratio of flux in our region of interest F_Ω to flux in another analysis's region of interest $F_{\Omega'}$ via

$$\frac{F_{\Omega'}}{F_\Omega} = \left[\int_{\Omega'} d\Omega \int_0^\infty ds \rho_{\text{GCE}}(r) \right] \left[\int_\Omega d\Omega \int_0^\infty ds \rho_{\text{GCE}}(r) \right]^{-1}. \quad (\text{A.1})$$

Here, s represents the distance between Earth and the point of integration, and r represents the galactocentric distance. They are related by $r^2 = s^2 + r_c^2 - 2sr_c \cos \ell \cos b$, where ℓ and b are the Galactic longitude and latitude and $r_c = 8.5$ kpc is the distance between Earth and the center of the Galaxy.

The flux ratio as computed by eq. A.1 between our region and a $40^\circ \times 40^\circ$ square region without the Galactic disk mask is 1.9. Similarly, the ratio is 1.8, 1.5, 1.3, 1.1, and 0.92 for regions centered on the Galactic center of side length $30^\circ \times 30^\circ$, $20^\circ \times 20^\circ$, $15^\circ \times 15^\circ$, $10^\circ \times 10^\circ$, and $7^\circ \times 7^\circ$ respectively, all without a disk mask. The $15^\circ \times 15^\circ$ square is the ROI used by Ref. [7], and the $7^\circ \times 7^\circ$ square is used by Ref. [32]. For Ref. [6], we reproduce their ROI from the 3FGL PS catalog [5] and obtain a flux ratio

of 0.56.¹ Finally, Ref. [2] uses the same $7^\circ \times 7^\circ$ ROI as Ref. [32], but with $\gamma = 1.1$. The ratio for this ROI to our ROI with $\gamma = 1.1$ is also 0.56.

We perform the ROI conversion by taking the flux emitted from the entire ROI Ω' of the study in question and multiplying by $\frac{F_\Omega}{F_{\Omega'}}$ to get the flux in our ROI Ω . To make figure 2-1, we perform this multiplication for the flux in every energy bin.

¹This ROI is a 10° -radius disk, pixelated with pixel size $a \approx 0.47^\circ$. We place $\approx 1^\circ + a \sin(\pi/4)$ -radius circular masks around the 200 brightest 3FGL sources in the sky.

Appendix B

Broken power-law fits to GCE spectra

Section 2.1.2 relied on broken power-law fits to nine energy spectra found in the literature. Those fits were done via χ^2 minimization, with

$$\chi^2 = \sum_i \begin{cases} \frac{(y_i - y_i^*)^2}{\sigma_{+,i}^2} & y_i^* \geq y_i \\ \frac{(y_i - y_i^*)^2}{\sigma_{-,i}^2} & y_i^* \leq y_i \end{cases}$$

where y_i^* is the model data given the fit parameters, y_i is the observed data, and $\sigma_{\pm,i}$ are the upper and lower error bar lengths.

As mentioned in section 2.1.2, the fit is done to a one-parameter broken power-law model with n_1, n_2, L_b fixed at parameters determined by Ref. [16], as well as a four-parameter broken power-law with all parameters floated. The resulting fits are shown in figure B-1.

The best-fit results for the floated parameters do not agree perfectly between this analysis of data from Ref. [16] and the best-fit parameters reported in that work, although they are similar. This is likely due to the fact that this work approximates the error bars as uncorrelated and adds systematic and statistical uncertainties (where they are separated) in quadrature, whereas Ref. [16] used a full covariance matrix.

Several of the studies report multiple spectra for the GCE. For the spectrum




Figure B-1: Broken power-law fits to the nine spectra studied in section 2.1.2 to extract the total GCE luminosity. Arrows on error bars denote bars that would descend past the lower edge of the plot. See text for which functions are being fitted. Fits are performed approximating the uncertainties at different energies as uncorrelated, which is why the results of the two fits differ slightly for the “Calore 2015” spectrum taken from Ref. [16]; Ref. [16] took the full covariance matrix into account.


labeled “Abazajian 2014”, we take the GCE spectrum extracted from the “2FGL + 2PS + I + MG + ND + GCE” fit, which provides the best fit [2]. For the spectrum labeled “Ajello 2017”, we use the “Sample” spectrum and recreate the ROI from the 4FGL catalog [6]. For the spectrum labeled “Ajello 2016”, we use the plotted spectra in Figure 13 of Ref. [7].

Appendix C


Fits to GCE and AIC luminosity functions

We claimed in section 2.3 that the AIC and GCE luminosity functions can be reasonably well-described by a log-normal function, which motivates treating the log-normal form as a plausible parameterization for a range of physically reasonable luminosity functions. In this appendix, we explain how the parameters were derived for this log-normal approximation to the relevant luminosity functions.

In both cases, the numerical luminosity functions were extracted from the papers where they were presented: Refs. [53] and [29] respectively. The functions are not normalized, which is why the vertical axes are marked with an arbitrary additional term $+C$. We sampled the functions at 100 evenly log-spaced values of luminosity spanning the domain of the original functions (this number was chosen arbitrarily), and used the width of the band at each luminosity value to estimate an uncertainty on the value. We then performed a least-squares fit to these extracted data for log-normal (eq. 2.7) and cutoff power-law (eq. 2.6) luminosity function models. The best-fitting models are shown in Figures C-1a and C-1b respectively; we observe that for the GCE luminosity function the log-normal fit very accurately tracks the numerical result, and for the AIC luminosity function the log-normal model lies consistently within the uncertainty band except at very low and high luminosities. This is not the case for the cutoff power-law fit in either case.



(a) Fit to the GCE luminosity function



(b) Fit to the AIC luminosity function

Figure C-1: Log normal fits to two numerical luminosity functions: the boxy bulge luminosity function from Ref. [53] and the luminosity function from [29]. Uncertainties are assumed to be uncorrelated and represent one standard deviation from the mean. The best fit for a broken power-law with an exponential cutoff is also shown.

Appendix D

Conversion between GCE luminosity and flux

This analysis requires luminosity functions to be expressed as probability distributions as a function of luminosity. Yet several papers referenced in this study express luminosity functions as a function of flux, and sometimes also assume that the flux distribution is constant along all lines of sight (rather than assuming a consistent luminosity distribution, which translates into different flux distributions depending on the distribution of distances to the sources). This appendix describes how the conversion to luminosity is done.

As discussed in section 2.1.1, we represent the MSP population as distributed according to an gNFW squared distribution with $\gamma = 1.2$. One way to convert a function of flux to luminosity would be to use the luminosity function which, when integrated over the gNFW squared spatial distribution, would reproduce the observed function of flux. But this method would change the functional form of the luminosity function so that, for example, a luminosity function that is log-normal when written in terms of flux would no longer be log-normal when written in terms of luminosity; it would also require a spatially-dependent luminosity function. For source populations distributed through the Galaxy, this difference is crucial; an intrinsically faint source may be very visible simply because it is close to the Earth. However, for the GCE population, we expect most sources to be rather close to the GC and thus to

have an approximately consistent distance from Earth, suggesting we may be able to approximate flux and luminosity as being proportional. This is the tacit assumption made in studies which treat the flux distribution of GCE sources as being the same along all lines of sight (as in general this would require a peculiar coincidence in the spatial evolution of the luminosity function).

Consequently, we simply convert the flux value of every bin to a luminosity according to the following method. Suppose the entire population of MSPs only contains pulsars with luminosity L ; then we can compute the average flux F per pulsar. Integrated over the region of interest Ω , this yields a constant flux to luminosity ratio of

$$\frac{F}{L} = \frac{1}{4\pi} \left[\int_{\Omega} d\Omega \int_0^{\infty} ds \rho_{\text{GCE}}(r) \right] \left[\int_{\Omega} d\Omega \int_0^{\infty} s^2 ds \rho_{\text{GCE}}(r) \right]^{-1} = 1.11 \times 10^{-46} \text{ cm}^{-2}. \quad (\text{D.1})$$

Here, s represents the distance from Earth to the point of integration, and r represents the distance from the GC to the point of integration. They are related by the law of cosines: $r^2 = s^2 + r_c^2 - 2sr_c \cos b \cos \ell$, where ℓ is the Galactic longitude. The numerical value reported was computed for $r_c = 8.5 \text{ kpc}$. It is slightly lower than the naïve value of $\frac{F}{L} = \frac{1}{4\pi r_c^2} = 1.16 \times 10^{-46} \text{ cm}^{-2}$, which assumes that all the MSPs are at the Galactic center, and therefore does not rely on a choice of γ . The similarity between the flux-to-luminosity ratio at the GC and the average flux-to-luminosity ratio makes it seem probable that it is quite a good approximation to treat the sources as being close to the GC, and the error due to this approximation is likely to be small. However, when computing the number of resolvable sources, which has a strong dependence on each source's distance to the GC, we do not use this approximation.

Appendix E

Calculation of photon energy

The break flux of the NPTF luminosity function benchmark of Ref. [44] is given with flux units of photons per square centimeter per second ($F_b^c = 1.76 \times 10^{-10}$ photon $\text{cm}^{-2} \text{s}^{-1}$), evaluated in the 1.893–11.943 GeV energy range. Our analysis requires a break luminosity L_b in units of ergs per second, in the 0.1–100 GeV energy range. To do this conversion, we use the GCE spectrum inferred from Ref. [25] by floating all four broken power-law parameters as described in section 2.1.2 and plotted in figure 2-2.

The conversion between F_b^c and L_b is then achieved by

$$L_b = F_b^c \left[\int_{0.1 \text{ GeV}}^{100 \text{ GeV}} E_\gamma N(E_\gamma) dE_\gamma \right] \left[\int_{1.893 \text{ GeV}}^{11.943 \text{ GeV}} N(E_\gamma) dE_\gamma \right]^{-1} \frac{L}{F} \quad (\text{E.1})$$

where $N(E_\gamma) = dN/dE$ is the spectrum of the GCE, and the fraction $\frac{L}{F}$ is given in appendix D. This calculation gives $L_b = 4.4 \times 10^{36}$ erg s^{-1} (in the 0.1-100 GeV energy band relevant for our analysis).

Appendix F

Analysis of L_{th} for the simplified sensitivity model

The simplified sensitivity model outlined in section 2.2.2 represented all PSs with luminosity $L > L_{\text{th}}$ as resolved, and all with $L < L_{\text{th}}$ as unresolved, where L_{th} is pixel-independent. The key quantities N_{GCE} , R_r , and N_r of a population of MSPs necessary to reproduce the GCE were then given by eq. 2.4.

The value $L_{\text{th}} = 10^{34} \text{ erg s}^{-1}$ was used in Ref. [60], and therefore was used in this paper for the simplified sensitivity model. However, the fact that the N_{GCE} , R_r , and N_r values produced by the simplified sensitivity model are almost always larger than those produced by the more detailed standard sensitivity model (table 3.1) indicates that $L_{\text{th}} = 10^{34} \text{ erg s}^{-1}$ is an underestimate. The fact that the none of the 47 PSs in the subpopulation identified by Ref. [60] have L below or near $10^{34} \text{ erg s}^{-1}$ (figure 3-3) also demonstrates that the true average threshold is larger; we would expect some PSs to be observed at or even below the true threshold due to uncertainty in the threshold and the fact that some regions of the sky have greater sensitivity than others.

A more accurate estimate for a constant L_{th} could be gained from an average over the per-pixel threshold sensitivities provided by Refs. [4, 9] (figure 2-6). We weight our average by the amount of flux predicted to emanate from each pixel with a gNFW-squared-distributed population of PSs. The new weighted-average flux threshold is

then

$$F_{\text{th}} = \left[\sum_{\mathcal{P}(\ell,b)} L_{\text{th}}(\ell,b) \int_0^\infty ds \rho_{\text{GCE}}(r) \int_{\mathcal{P}(\ell,b)} d\Omega \right] \left[\sum_{\mathcal{P}(\ell,b)} \int_0^\infty ds \rho_{\text{GCE}}(r) \int_{\mathcal{P}(\ell,b)} d\Omega \right]^{-1} \quad (\text{F.1})$$

where $\mathcal{P}(\ell, b)$ represents the pixel at galactic coordinates (ℓ, b) . As before, r represents the distance from the line of sight distance s to the center of the Galaxy, defined by $r^2 = s^2 + r_c^2 - 2sr_c \cos(\ell) \cos(b)$. We assume here that the average luminosity emitted from a PS is not correlated with its position, so that the luminosity emanating from a region of space is proportional to the number density ρ_{GCE} of PSs in that region. Then the s^2 of the volume element and s^{-2} required to convert luminosity to flux cancel out.

The result of eq. F.1 is $L_{\text{th}} = 3.8 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$, which corresponds to $L_{\text{th}} = 3.4 \times 10^{34} \text{ erg s}^{-1}$ via the flux-luminosity conversion defined in appendix D. Note that a few of the 47 wavelets have flux below this value of F_{th} (figure 3-3a), demonstrating that it is a better candidate for *Fermi*'s sensitivity.

Regenerating previous results with this new $L_{\text{th}} = 3.4 \times 10^{34} \text{ erg s}^{-1}$ value, figure F-1 displays the configurations of luminosity functions with the power-law or log-normal functional form that obey the observational constraints, according to the simplified sensitivity model with the new threshold. It is the analog of figure 3-1b. Table F.1 shows N_r , R_r , and N_{GCE} for specific benchmarks drawn from other works, again with the simplified model and the new threshold. It is the analog of table 3.1b.

We see from table F.1 that the new step-function threshold values N_r and R_r for $L_{\text{th}} = 3.4 \times 10^{34} \text{ erg s}^{-1}$ are generally closer to the standard values than the $L_{\text{th}} = 10^{34} \text{ erg s}^{-1}$ were, indicating that the averaged threshold $L_{\text{th}} = 3.4 \times 10^{34} \text{ erg s}^{-1}$ is a better estimate than $L_{\text{th}} = 10^{34} \text{ erg s}^{-1}$. The values for N_{GCE} do not change because they are not affected by the threshold. The NPTF luminosity function is still poorly represented by the new value of L_{th} because its strong peak just below the threshold sensitivity makes it behave very differently in the standard sensitivity model, where the peak is widened by the spatial distribution of the MSPs.

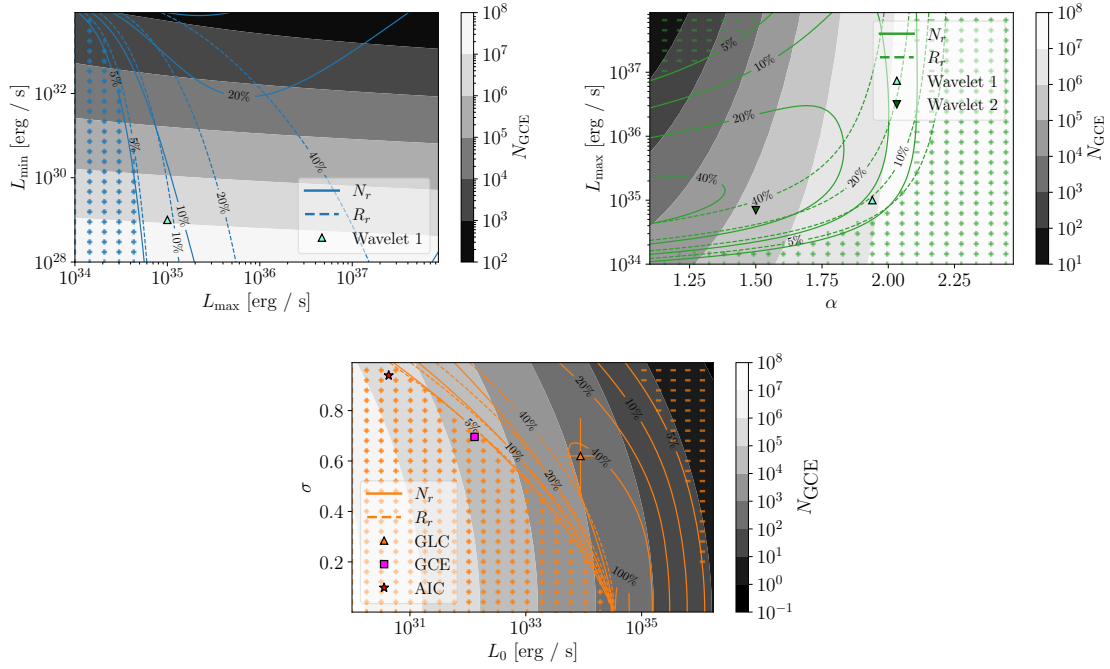


Figure F-1: As figure 3-1, but using the simplified sensitivity model with the average flux threshold of $L_{\text{th}} = 3.4 \times 10^{34} \text{ erg s}^{-1}$.

| Luminosity function | N_r | R_r | N_{GCE} |
|---------------------|-------|----------------------|-------------------|
| Wavelet 1 | 22 | 0.089 | 8.5×10^6 |
| Wavelet 2 | 79 | 0.32 | 2.2×10^5 |
| GLC | 113 | 0.68 | 660 |
| GCE | 8.6 | 0.030 | 3.4×10^4 |
| AIC | 5.9 | 0.023 | 3.6×10^5 |
| NPTF | 0.47 | 1.1×10^{-3} | 960 |
| Disk | 19 | 0.10 | 2.5×10^4 |

Table F.1: Number of resolved PSs, ratio of resolved flux to total flux, and total number of PSs predicted to make up the GCE based on the new $L_{\text{th}} = 3.56 \times 10^{34} \text{ erg s}^{-1}$ simplified sensitivity threshold produced by a flux-weighted average of per-pixel sensitivity provided by Refs. [4, 9]. Compare to table 3.1b, which produces higher values.

The abrupt change caused by the new value of L_{th} (compare Figs. 3-1b and F-1) demonstrates the large impact of L_{th} on the range of allowed luminosity function configurations. In particular, with the more accurate, averaged value of L_{th} , the observational limits are satisfied at larger L_{min} than with the old averaged value, which allows N_{GCE} to be lower. For the log-normal case, slightly higher L_0 values are allowed by the averaged L_{th} , which corresponds to lower N_{GCE} .

Appendix G

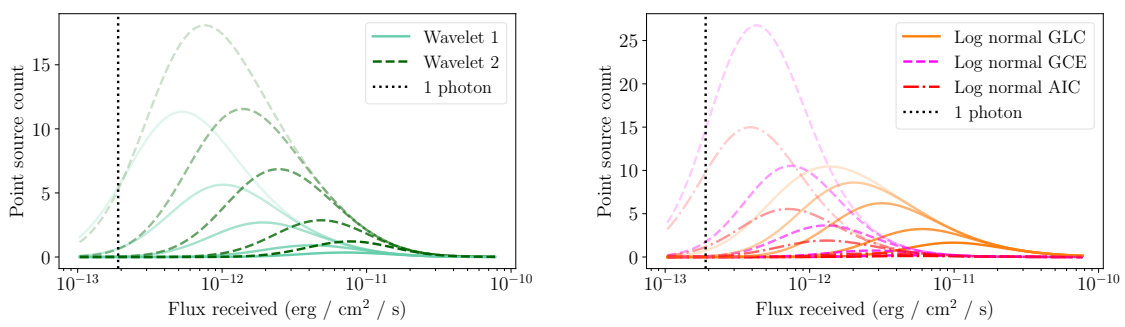
Flux histograms for the smoothed sensitivity model

The smoothed sensitivity model was not used in section 4.2 because its parameters were extracted to represent the probability that a PS is resolved as an MSP [53, 29], whereas in this study we focus on the probability for a PS to be resolved at all. Therefore, the resolved PS flux distributions using the smoothed sensitivity model cannot be compared directly to the current observational data. However, since the smoothed sensitivity model is more detailed and accounts for uncertainty in the thresholds, and might better represent the source populations we could positively identify with the GCE (via identifying the sources as pulsars), we present the flux distributions of resolved MSPs using the smoothed sensitivity model in this appendix for reference.

Figure G-1 is the analog of figure 4-2, showing the flux distributions at different sensitivities of resolved MSPs, assuming different luminosity functions, using the smoothed sensitivity model. Again, the histogram uses 100 flux bins distributed uniformly over the range $1 \times 10^{-13} - 8 \times 10^{-11} \text{ erg s}^{-1} \text{ cm}^{-2}$


Figure G-2 is the analog of figure 4-3, showing λ values (defined in eq. 4.2) attained from fitting the distributions to each other. N_r has been re-scaled to be equal to the N_r predicted by the assumed-true luminosity function. The pairs of luminosity functions that were similar for figure 4-3 are also similar for figure G-2.

As expected, figure G-1 shows that the smoothed sensitivity model visibly smooths



(a) Power law luminosity function

(b) Log normal luminosity functions



(c) NPTF luminosity function

Figure G-1: Histogram of the predicted number of resolvable pulsars required by each luminosity function benchmark to reproduce the GCE flux, for different models of *Fermi* sensitivity under the smoothed sensitivity model. Darker lines indicate the current sensitivity; lighter lines represent increases in sensitivity by factors of two, five, ten, and twenty. See text for details. The bins are spaced evenly in log flux space, with width $\Delta F/F = 0.069$.

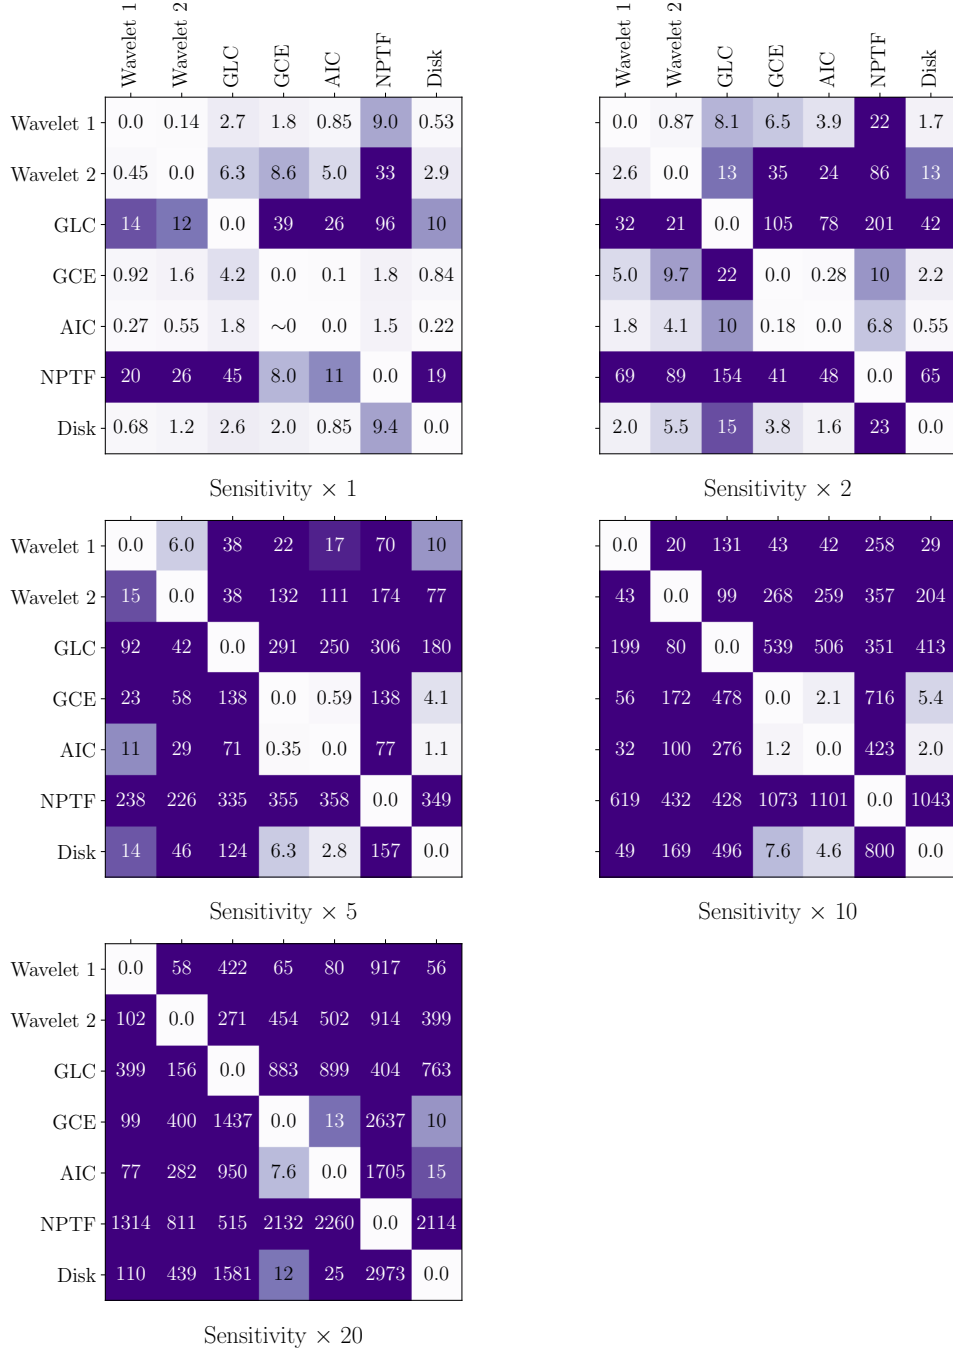


Figure G-2: Map of the expected value of λ resulting from assuming one luminosity function (horizontal axis) as the hypothesis while the other (vertical axis) is the true luminosity function. Each luminosity function is normalized to produce the same number of resolved PSs as the assumed-true luminosity function. Each plot represents a different multiplicative increase in the *Fermi* sensitivity level. Diagonals indicate the fit of a luminosity function to itself. The smoothed sensitivity model is used.

the low-flux falloff of the histograms of MSPs resolved. This smoothing nearly always makes flux distributions more similar to each other, as can be seen by comparing Figs. 4-3 and G-2 and noting which figure contains the higher value of λ for each pair. This is likely due to the fact that the smoothed sensitivity model produces fewer MSPs than the standard model, thereby increasing the relative size of error bars on the smoothed model data.

Bibliography

- [1] Kevork N. Abazajian. The Consistency of Fermi-LAT Observations of the Galactic Center with a Millisecond Pulsar Population in the Central Stellar Cluster. *JCAP*, 03:010, 2011.
- [2] Kevork N. Abazajian, Nicolas Canac, Shunsaku Horiuchi, and Manoj Kaplinghat. Astrophysical and Dark Matter Interpretations of Extended Gamma-Ray Emission from the Galactic Center. *Phys. Rev. D*, 90(2):023526, 2014.
- [3] Kevork N. Abazajian, Shunsaku Horiuchi, Manoj Kaplinghat, Ryan E. Keeley, and Oscar Macias. Strong constraints on thermal relic dark matter from Fermi-LAT observations of the Galactic Center. *Phys. Rev. D*, 102(4):043012, 2020.
- [4] S. Abdollahi et al. *Fermi* Large Area Telescope Fourth Source Catalog. *Astrophys. J. Suppl.*, 247(1):33, 2020.
- [5] F. Acero et al. *Fermi* Large Area Telescope Third Source Catalog. *Astrophys. J. Suppl.*, 218(2):23, 2015.
- [6] M. Ackermann et al. The *Fermi* Galactic Center GeV Excess and Implications for Dark Matter. *Astrophys. J.*, 840(1):43, 2017.
- [7] M. Ajello et al. *Fermi*-LAT Observations of High-Energy γ -Ray Emission Toward the Galactic Center. *Astrophys. J.*, 819(1):44, 2016.
- [8] Walter Baade and Fritz Zwicky. Remarks on super-novae and cosmic rays. *Physical Review*, 46(1):76, 1934.
- [9] J. Ballet, T. H. Burnett, S. W. Digel, and B. Lott. *Fermi* Large Area Telescope Fourth Source Catalog Data Release 2. 5 2020.
- [10] R. T. Bartels, T. D. P. Edwards, and C. Weniger. Bayesian model comparison and analysis of the Galactic disc population of gamma-ray millisecond pulsars. *Mon. Not. Roy. Astron. Soc.*, 481(3):3966–3987, 2018.
- [11] Richard Bartels, Suraj Krishnamurthy, and Christoph Weniger. Strong support for the millisecond pulsar origin of the Galactic center GeV excess. *Phys. Rev. Lett.*, 116(5):051102, 2016.

- [12] Richard Bartels, Emma Storm, Christoph Weniger, and Francesca Calore. The Fermi-LAT GeV excess as a tracer of stellar mass in the Galactic bulge. *Nature Astron.*, 2(10):819–828, 2018.
- [13] Malte Buschmann, Nicholas L. Rodd, Benjamin R. Safdi, Laura J. Chang, Siddharth Mishra-Sharma, Mariangela Lisanti, and Oscar Macias. Foreground Mis-modeling and the Point Source Explanation of the Fermi Galactic Center Excess. *Phys. Rev. D*, 102(2):023023, 2020.
- [14] F. Calore, F. Donato, and S. Manconi. Dissecting the inner Galaxy with γ -ray pixel count statistics. 2 2021.
- [15] Francesca Calore, Ilias Cholis, Christopher McCabe, and Christoph Weniger. A Tale of Tails: Dark Matter Interpretations of the Fermi GeV Excess in Light of Background Model Systematics. *Phys. Rev. D*, 91(6):063003, 2015.
- [16] Francesca Calore, Ilias Cholis, and Christoph Weniger. Background Model Systematics for the Fermi GeV Excess. *JCAP*, 03:038, 2015.
- [17] Sascha Caron, Germán A. Gómez-Vargas, Luc Hendriks, and Roberto Ruiz de Austri. Analyzing γ -rays of the Galactic Center with Deep Learning. *JCAP*, 05:058, 2018.
- [18] S. Chandrasekhar. The density of white dwarf stars. *Philosophical Magazine*, 11(70):592–596, 1931.
- [19] Laura J. Chang, Siddharth Mishra-Sharma, Mariangela Lisanti, Malte Buschmann, Nicholas L. Rodd, and Benjamin R. Safdi. Characterizing the nature of the unresolved point sources in the Galactic Center: An assessment of systematic uncertainties. *Phys. Rev. D*, 101(2):023014, 2020.
- [20] Ilias Cholis, Dan Hooper, and Tim Linden. A New Determination of the Spectra and Luminosity Function of Gamma-Ray Millisecond Pulsars. 7 2014.
- [21] Ilias Cholis, Dan Hooper, and Tim Linden. Challenges in Explaining the Galactic Center Gamma-Ray Excess with Millisecond Pulsars. *JCAP*, 06:043, 2015.
- [22] Brendan Coleman, Dylan Paterson, Chris Gordon, Oscar Macias, and Harrison Ploeg. Maximum Entropy Estimation of the Galactic Bulge Morphology via the VVV Red Clump. *Mon. Not. Roy. Astron. Soc.*, 495(3):3350–3372, 2020.
- [23] Gabriel H. Collin, Nicholas L. Rodd, Tyler Erjavec, and Kerstin Perez. A Compound Poisson Generator approach to Point-Source Inference in Astrophysics. 4 2021.
- [24] Tansu Daylan, Douglas P. Finkbeiner, Dan Hooper, Tim Linden, Stephen K.N. Portillo, Nicholas L. Rodd, and Tracy R. Slatyer. The characterization of the gamma-ray signal from the central milky way: A case for annihilating dark matter. *Physics of the Dark Universe*, 12:1–23, 2016.

- [25] Mattia Di Mauro. Characteristics of the Galactic Center excess measured with 11 years of *Fermi*-LAT data. *Phys. Rev. D*, 103(6):063029, 2021.
- [26] Jack T. Dinsmore and Tracy R. Slatyer. Luminosity functions consistent with a pulsar-dominated Galactic Center Excess. 12 2021.
- [27] Claude-Andre Faucher-Giguere and Victoria M. Kaspi. Birth and evolution of isolated radio pulsars. *Astrophys. J.*, 643:332–355, 2006.
- [28] Kenneth C Freeman. On the disks of spiral and s0 galaxies. *The Astrophysical Journal*, 160:811, 1970.
- [29] Anuj Gautam, Roland M. Crocker, Lilia Ferrario, Ashley J. Ruiter, Harrison Ploeg, Chris Gordon, and Oscar Macias. Millisecond Pulsars from Accretion Induced Collapse naturally explain the Galactic Center Gamma-ray Excess. 6 2021.
- [30] Peter L. Gonthier, Alice K. Harding, Elizabeth C. Ferrara, Sara E. Frederick, Victoria E. Mohr, and Yew-Meng Koh. Population syntheses of millisecond pulsars from the Galactic Disk and Bulge. *Astrophys. J.*, 863(2):199, 2018.
- [31] Lisa Goodenough and Dan Hooper. Possible Evidence For Dark Matter Annihilation In The Inner Milky Way From The Fermi Gamma Ray Space Telescope. 10 2009.
- [32] Chris Gordon and Oscar Macías. Dark matter and pulsar model constraints from galactic center fermi-lat gamma-ray observations. *Phys. Rev. D*, 88:083521, Oct 2013.
- [33] Chris Gordon and Oscar Macias. Dark Matter and Pulsar Model Constraints from Galactic Center Fermi-LAT Gamma Ray Observations. *Phys. Rev. D*, 88(8):083521, 2013. [Erratum: Phys.Rev.D 89, 049901 (2014)].
- [34] Daryl Haggard, Craig Heinke, Dan Hooper, and Tim Linden. Low Mass X-Ray Binaries in the Inner Galaxy: Implications for Millisecond Pulsars and the GeV Excess. *JCAP*, 05:056, 2017.
- [35] G. Hobbs, R. Manchester, A. Teoh, and M. Hobbs. The ATNF Pulsar Catalog. In Fernando Camilo and Bryan M. Gaensler, editors, *Young Neutron Stars and Their Environments*, volume 218, page 139, January 2004.
- [36] Dan Hooper and Lisa Goodenough. Dark matter annihilation in the galactic center as seen by the fermi gamma ray space telescope. *Physics Letters B*, 697(5):412–428, 2011.
- [37] Dan Hooper and Tim Linden. The gamma-ray pulsar population of globular clusters: Implications for the GeV excess. *JCAP*, 2016(08), 8 2016.

- [38] Dan Hooper and Gopolang Mohlabeng. The Gamma-Ray Luminosity Function of Millisecond Pulsars and Implications for the GeV Excess. *JCAP*, 03:049, 2016.
- [39] Constantinos Kalapotharakos, Alice K. Harding, Demosthenes Kazanas, and Zorawar Wadiasingh. A Fundamental Plane for Gamma-Ray Pulsars. *Astrophys. J. Lett.*, 883(1):L4, 2019.
- [40] Rebecca K. Leane and Tracy R. Slatyer. Revival of the Dark Matter Hypothesis for the Galactic Center Gamma-Ray Excess. *Phys. Rev. Lett.*, 123(24):241101, 2019.
- [41] Rebecca K. Leane and Tracy R. Slatyer. Spurious Point Source Signals in the Galactic Center Excess. *Phys. Rev. Lett.*, 125(12):121105, 2020.
- [42] Rebecca K. Leane and Tracy R. Slatyer. The enigmatic Galactic Center excess: Spurious point sources and signal mismodeling. *Phys. Rev. D*, 102(6):063019, 2020.
- [43] Samuel K. Lee, Mariangela Lisanti, and Benjamin R. Safdi. Distinguishing Dark Matter from Unresolved Point Sources in the Inner Galaxy with Photon Statistics. *JCAP*, 05:056, 2015.
- [44] Samuel K. Lee, Mariangela Lisanti, Benjamin R. Safdi, Tracy R. Slatyer, and Wei Xue. Evidence for Unresolved γ -Ray Point Sources in the Inner Galaxy. *Phys. Rev. Lett.*, 116(5):051103, 2016.
- [45] Florian List, Nicholas L. Rodd, and Geraint F. Lewis. Dim but not entirely dark: Extracting the Galactic Center Excess' source-count distribution with neural nets. 7 2021.
- [46] Florian List, Nicholas L. Rodd, Geraint F. Lewis, and Ishaan Bhat. The GCE in a New Light: Disentangling the γ -ray Sky with Bayesian Graph Convolutional Neural Networks. *Phys. Rev. Lett.*, 125:241102, 2020.
- [47] Oscar Macias, Chris Gordon, Roland M. Crocker, Brendan Coleman, Dylan Paterson, Shunsaku Horiuchi, and Martin Pohl. Galactic bulge preferred over dark matter for the Galactic centre gamma-ray excess. *Nature Astron.*, 2(5):387–392, 2018.
- [48] Oscar Macias, Shunsaku Horiuchi, Manoj Kaplinghat, Chris Gordon, Roland M. Crocker, and David M. Nataf. Strong Evidence that the Galactic Bulge is Shining in Gamma Rays. *JCAP*, 09:042, 2019.
- [49] Siddharth Mishra-Sharma and Kyle Cranmer. A neural simulation-based inference approach for characterizing the Galactic Center γ -ray excess. 10 2021.
- [50] Siddharth Mishra-Sharma, Nicholas L. Rodd, and Benjamin R. Safdi. NPTFit: A code package for Non-Poissonian Template Fitting. *Astron. J.*, 153(6):253, 2017.

- [51] Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. The Structure of cold dark matter halos. *Astrophys. J.*, 462:563–575, 1996.
- [52] Jovana Petrović, Pasquale D. Serpico, and Gabrijele Zaharijas. Millisecond pulsars and the Galactic Center gamma-ray excess: the importance of luminosity function and secondary emission. *JCAP*, 02:023, 2015.
- [53] Harrison Ploeg, Chris Gordon, Roland Crocker, and Oscar Macias. Comparing the Galactic Bulge and Galactic Disk Millisecond Pulsars. *JCAP*, 12:035, 2020.
- [54] Vera C Rubin and W Kent Ford Jr. Rotation of the andromeda nebula from a spectroscopic survey of emission regions. *The Astrophysical Journal*, 159:379, 1970.
- [55] Sarah A. Story, Peter L. Gonthier, and Alice K. Harding. Population synthesis of radio and gamma-ray millisecond pulsars from the Galactic disk. *Astrophys. J.*, 671:713–726, 2007.
- [56] Meng Su, Tracy R Slatyer, and Douglas P Finkbeiner. Giant gamma-ray bubbles from fermi-lat: active galactic nucleus activity or bipolar galactic wind? *The Astrophysical Journal*, 724(2):1044, 2010.
- [57] W. Wang, Z. J. Jiang, and K. S. Cheng. Contribution to diffuse gamma-rays in the Galactic Center region from unresolved millisecond pulsars. *Mon. Not. Roy. Astron. Soc.*, 358:263–269, 2005.
- [58] J. Woltjer. X-rays and type i supernova remnants. *The Astrophysical Journal*, 140:1309–1313, 1964.
- [59] Qiang Yuan and Bing Zhang. Millisecond pulsar interpretation of the Galactic center gamma-ray excess. *JHEAp*, 3-4:1–8, 2014.
- [60] Yi-Ming Zhong, Samuel D. McDermott, Ilias Cholis, and Patrick J. Fox. Testing the Sensitivity of the Galactic Center Excess to the Point Source Mask. *Phys. Rev. Lett.*, 124(23):231103, 2020.