

An unbiased statistic for pulsar parameter estimation via dual-band light curve fitting

AS Seyffert

ORCID: 0000-0002-2033-2999

Thesis accepted in fulfilment of the requirements for the degree
Doctor of Philosophy in Space Physics at the North-West
University

Promoter: Prof C Venter

Co-promoter: Dr AK Harding

Graduation December 2021

20126999

Abstract

The wealth of multiwavelength pulsar data has stimulated the development of emission models that predict light curves (LCs) over multiple wavebands, most notably radio and gamma-ray. Using established statistical methods to fit these model LCs to data can prove ineffectual if the data from one waveband are substantially more precise. This waveband—typically radio—dominates the fit and biases the inferred model parameters. We re-examine the use of Pearson’s chi-squared statistic for joint fits, and introduce a new, derived statistic. Alongside the intuitive reasoning we provide in support of this new statistic, we also construct a formal mathematical framework for pulsar LC fitting, within which our new statistic arises naturally. This framework also provides significant geometric tools by means of which pulsar LC fitting can be better understood intuitively, and the quality of dual-band fits can be evaluated objectively. The core insight that this statistic encodes is that the component single-band chi-squared values (for each waveband) implicitly express goodness of fit in units of the respective LC uncertainties. The resulting implicit weighting that the dual-band chi-squared carries is eliminated by expressing these values in a shared unit before calculating their sum, derived by effectively standardising the scaled pulsar-associated flux across the two wavebands. While our new statistic tends to yield dual-band fits that are not formally good (according to the standard chi-squared statistic), it also tends to yield dual-band fits that better reproduce the broad structure of the observed radio and gamma-ray LCs; this means that the parameter estimates it yield are also better. We use a Monte Carlo method to derive constraints on our parameter estimates, in the form of inclusion contours in the model’s parameter space. Using newly developed quantities to characterise non-colocation of best fits for each band as well as the relative dominance of the respective bands as a function of each band’s precision, we show that chi-squared and our new statistic converge to the same constraints as the precision disparity dissipates. As a first test, we fit two amalgamated dual-band pulsar models to 23 *Fermi* LAT pulsars and compare the resulting parameter constraints to earlier independent results derived using the same data and similar models. Our fits consistently show no radio dominance, and our constraints more strongly correlate with those derived by eye. We next use PSR J2039-5617 as a case study where we perform joint fits on its radio and gamma-ray LCs, the constraints of which can then be used to infer the pulsar mass. Lastly, we explore two more applications of the new statistic: attempting to recover a preliminary trend between macroscopic conductivity vs. spin-down luminosity for the FIDO pulsar model, and performing a joint fit to three datasets (spectrum and brightness / spectral index profiles) associated with pulsar wind nebulae. For the first application, we conclude that spectral data are needed to solidify this trend, in addition to LC data, while for the second, we accept that sub-optimal fits to the available datasets are not due to the statistical fitting method, but rather point to model revision.

Keywords— statistics: model parameter estimation; pulsars: multiwavelength modelling

Contents

	Page
List of Figures	10
List of Tables	11
Glossary	13
Introduction	15
Problem Statement	15
This Work	16
Thesis, Aims and Objectives	16
Document Structure	17
Research Outputs	18
 I BACKGROUND	 19
1 Pulsar Astrophysics	21
1.1 A Brief History of Pulsar Observation	22
1.1.1 The <i>Fermi</i> LAT Era	22
1.2 The Basic Pulsar Picture	24
1.2.1 Pulsar Observables	24
1.2.2 The Formation of Neutron Stars	25
1.2.3 Some General Properties of Pulsars	26
1.2.4 Types of Pulsars	27
1.3 High-energy Emission in Pulsars	28
1.3.1 The Goldreich-Julian Pulsar Model	29
1.3.2 Emission Processes	30
1.4 The Physical Models Relevant to this Work	32
1.4.1 The Polar Cap Model	32
1.4.2 The Slot Gap Model	34
1.4.3 The Outer Gap Model	35
1.5 Overview of Current Models	36
 2 The Geometric Models	 39
2.1 The Single-band Models	39
2.1.1 The Cone Model	39
2.1.2 The γ -ray Models	40
2.2 The Dual-band Models	40

3	Light Curve Fitting	43
3.1	Pearson's χ^2 Statistic	43
3.1.1	Parameter Estimation	44
3.1.2	Example Synthetic Fit	45
3.1.3	The Reduced χ^2 Statistic	47
3.1.4	Dual-band Extension	47
3.2	Real-World Fits	50
3.2.1	Single-band Fits	50
3.2.2	Dual-band Fits	53
3.2.3	Attempts at Eliminating Radio Dominance	56
II	DEVELOPMENT	57
4	Finding a More Appropriate Statistic	59
4.1	The Cause of the Failure	59
4.1.1	Scrutinising the χ_c^2 Statistic using a By-eye Fit	59
4.2	The Scaled Flux	61
4.3	Scaled-flux Normalised GOF	62
4.4	Revisiting the χ_c^2 Statistic	63
4.4.1	Single-band Dominance	64
4.4.2	Implied Model Usefulness	64
4.5	The Scaled-flux Standardised $\Psi_{\Phi,c}^2$ Statistic	65
4.5.1	Scrutinising $\Psi_{\Phi,c}^2$ Using a By-eye Fit	65
5	The Vector Space Framework	67
5.1	Single-band Fits	67
5.1.1	Associating LCs with Vectors	67
5.1.2	Identifying the Best-fit LC	68
5.1.3	The Relationship between S^2 and χ^2	70
5.1.4	Scaled-flux Normalised GOF	72
5.2	Extended to Dual-band Fits	73
5.2.1	Associating LC Pairs with Vectors	74
5.2.2	Identifying the Joint Best-fit LC Pair using χ_c^2	75
5.2.3	Component Perturbations and Dominance	75
5.2.4	The Scaled-flux Standardised $\Psi_{\Phi,c}^2$ Statistic	77
6	Gauging a Statistic's Effectiveness	79
6.1	The Relationship between χ_c^2 and $\Psi_{\Phi,c}^2$	79
6.1.1	Understood Geometrically	81
6.2	Characterising Dual-band Fits	81
6.2.1	The Causative Factors	81
6.2.2	The Magnitude of their Effect	83
6.3	Characterising Dual-band Statistics	84
6.3.1	Precision-determined Priority	85
6.4	The Convergence of χ_c^2 and $\Psi_{\Phi,c}^2$	86

III	IMPLEMENTATION	87
7	Conceptual Design	89
7.1	The Observed LC Pairs	89
7.2	Extracting LCs from Geometric Models	89
7.2.1	Finding the Optimal Amplitude	90
7.2.2	Accounting for Phase Ambiguity	91
7.3	The Sets of Model LC Pairs	92
7.4	Deriving Parameter Constraints	93
7.4.1	A Generalised Procedure	93
7.4.2	The Monte Carlo Algorithm	95
7.4.3	Applied to Fits using $\Psi_{\Phi,c}^2$	96
8	Synthetic Fits	99
8.1	The Set of Synthetic Observed LC Pairs	99
8.1.1	Synthetic Observation	99
8.1.2	Synthetic Non-colocation	99
8.1.3	Synthetic Uncertainty Disparity	100
8.1.4	Coverage of $(\eta_s, \lambda_{r\gamma})$ -space	101
8.2	The Set of Model LC Pairs	101
8.3	Results	102
8.3.1	The Effectiveness of $\Psi_{\Phi,c}^2$	102
8.3.2	The Convergence of χ_c^2 and $\Psi_{\Phi,c}^2$	104
IV	APPLICATION	107
9	Case Study: Dual-band Fits for PSR J2039–5617	109
9.1	Background	109
9.2	The Single- and Dual-band Fits	110
9.3	Results	110
9.3.1	Single-band-only Fits	110
9.3.2	Dual-band Fits	111
9.3.3	Implications for the Orbital Inclination	114
9.4	Discussion	114
10	Dual-band Fits for 23 <i>Fermi</i> LAT 2PC Pulsars	117
10.1	Pulsar Selection	117
10.2	Pulsar Data	119
10.3	By-Eye Fits	120
10.4	Results	121
10.4.1	Single-band-only Fits	121
10.4.2	Dual-band Fits	121
10.4.3	Comparing Parameter Estimates	129
11	Probative Extensions	135
11.1	Dual-band LC Fits using a FIDO γ -ray Model	135
11.2	Joint Fits to Spectral and Spatial Data for Three Pulsar Wind Nebulae	136

Summary and Conclusion	145
A Broader Perspective	146
 APPENDICES	 149
A Best-fit LC pairs for 23 <i>Fermi</i> LAT 2PC pulsars	149
B Maps of $\Psi^2_{\Phi,c}$ for 23 <i>Fermi</i> LAT 2PC Pulsars	163
C The Evolution of the $\Psi^2_{\Phi,c}$ Statistic	173
C.1 Blazing the Trail	173
C.2 Crystalising our Understanding	174
 Bibliography	 178

List of Figures

1	Example: An LC pair identified by eye outperforms the best-fit LC pair identified using χ_c^2	16
1.1	The geometry of a basic pulsar system.	23
1.2	Example: γ -ray and radio LCs for PSR J0659+1414.	24
1.3	The P - $\dot{P}$ diagram.	27
1.4	The standard model of the pulsar magnetosphere.	30
1.5	Schematic representation of a pair formation cascade.	33
1.6	Meridional cross section of an axisymmetric pair formation front.	34
1.7	The structure and location of the outer gap.	35
1.8	The geometry of the current sheet in a striped wind model	36
2.1	Example: The LC pair and radio and γ -ray phaseplots produced by the TPC+Cone dual-band model for $(\alpha, \zeta) = (54^\circ, 45^\circ)$	41
2.2	TPC model gap location.	42
3.1	Example: Ideal single-band fit using the χ^2 statistic.	46
3.2	Example: Ideal dual-band fit using the χ_c^2 statistic.	49
3.3	Example: A single-band fit that is bad but favourable.	50
3.4	The effectiveness of the χ^2 statistic for real-world single-band fits.	51
3.5	How single-band best-fit non-colocation and uncertainty disparity conspire to cause radio dominance in joint, dual-band fits.	52
3.6	The effectiveness of parameter estimation for real-world dual-band fits.	55
4.1	Example: An LC pair identified by eye outperforms the best-fit LC pair identified using χ_c^2	60
5.1	A diagrammatical representation of the n -dimensional geometry associated with LC fitting in $\mathbb{K}^n$	68
5.2	The effective mappings of $\mathbb{K}^n$ onto $\mathbb{R}^n$ under the operators $[\cdot]_{\mathcal{U}}$ and $[\cdot]_{\mathcal{E}}$	71
5.3	The effective mappings of $\mathbb{K}^n$ onto $\mathbb{R}^n$ under the operators $[\cdot]_{\mathcal{E}}$ and $[\cdot]_{\mathcal{E}_\Phi}$	71
5.4	The effective mappings of $\mathbb{K}^{n_c}$ onto $\mathbb{R}^{n_c}$ under the operators $[\cdot]_{\mathcal{E}_c}$ and $[\cdot]_{\mathcal{E}_c^*}$	76
6.1	Example: How dual-band statistics identify best-fit LC pairs, and how to characterise dual-band fits.	80
6.2	Archetypical examples of plots for radio priority (f_r) as function of scaled flux ratio ($\lambda_{r\gamma}$) for four types of dual-band statistic.	85
7.1	Example: Constant- ζ phaseplot cuts that produce the same LC as a result of model symmetry.	92
7.2	Comparing the n -dimensional geometry associated with two single-band fits: one using a valid model and one using an invalid model.	94
7.3	A geometric representation of the quantities relevant to the k th iteration of the Monte Carlo simulation.	96
8.1	Example: The synthetic LC pair (alongside the accompanying radio and γ -ray phaseplots) created by choosing $(\alpha_r, \zeta_r) = (6^\circ, 8^\circ)$ and $(\alpha_\gamma, \zeta_\gamma) = (48^\circ, 36^\circ)$	100
8.2	Single-band scaled flux as function of viewing geometry for the set of synthetic LC pairs.	101

8.3	Joint frequency distributions of f_r and $\lambda_{r\gamma}$ for the synthetic fits.	102
8.4	Precision-determined priority as function of uncertainty disparity for the synthetic fits.	103
8.5	Dual-band best-fit non-colocation (η_c) as function of single-band best-fit non-colocation (η_s) and scaled flux ratio ($\lambda_{r\gamma}$) for the synthetic fits.	104
8.6	Component-wise breakdown of the plots in Figure 8.5.	105
9.1	Single-band best-fit LCs for PSR J2039–5617.	111
9.2	Dual-band best-fit LC pairs for PSR J2039–5617.	112
9.3	Phaseplots associated with the dual-band best-fit LC pairs for PSR J2039–5617.	113
9.4	Joint GOF ($\Psi_{\Phi,c}^2$) as function of viewing geometry for PSR J2039–5617.	114
10.1	Plots of radio priority vs. scaled flux ratio for 23 <i>Fermi</i> LAT pulsars.	128
10.2	Plots of the difference between the parameter point estimates derived by Johnson et al. (2014) and Pierbattista et al. (2015) and those derived here.	129
10.3	Smoothed distributions for the differences $\Delta\alpha$, $\Delta\zeta$, $\Delta \beta $ and $\Delta(\alpha + \zeta)$ for the OG+Cone fits.	131
10.4	Smoothed distributions for the differences $\Delta\alpha$, $\Delta\zeta$, $\Delta \beta $ and $\Delta(\alpha + \zeta)$ for the TPC+Cone fits.	131
10.5	Correlation plots for α , ζ , $ \beta $ and $(\alpha + \zeta)$ for the OG+Cone dual-band fits.	133
10.6	Correlation plots for α , ζ , $ \beta $ and $(\alpha + \zeta)$ for the TPC+Cone dual-band fits.	134
11.1	Plot of γ -ray-band GOF (as measured using $\chi_{\text{sfm},\gamma}^2$) as function of γ -ray-band LC sensitivity for the first round of γ -ray-only FIDO fits.	140
11.2	Plots of best-fit σ vs. $\dot{E}$ for the γ -ray-only and joint FIDO fits.	141
11.3	SED, XPI and XSB fits for PWN 3C 58.	142
11.4	SED, XPI and XSB fits for PWN G21.5–0.9.	143
11.5	SED, XPI and XSB fits for PWN G09+0.1.	144

List of Tables

4.1	Example: Single- and dual-band indirect GOF values for a radio-biased fit.	60
4.2	Example: Single- and dual-band direct GOF values for a radio-biased fit.	66
6.1	Example: Single- and dual-band GOF values and characterising statistic values for a radio-biased fit. . .	83
9.1	Single-band-only GOF values and parameter estimates for PSR J2039–5617.	110
9.2	Dual-band GOF values and parameter estimates for PSR J2039–5617.	111
10.1	Simulation parameters, LC characteristics and previously derived parameter estimates for 23 <i>Fermi</i> LAT pulsars.	118
10.2	Results for the Cone, OG and TPC single-band-only fits.	122
10.3	Results for the OG+Cone dual-band fits.	124
10.4	Results for the TPC+Cone dual-band fits.	125
10.5	Fit characteristics for the dual-band fits.	127
11.1	Single-dataset-only and joint GOF values for PWN 3C 58.	138

Glossary

DFO Deviation from observation. [59](#)

GOF Goodness of fit. [17](#)

LC Light curve. [15](#)

Introduction

To date, the Large Area Telescope (LAT) aboard the *Fermi* Gamma-ray Space Telescope (*FGST*; Atwood et al., 2009, launched in 2008) has detected pulsed γ -ray emission from more than 250 pulsars, the majority of which had previously only been detected as pulsars in the radio or X-ray bands.¹ This crop of γ -ray pulsars, alongside similar detections made by other, contemporaneous instruments (such as the GRID instrument aboard the *AGILE* satellite; Tavani et al., 2009), constitutes a new population of (typically) multi-band objects that is ripe for study. Prior to the launch of the current generation of γ -ray instruments (to which those aboard *AGILE* and *FGST* belong), only 7 γ -ray pulsars had been detected, primarily using the EGRET and COMPTEL instruments aboard the Compton Gamma-ray Observatory (*CGRO*; Gehrels et al., 1993, the predecessor to *FGST*).

This tremendous increase in the number of known γ -ray pulsars (by a factor of almost 40 in a span of just over a decade), has stimulated renewed interest in the study of high-energy pulsar astrophysics, and has been accompanied by the development (or further development) of a host of sophisticated pulsar models (a number of which date back to even before the *CGRO* era). Since the majority of the γ -ray pulsars that have been observed also pulsate in other wavebands (primarily radio), these models are increasingly leveraged to not only account for the pulsed γ -ray emission such a pulsar produces, but also (simultaneously) for the pulsed emission it produces in the other wavebands (e.g., Venter et al., 2009; Weltevrede et al., 2010; Johnson et al., 2014; Pierbattista et al., 2015; Pétri & Mitra, 2020). While such accounts can take many forms, a prominent one is that of a joint, dual-band fit to the (typically phase-matched) radio and γ -ray light curves (LCs) associated with a given pulsar, wherein the model (which takes some combination of parameters) is required to reproduce both LCs simultaneously. Such dual-band fits are typically substantially more constraining than their single-band-only counterparts (i.e., than fits where the model is only required to reproduce one LC; see Section 3.1 for a demonstration of this) since the radio emission likely originates from a different location in the magnetosphere than does the γ -ray emission.

Problem Statement

Unfortunately, as quickly became evident, conducting such fits is often challenging since it is typically the case that a pulsar’s radio LC is accompanied by much smaller relative uncertainties than is its γ -ray LC; i.e., since there is typically a large disparity between the radio and γ -ray uncertainties.² When such fits are conducted using established statistical fitting techniques (e.g., by using Pearson’s χ^2 goodness of fit statistic or by calculating some other joint likelihood statistic; see Chapter 3), the radio-band fit typically dominates the result. This has led to a number of temporary and candidate solutions that range from simply conducting the joint fits by eye (e.g., Venter et al., 2009) to artificially inflating the radio LC uncertainties to resolve the disparity (e.g., Johnson et al., 2014; Pierbattista et al., 2015). While these solutions have yielded a number of excellent results, it can be argued that none of them satisfactorily resolve the radio dominance issue.

Figure 1 shows an example of such a radio-dominated dual-band fit, conducted using the χ_c^2 statistic (the dual-band instantiation of Pearson’s χ^2 statistic; see Section 3.1). The dark grey lines in the *bottom* and *top* panels (respectively) are the observed radio and γ -ray LCs for PSR J1048–5832. These are each accompanied by a set of LC uncertainties, depicted as short, vertical black lines in Figure 1. Notice that there are, indeed, such black lines plotted for the radio-band LC; these are so short that they almost disappear in this plot, however, because the relative uncertainties on the radio LC

¹As reported in the public list of LAT-detected γ -ray pulsars, available at <https://confluence.slac.stanford.edu/display/GLAMCOG/Public+List+of+LAT-Detected+Gamma-Ray+Pulsars> (last accessed 2020-10-05).

²In the interest of clarity we have opted to use the term ‘uncertainty’ instead of ‘error’ when referring to measurement uncertainties.




Figure 1: The joint best-fit LC pair identified using χ_c^2 (plotted in yellow) and an LC pair (identified by eye; plotted in brown) that is arguably a better joint fit to the LC pair observed for PSR J1048–5832. See Section 4.1.1 for a more detailed account of this fit.

intensities (derived from the variability in the off-peak region of the radio LC) are extremely small. The pair of LCs (or ‘LC pair’) plotted in yellow is the best-fit LC pair identified using χ_c^2 (by minimising its value over the model’s parameter space; see Section 3.1), and the LC pair plotted in brown is the best-fit LC pair identified by eye (see Section 10.3). As is evidenced by how much better the γ -ray-band fit is for the LC pair identified by eye than for the LC pair identified using χ_c^2 (dispite the radio-band LCs being almost identical), the fit conducted using the χ_c^2 statistic is radio-dominated. See Section 4.1.1 for a more detailed account of this fit.

This Work

The problem we intend to address in this work, therefore, is that of radio-dominated joint, dual-band LC fits; or, more generally, that of single-band dominated joint, dual-band LC fits. To do so, we develop and implement the $\Psi_{\Phi,c}^2$ statistic; an alternative to the χ_c^2 statistic (for joint, dual-band LC fits aimed at pulsar parameter estimation) that yields undominated joint dual-band fits even if there is a large uncertainty disparity. Like χ_c^2 , this new statistic is an extension of Pearson’s χ^2 statistic for single-band fits, and is a function of the radio and γ -ray single-band statistic values. Unlike χ_c^2 , it takes into account the effects that give rise to the single-band dominance that sometimes afflicts joint, dual-band fits that are conducted using χ_c^2 (primarily the effect of uncertainty disparity; see Section 3.2.2 for a more detailed account of these effects).

An interesting property of this new statistic is that since it should yield identical results to those of χ_c^2 for fits where there is no uncertainty disparity, it is in effect equivalent to χ_c^2 when there is no uncertainty disparity. More generally, the degree to which the results that these two statistics yield are equivalent tends to increase as uncertainty disparity decreases; i.e., these two statistics converge in their effect as uncertainty disparity decreases.

Notably, this new statistic does not supplant the χ_c^2 statistic entirely. Instead, as will be demonstrated in Chapter 9, these two statistics complement each other in those cases where they are not equivalent: the former is sensitive to the quality of the fine-structure fit while the latter is sensitive to that of the broad-structure fit.

Thesis, Aims and Objectives

The main thesis of this work is that the $\Psi_{\Phi,c}^2$ statistic is a more appropriate extension of Pearson’s χ^2 statistic to joint, dual-band LC fits aimed at pulsar parameter estimation than is χ_c^2 ; one that does not yield single-band dominated parameter estimates when there is a large uncertainty disparity. To demonstrate that this is the case, we aim to:

- Demonstrate that the procedure by which pulsar parameter estimates are typically derived from joint, dual-band LC fits does, indeed, fail to effectively derive meaningful parameter estimates for some joint, dual-band LC fits when it is equipped with the χ_c^2 statistic,
- Construct a framework within which both the cause of this failure and a more appropriate extension of Pearson's χ^2 statistic (i.e., the $\Psi_{\Phi,c}^2$ statistic) may be identified,
- Leverage this framework to develop tools with which to gauge the effectiveness of the parameter estimation procedure with regard to its ability to yield undominated joint, dual-band LC fits,
- Demonstrate that the parameter estimation procedure is rendered more effective when it is equipped with the $\Psi_{\Phi,c}^2$ statistic instead of the χ_c^2 statistic for pulsars where there is a large uncertainty disparity, and
- Demonstrate that the parameter estimation procedure yields exactly the same estimates when equipped with either statistic when there is no uncertainty disparity (i.e., that the χ_c^2 and $\Psi_{\Phi,c}^2$ statistic do, indeed, converge in their effect with regard to parameter estimation).

In order to achieve these aims, our objectives are as follows:

1. To show that Pearson's χ^2 statistic expresses goodness of fit (GOF) in such a manner that the single-band instantiations of it effectively have differing units within the pulsar fitting context; specifically, units determined by the precision associated with the respective observed LCs.
2. To show that the χ_c^2 statistic therefore effectively weights the single-band GOF contributions (to the joint GOF) according to the relative precision of the observed LCs.
3. To identify an alternative statistic (the $\Psi_{\Phi,c}^2$ statistic) that does not weight the single-band GOF contributions.
4. To compare the effectiveness of the parameter estimation procedure when it is equipped with this new statistic to its effectiveness when it is equipped with the χ_c^2 statistic.

Document Structure

In the interest of clarity, this thesis is broken up into four parts:

Part I (titled “**BACKGROUND**”) gives a brief account of pulsar astrophysics in general (in Chapter 1; the context within which joint, dual-band LC fits are conducted) and introduces the various components that are involved in such fits. Specifically, Chapter 2 introduces the geometric models relevant to this work and Chapter 3 gives a brief overview (for both the single-band case and the dual-band case) of how pulsar parameter estimates are derived from LC fits conducted using Pearson's χ^2 statistic.

Part II (titled “**DEVELOPMENT**”) details the development of the $\Psi_{\Phi,c}^2$ statistic and the concepts that surround it. In Chapters 4 and 5 we construct a more robust conceptual and mathematical framework within which to understand pulsar LC fitting in general, and leverage it both to show precisely how and why the parameter estimation procedure fails when it is equipped with the χ_c^2 statistic and to identify a more effective statistic (the $\Psi_{\Phi,c}^2$ statistic). Of particular interest is the work done in the latter of these two chapters, where an exhaustive and rigorous account of this framework is given, which serves as a solid mathematical foundation for the entirety of the work done here. See Appendix C for an account of how the development of this foundation influenced (and greatly enhanced) the understanding of LC fitting presented in this work. In Chapter 6 this framework is leveraged to develop a set of tools with which to test the effectiveness of the parameter estimation procedure when it is equipped with a given dual-band GOF statistic.

Part III (titled “**IMPLEMENTATION**”) discusses the implementation of the χ_c^2 and $\Psi_{\Phi,c}^2$ statistics within the context of pulsar parameter estimation and the framework developed in Part II (specifically, in Chapter 7), and tests this implementation in Chapter 8 by applying the tools developed in Chapter 6 to a set of synthetic dual-band fits as a first test of the effectiveness of the $\Psi_{\Phi,c}^2$ statistic (as compared to that of the χ_c^2 statistic).

Part IV (titled “**APPLICATION**”; in concert with Appendices [A](#) and [B](#)) documents the results of two applications of the $\Psi_{\Phi,c}^2$ statistic: the derivation of parameter estimates for PSR J2039–5617 in Chapter [9](#), as a case study in the application of $\Psi_{\Phi,c}^2$, and the derivation of parameter estimates for 23 *Fermi* LAT pulsars in Chapter [10](#), as a means of testing the effectiveness of $\Psi_{\Phi,c}^2$ for real-world fits. Alongside these two applications, it also documents (in Chapter [11](#); in concert with Appendix [C](#)) the results of two probative extensions of the implementation of the $\Psi_{\Phi,c}^2$ statistic presented in Part [III](#).

These four parts are followed up by a [Summary and Conclusion](#) and a look at possible [Future Work](#).

Research Outputs

To date, the work presented here has led to or contributed to two peer-reviewed papers:

[van Rensburg et al. \(2020\)](#) uses a spatio-leptonic transport code to model the morphologies of three pulsar wind nebulae, and the results are fitted to morphological and spectral data for those nebulae. These fits were conducted using an older, now-defunct incarnation of the $\Psi_{\Phi,c}^2$ statistic. See Section [11.2](#) for a brief description of this work.

[Corongiu et al. \(2020\)](#) reports the detection of radio pulsations for PSR J2039–5617, for which only γ -ray pulsations had previously been detected ([Clark et al., 2020](#)). Alongside this detection, a dual-band fit to the observed radio and γ -ray LC pair is conducted using the $\Psi_{\Phi,c}^2$ statistic, which leads to an estimate for the pulsar mass (derived in concert with the results obtained by [Clark et al.](#)). See Chapter [9](#) for a brief overview of the broader details of this study and a more detailed look at the dual-band fits that form part of it.

In addition, two more focussed papers (that effectively condense the contents of this work) are in process.

- The first of these papers will document the development of the $\Psi_{\Phi,c}^2$ statistic, with specific focus on its effectiveness with regard to the derivation of parameter point estimates (see Section [3.1.1](#)), and will roughly comprise Chapters [4](#), [5](#), [6](#) and [8](#), as well as parts of Chapters [3](#) and [7](#).
- The second of these papers will focus on the effectiveness of the $\Psi_{\Phi,c}^2$ statistic with regard to the derivation of parameter constraints (see Section [3.1.1](#)), and will roughly comprise an expansion of Chapter [10](#) along with parts of Chapters [3](#) and [7](#).

Beyond these papers, this work has also given rise to or formed part of the work contained in three conference proceedings ([Seyffert et al., 2015, 2016](#); [Bezuidenhout et al., 2017](#)), as well as a poster at the *8th International Fermi Symposium* (2018).

Part I

BACKGROUND

Chapter 1

Pulsar Astrophysics

The existence of an astrophysical object composed mainly of neutrons was first proposed by [Baade & Zwicky \(1934\)](#), who coined the term ‘neutron star’ (NS) to describe it. They also proposed that these stars are supported against gravitational collapse by degenerate neutron pressure, in the same way white dwarf stars are supported by degenerate electron pressure, since both electrons and neutrons obey Pauli’s Exclusion Principle. Furthermore, they postulated that these objects would be the remnants of stars destroyed in Type II supernovae. A supernova explosion occurs at the end of a sufficiently massive star’s lifespan, when all the star’s nuclear fuel has been spent ([Padmanabhan, 2001](#), see Section 1.2.3 for a more detailed discussion of the formation and properties of NSs).

The first detailed models of the interior structure of NSs were proposed by [Oppenheimer & Volkoff \(1939\)](#), based on earlier work by [Tolman \(1939\)](#). They based their models on the equation of state of a cold Fermi gas in the context of general relativity. One of the major results from this work was that only NSs less massive than $\sim 3M_{\odot}$ could successfully support themselves against gravitational collapse; the Tolman-Oppenheimer-Volkoff limit. Their work was not taken seriously until after the discovery of pulsars in 1968, but has subsequently become one of the foundational pieces of the general relativistic theory of stellar structure ([Misner et al., 1973](#)).

The small expected size of NSs—the diameter of an NS is 10 to 30 km ([Padmanabhan, 2001](#))—coupled with the then contemporary belief that NSs do not radiate much other than thermal radiation, meant that the chances of identifying a source as an NS were slim. Given that NSs were thought to form in Type II supernovae, astronomers knew roughly where to look though: at the centres of supernova remnants. The first source to be identified as being a supernova from historical records was the Crab Nebula. [Duyvendak \(1942\)](#) and [Mayall & Oort \(1942\)](#) concluded that the Crab Nebula had to be the remnant of the supernova explosion witnessed by Chinese astronomers in 1054 AD, based on its current rate of expansion, location in the sky, and luminosity. It is worth noting here that, while the research was based on the Chinese accounts, other accounts of the same event were later found in Indian, Japanese and Arab historical records. The association of the Crab Nebula with a supernova explosion was acknowledged by [Baade \(1942\)](#), who added that the nova of 1054 AD was a Type II supernova. [Minkowski \(1942\)](#) agreed with this assessment, and suggested that the central star of the Crab Nebula could be the remnant of the exploded star.

Less than two years after the discovery of the first extra-solar X-ray source (Sco X-1) by [Giacconi et al. \(1962\)](#) (thought to be synchrotron radiation from cosmic electrons), [Bowyer et al. \(1964\)](#) reported the detection of a second source in the direction of the Crab Nebula. The claim that an NS had been found could not be made, however, as the source could not be identified as being a point source. NSs were expected to appear as point sources, owing to their very small size compared to other stellar mass objects.

[Wheeler \(1966\)](#) was the first to speculate that a rotating NS (with its magnetic field strongly coupled to the surrounding cloud) may be the engine that drives the luminous emission from the Crab Nebula. [Pacini \(1967, 1968\)](#) also investigated the conversion of rotational energy into magnetic-dipole radiation (and ultimately into particle motions) in a system such as the Crab Nebula. The idea that the rotation of such an object may have important observational effects was soon ratified by the accidental discovery of the first four radio pulsars by a Cambridge research group, headed by Anthony Hewish, and [Gold’s \(1968\)](#) subsequent suggestion that these sources were NSs rotating at the pulsation frequency. The actual discovery was made by dame Jocelyn Bell Burnell, a graduate student at the time, and came from data intended to shed

light on the angular structure of compact radio sources by observing the scintillation caused by the irregular structure of the interplanetary medium. The first pulsar they observed (later labelled CP 1919, or PSR J1921+2153) manifested as a series of pulses with a period of about 1.337 s, and was promptly followed by the discovery of three more pulsars in a preliminary search conducted by [Hewish et al. \(1968\)](#).

[Gold \(1968, 1969\)](#) made his suggestion that pulsars are NSs rotating at the pulsation frequency on the grounds that no other theoretically known astronomical object was capable of such short and consistent periodicities as those observed by [Hewish et al.](#) He further speculated that the radiation originated from the plasma in the corotating magnetosphere. He suggested that NSs spin down as their rotational energy is radiated away by relativistic particles. The connection between supernovae, NSs and pulsars was finally cemented by [Cocke et al. \(1969\)](#) after they observed strong optical pulsations from the central star of the Crab Nebula. This observation was preceded by the detection of non-periodic radio pulses from the star by [Staelin & Reifenstein \(1968\)](#), with the detection of pulsed radiation in X-rays ([Fritz et al., 1969](#)) and low-energy γ -rays ([Hillier et al., 1970](#)) coming soon thereafter.

1.1 A Brief History of Pulsar Observation

The first satellite mission to make a substantial contribution to the number of known pulsars was the Small Astronomy Satellite 1 (or SAS-1; later redubbed the “Uhuru” satellite), which launched in 1970 and observed the sky in X-rays. The fourth (and final) Uhuru catalogue ([Forman et al., 1978](#)) contained a total of 339 new X-ray sources, the majority of which were binary systems containing accretion-powered pulsars (i.e., binary systems wherein an NS is accreting matter stripped from its companion). The SAS-2 instrument was launched in 1972 and was the first satellite devoted to γ -ray astronomy, observing in the 20 MeV to 1 GeV range ([Fichtel et al., 1975](#)). The first detections of γ -ray pulsations from the Crab and Vela pulsars were made using data generated by this instrument, as well as observations of some important non-periodic unidentified sources (one of which would turn out to be Geminga). [Kanbach et al. \(1980\)](#) showed (using COS-B data) that the Vela pulsar’s γ -ray spectrum could be represented by a power law of the form $dN/dE \propto E^{-\alpha}$, with $\alpha = 1.89 \pm 0.06$ for the phase-averaged spectrum. This provided clear evidence for the existence of non-thermal acceleration processes, as thermal processes are characterised by black-body-like spectra.

Confirmation that the magnetic fields of accreting pulsars were as large as is expected from the simple conservation laws (see Section 1.2.3) came soon after. [Truemper et al. \(1978\)](#) inferred a magnetic field strength of 5.3×10^{12} G for HER X-1 based on the observation of a strong line feature at ~ 58 keV in the object’s pulsed X-ray spectrum, which they interpreted as electron-cyclotron emission at the basic frequency of the hot plasma in the rotating NS’s magnetosphere.

The majority of pulsar astrophysics research would be focused on the radio and X-ray bands for the next few decades, with a large number of discoveries being made. Some of the more noteworthy missions during that time include: the Röntgensatellit (ROSAT; [Truemper, 1982](#)), which leveraged its factor ~ 1000 improvement in sensitivity over Uhuru to provide new information on X-ray fluxes for all of the then-known pulsars; the Extreme Ultraviolet Explorer (EUVE; [Malina et al., 1982](#)), which observed several NSs in soft X-rays; and the Rossi X-ray Timing Explorer (RXTE; [Bradt et al., 1993](#)), which studied X-ray binaries.

The next significant γ -ray discoveries would only come after the launch of the EGRET instrument aboard the Compton Gamma Ray Observatory (CGRO; [Kanbach et al., 1989](#)) in 1991, which studied the γ -ray sky between 1991 and 2000. The five new pulsars detected by EGRET (including one weak detection), along with the γ -ray pulsars in the Crab and Vela Nebulae, would constitute the entirety of the known γ -ray pulsar population until the launch of the current generation of γ -ray missions, most notably the Large Area Telescope (LAT) aboard the *Fermi* mission.

1.1.1 The *Fermi* LAT Era

The LAT is the primary instrument on board the *Fermi* Gamma-ray Space Telescope (FGST; [Atwood et al., 2009](#)), launched on 11 June 2008. *Fermi* LAT is an imaging, wide field-of-view, high-energy γ -ray telescope covering the energy range from 20 MeV to > 300 GeV. The scientific objectives addressed by *Fermi* LAT include:

- Understanding the high-energy behaviour of γ -ray bursts.
- Using γ -ray observations as a probe of dark matter.




Figure 1.1: A diagram showing the geometric parameters associated with an individual pulsar: The inclination angle, observer angle, and impact angle. The rotation (Ω) and magnetic (B) axes, as well as the light cylinder (i.e., the cylinder with radius R_{LC}) are shown.

- Using high-energy γ -rays to probe the early universe.
- Understanding the mechanisms of particle acceleration operating in celestial sources, particularly active galactic nuclei, pulsars and supernova remnants.

To achieve these objectives, *Fermi* LAT scans the entire sky every three hours (the “all-sky scanning” mode of observation), taking advantage of its large field of view and high sensitivity (Atwood et al., 2009). One of the important results to date of this all-sky survey is the first *Fermi* LAT catalogue of γ -ray pulsars (1PC; Abdo et al., 2010), produced using the first six months of *Fermi* LAT data. Upon its publication, the number of known γ -ray pulsars increased from (at least; Thompson, 2001) 6 to 46. This catalogue was succeeded by the Second *Fermi* LAT Pulsar Catalogue (2PC; Abdo et al., 2013), which again increased the number of known γ -ray pulsars, to 117. This number continues to grow, however, and has breached the 250 pulsar mark.¹

The planned duration of the *Fermi* mission was five years, with a goal of ten years of operations. At present, more than twelve years have passed since the launch of the *Fermi* satellite. The spacecraft experienced its most significant mechanical failure to date on 16 March 2018, when the drive assembly responsible for orienting one of its two solar arrays began to malfunction. As a result of this failure, the observation strategy by which the continual all-sky survey is conducted had to be somewhat modified (due, in part, to thermal constraints), and support for target-of-opportunity observations was effectively ended.




Figure 1.2: Example γ -ray (top panel) and radio (bottom panel) LCs for PSR J0659+1414 generated from *Fermi* LAT and Nançay Radio Telescope data, respectively. In both cases two rotations are shown, with the start of both profiles aligned according to the position of the centre of the radio peak. This means that the actual lag in phase between the radio and γ -ray peaks is also accurately depicted. The γ -ray profile consists of twenty-five bins with the estimated background emission level indicated by the dashed line. The errors on the γ -ray counts are also shown. Adapted from Weltevrede et al. (2010).

1.2 The Basic Pulsar Picture

A pulsar is a rapidly spinning neutron star (NS) with a (predominantly) dipole magnetic field inclined relative to the rotation axis. In Figure 1.1 a basic representation of a pulsar system is shown. Corotation with the NS is only possible up to a distance R_{LC} from the rotation axis. The boundary between the region where corotation is possible and the region where it is not, is called the light cylinder and is indicated by the two vertical lines in Figure 1.1; at the light cylinder, the corotation speed is the speed of light in vacuum. The most important consequence of this limit on corotation is that magnetic field lines can only close inside the light cylinder. The dashed line in Figure 1.1 shows the shape of the last closed field line.

The two most important parameters associated with a pulsar, in the context of this study, are the magnetic inclination and observer angles. The magnetic inclination angle α is the angle at which the magnetic axis of a pulsar is inclined with respect to its rotation axis. The observer angle ζ is the angle between the observer's line of sight and the rotation axis of the pulsar. These two quantities together comprise the 'geometric parameters' associated with a pulsar. A third quantity can be derived from these two angles: the impact angle $\beta = \zeta - \alpha$. It is also commonly referred to when discussing the geometric parameters of pulsar systems, especially in the context of radio observations (e.g., Weltevrede et al., 2010).

1.2.1 Pulsar Observables

Pulsars are observed as sources of pulsed radiation displaying distinct periodicity covering a wide range of energy bands. These observations are usually organised in the form of so-called pulse profiles, or LCs, which relate how the intensity of the emission changes with time during a single rotation of the pulsar. Given a lack of statistics, or to increase stability, these profiles are produced by the stacking of observed emission from numerous rotations of the pulsar. Figure 1.2 shows the radio and γ -ray pulse profiles for one of the pulsars relevant to this study, PSR J0659+1414. The top panel is the γ -ray

¹As reported in the public list of LAT-detected γ -ray pulsars, available at <https://confluence.slac.stanford.edu/display/GLAMCOG/Public+List+of+LAT-Detected+Gamma-Ray+Pulsars> (last accessed 2020-10-05).

LC and the bottom panel is the radio LC. In both panels two rotations of the pulsar are shown for the sake of clarity. By convention, for pulsars with a single-peak radio profile such as the one shown in Figure 1.2, the peak of the radio profile marks the boundary between pulsar rotations; hence, the normalised rotation phase $\phi = 0$ at the radio LC's peak. For pulsars with more complicated radio profiles, the location of $\phi = 0$ is determined via a Fourier transform of the profile. Crucially, the radio and γ -ray profiles plotted in Figure 1.2 (and used throughout) are phase-matched; i.e., the actual radio-to- γ phase lag of the observed peaks is preserved in Figure 1.2.

Alongside LCs it is also possible to obtain both phase-averaged and phase-resolved spectra of the observed emission in several energy bands. These spectra are useful in the study of the magnetospheric environment of the pulsar (specifically, the acceleration mechanism), as well as providing insight into the radiation mechanisms that are responsible for the observed emission. Polarisation observations are also made in some energy bands, with the most relevant example being the observations used by [Weltevrede et al. \(2010\)](#) to obtain constraints on the geometric parameters of five of the 23 *Fermi* LAT pulsars considered in this study.

1.2.2 The Formation of Neutron Stars

An NS is the remnant of a star that has destroyed itself in a Type II supernova explosion. Type II supernovae are the result of stellar core collapse in massive stars that have not yet shed their outer hydrogen layer. They differ from Type Ib and Type Ic supernovae, and are characterised by the presence of hydrogen absorption lines in their spectra. The latter two subtypes are also the result of stellar core collapse, but are associated with stars that have already shed their hydrogen layer before the explosion (in Type Ic supernovae the progenitor star has also shed its helium layer before the explosion). Type Ia supernovae occur when white dwarfs in binary systems accrete enough matter to breach the Chandrasekhar limiting mass of about $1.44M_{\odot}$, beyond which the star can no longer support itself via electron degeneracy pressure against gravitational collapse ([Padmanabhan, 2001](#)). What is left of the star after a Type II explosion is either an NS or a black hole, depending on the mass of the progenitor star: Stars between approximately $8M_{\odot}$ and $30M_{\odot}$ leave an NS behind, while more massive stars leave a black hole behind.

Type II supernova explosions occur as follows. For most of its life, a star is supported against gravitational collapse by radiation pressure. The energy necessary to maintain this pressure is generated by nuclear fusion reactions at the star's core. Initially the fuel for these reactions is hydrogen, but as the star gets older it is forced to burn increasingly heavier elements. In stars lighter than $\sim 8M_{\odot}$ this process only continues to helium burning, and the star ends its life as a white dwarf; supported against gravitational collapse by electron degeneracy pressure and with a core composed mainly of carbon and oxygen. In heavier stars, however, successively heavier elements will be burned until the star's core is mainly composed of iron. The endothermic nature of the iron burning process means that no further energy production is possible, and that the star is unable to support itself against gravitational collapse. The star's core subsequently collapses, with a large amount of energy being released in the form of neutrinos as the process of electron capture (or inverse β -decay) becomes important. A $20M_{\odot}$ star will have a photon luminosity of 4.4×10^{38} ergs/s and a neutrino luminosity of 3.1×10^{45} ergs/s during this phase. The stiffening of the collapsing core material, as Pauli's Exclusion Principle becomes important (and the strong force becomes repulsive), may lead to a 'hydrodynamic bounce'. The resultant outwardly propagating shock blows the outer layers of the star away in a spectacular supernova explosion. In cases where the neutron degeneracy pressure at the core provides insufficient support for the remaining mass, gravitational collapse continues and a black hole is formed. If this is not the case, though, electron capture continues and nuclei disintegrate into nucleons and electrons. The resulting object is an NS and is supported against gravitational collapse by neutron degeneracy pressure ([Carroll & Ostlie, 1996](#)).

Neutron Star Masses and Equations of State Unfortunately, the exact upper limit on the mass of an NS is still uncertain. While a number of models for the interiors of pulsars have been developed (see, e.g., [Baym et al., 2018](#), typically centering on the equations of state that describe the conditions inside the NS), they tend to disagree on the question of what this upper limit is. Neutron star masses inferred from observation give lower bounds on what this upper limit might be, and in so doing serve as a constraint on the models for pulsar interiors. To date, the most massive NS observed is likely PSR J0740+6620, with a mass of $2.14^{+0.10}_{-0.09} M_{\odot}$ ([Cromartie et al., 2019](#)).

1.2.3 Some General Properties of Pulsars

If it is assumed that the stellar core rotates more or less rigidly during the gravitational collapse, and that its angular momentum is conserved, then the angular speed Ω of the core after the collapse will be

$$\Omega_f \sim \Omega_i \left(\frac{R_i}{R_f} \right)^2, \quad (1.1)$$

where R is the star's radius, and the subscripts “i” and “f” denote the initial (before the collapse) and final (after the collapse) values, respectively. For typical values of $R_i \sim 10^{11}$ cm and $R_f \sim 10^6$ cm, this amounts to an increase in angular speed by a factor of $\sim 10^{10}$, yielding rotational periods (P) well within the millisecond to second range. The magnetic field is similarly amplified during the collapse under the assumption of the conservation of magnetic flux, with the magnitude of B after the collapse being

$$B_f \sim B_i \left(\frac{R_i}{R_f} \right)^2. \quad (1.2)$$

This relation yields a typical magnitude of $\sim 10^{12}$ G for NSs. Stellar core collapse as described above thus leads to a rapidly spinning, highly magnetised NS. These strong magnetic fields play a very important role in the production of radiation in pulsar magnetospheres (see Section 1.3).

The magnitude of the magnetic field at the pole (B_0) in terms of the pulsar's rotational period P and its time derivative $\dot{P}$ can be derived by equating the rate at which the pulsar's rotation is slowing down and the vacuum magnetic dipole radiation loss rate (Ostriker & Gunn, 1969):

$$\dot{E}_{\text{rot}} = \frac{d}{dt} \left(\frac{1}{2} I \Omega^2 \right) = I \Omega \dot{\Omega} = -\frac{2}{3c^3} \mu^2 \Omega^4 \sin^2 \alpha, \quad (1.3)$$

where $\mu = B_0 R^3 / 2$ is the magnetic moment, $I \sim MR^2$ the moment of inertia and c the speed of light in a vacuum. Solving for B_0 and inserting typical values of $I = 10^{45}$ g.cm², $R = 10^6$ cm, and $\alpha \sim 90^\circ$ yields

$$B_0 = \frac{\sqrt{6c^3 I}}{2\pi R^3 \sin \alpha} \cdot \sqrt{P\dot{P}} \simeq 6.4 \times 10^{19} \sqrt{P\dot{P}}, \quad (1.4)$$

in Gaussian units.

To put this relation into the proper context, it is useful to consider how constant- B_0 lines manifest themselves on a P - $\dot{P}$ diagram such as that shown in Figure 1.3. The P - $\dot{P}$ diagram is indicative of pulsar population properties and evolution. By squaring both sides of the relation obtained in Eq. (1.4) and taking the logarithm (base 10), one obtains an expression for the shape of the constant- B_0 lines:

$$\log(\dot{P}) = -\log(P) + C_1, \quad (1.5)$$

where $C_1 = 2 \log(B_0 / (6.4 \times 10^{19}))$ is a constant if B_0 is a constant. Lines of constant B_0 are thus straight lines with negative slope on the P - $\dot{P}$ diagram. In Figure 1.3 it is evident that the estimation made for the value of B_0 from Eq. (1.2) agrees well with the values derived from observations for the canonical pulsars using Eq. (1.4) (the largest clump in Figure 1.3, see Section 1.2.4 for more details).

Another quantity that can be derived from Eq. (1.3), is the characteristic age τ . If it is assumed that the component of μ that is perpendicular to the pulsar's rotation axis (i.e., $\mu_\perp = \mu \sin \alpha$) remains constant throughout the lifetime of the pulsar, then the quantity $P\dot{P}$ also remains constant (from Eq. (1.4)). Rewriting the identity $P\dot{P} = P\dot{P}$ as $PdP = P\dot{P}dt$, and integrating over the lifetime of the pulsar, the rotational age is (assuming that $P_0 \ll P$, where P_0 is the initial period)

$$\tau = \frac{(P^2 - P_0^2)}{2P\dot{P}} \simeq \frac{P}{2\dot{P}}. \quad (1.6)$$

Predictions made using this relation are in good agreement with the estimated or known age of the pulsar in many cases; e.g., the Crab pulsar, where it predicts an age of ~ 1257 yrs as compared to the known age of 967 yrs (Duyvendak, 1942;




Figure 1.3: A plot of $\dot{P}$ vs. P for all the pulsars in the ATNF pulsar catalogue at the time of writing (<http://www.atnf.csiro.au/people/pulsar/psrcat/>). Lines of constant B_0 (solid), τ (dashed), and $\dot{E}$ (dotted), as well as the so-called death line, are shown. The death line delimits the border of the so-called ‘pulsar graveyard,’ where $\dot{E}$ is too low for pair production to occur (see Section 1.3.2). For a discussion of each of the various types of pulsar indicated on this plot, see Section 1.2.4.

Mayall & Oort, 1942).² In a similar fashion to that by which Eq. (1.5) was constructed, it can be shown that for lines of constant age,

$$\log(\dot{P}) = \log(P) - C_2, \quad (1.7)$$

where $C_2 = \log(2\tau)$. Therefore, lines of constant age are straight and have a positive slope on the P - $\dot{P}$ diagram. Thus, under the assumptions made above, the evolutionary track of a pulsar on the P - $\dot{P}$ diagram roughly proceeds along a line of constant B_0 toward the lower right of the diagram, with the cessation of emission from the pulsar occurring once it crosses the death line. It should be noted, however, that this analysis only applies to rotation-powered pulsars (i.e., pulsars for which Eq. (1.3) holds), and consequently does not account for the properties of accretion-powered pulsars, including millisecond pulsars (MSPs) that are still being spun up.

1.2.4 Types of Pulsars

As the number of known pulsars grew, different populations of pulsars emerged with distinct characteristics. The distinction which is probably most obvious from the P - $\dot{P}$ diagram in Figure 1.3 is that between canonical pulsars and MSPs. The former of these two groups constitutes the majority of pulsars in the large cluster in the middle of Figure 1.3, while the latter constitutes the majority of pulsars in the smaller cluster in the bottom left corner. As the name of the latter group suggests, the distinction is made on the basis of rotational period, with pulsars that have periods below approximately

²As computed at the time of writing using the values for P ($= 0.0333924123$ s) and $\dot{P}$ ($= 4.20972 \times 10^{-13}$ s.s⁻¹) listed in the ATNF pulsar catalogue for the Crab pulsar (PSR J0534+2200).

50 ms being classified as MSPs. These more rapidly rotating pulsars are generally older, and have weaker spin-down rates (lower $\dot{P}$). MSPs attain their high rotational velocities through a spin-up process involving a binary companion (Alpar et al., 1982; Bhattacharya & van den Heuvel, 1991). The majority of pulsars start their lives as canonical pulsars (in the large cluster) from where their evolution takes them toward the lower right of the P - $\dot{P}$ diagram, where they eventually cross the death line. For some pulsars in binary systems this evolution is interrupted when the pulsar's companion expands enough to trespass outside of its own Roche lobe, resulting in the accretion of matter from the companion onto the NS's surface. The accretion imparts additional angular momentum onto the pulsar and it is spun up to very short periods. This accretion process not only affects the period of the pulsar, but may also affect the magnetic field, significantly reducing its strength (MSPs have magnetic fields of $\sim 10^8$ G), although this is yet to be confirmed. The short periods and low $\dot{P}$ values associated with MSPs mean that they lose their angular momentum much more slowly than canonical pulsars (via magnetic dipole radiation), and are consequently visible for much longer timespan. The pulsars in the bridge region between the canonical and MSP clusters are pulsars either undergoing spin-up, or pulsars in systems where accretion has ceased. The supernova death of the binary companion can result in the pulsar becoming an isolated star, either by the companion completely destroying itself, or by a fatal disruption of the binary system, possibly explaining the existence of isolated MSPs (Shapiro & Teukolsky, 1983). Other types of MSPs may ablate their companions to the point where they become isolated (Fruchter et al., 1990).

Beyond their periods and the strength of their magnetic fields, pulsars can also be classified by the energy band, or bands, in which they are visible. The simplest example of such a group (and the largest) is that of the radio pulsars (shown in blue). Radio pulsars make up the majority of the known (and of the canonical) pulsars; there are currently nearly 2900 pulsars listed in the ATNF Pulsar Catalogue. These pulsars are usually found by performing Fourier-space searches in large sets of radio data, and display a wide variety of pulse profile shapes.

Another such group is the γ -ray pulsars, with the majority of its members being discovered by using radio periods of known pulsars to fold data obtained from the *Fermi* LAT instrument. These pulsars display a wide variety of pulse profiles, ranging from sharp two-peak profiles to broad single-peak profiles. These pulsars also form part of a larger group: the high-energy pulsars (which also includes those pulsars visible in X-rays). The emission from these pulsars is predominantly curvature radiation (CR) and synchrotron radiation (SR). In Figure 1.3 the red filled triangles denote the high-energy pulsars.

The last group visible in Figure 1.3 (the green circles) is the binary pulsars. In most cases the systems in this class contain an NS and a normal star, but some instances of NS-NS binaries have been observed. Two important types of NS-normal-star binary system have been identified in recent years: black widow and redback pulsars. In both systems, an MSP is ablating material from its binary companion (via the NS's stellar wind and high-energy radiation). The main difference between the two is that black widow pulsars typically have lighter, degenerate stars as companions, which tend to be heavily ablated, and redback pulsars typically have heavier (typically an order of magnitude heavier than the companions to black widows) non-degenerate stars as companions which tend to be less heavily ablated. Binary systems composed of two compact objects have proved useful in tests of general relativity, and are thought to potentially be some of the strongest gravitational wave sources in the universe.

Another distinction may be drawn between rotation-powered pulsars and accretion-powered pulsars, but transitional pulsars occur when a pulsar changes from one state to the other. Magnetars are another special class of pulsar with exceedingly large magnetic fields of $\sim 10^{14}$ – 10^{15} G; the radiation from these are believed to be powered by magnetic field decay, since their luminosities far exceed the available spin-down luminosity.

For more information, refer to the review articles by Harding (2013) and Kaspi (2018).

1.3 High-energy Emission in Pulsars

In this section the 'standard model' of the pulsar magnetosphere is presented, as well as an account of the important emission mechanisms involved in the generation of γ -rays in pulsar magnetospheres.

1.3.1 The Goldreich-Julian Pulsar Model

Following the suggestion made by Gold (1968) (that the then newly-discovered objects subsequently known as pulsars were rotating magnetic NSs), Goldreich & Julian (1969) investigated the simple case of an aligned rotator. Goldreich & Julian concluded that a rotating magnetic NS cannot be surrounded by a vacuum.

Consider, as Goldreich & Julian did, an NS with a dipole external magnetic field continuous at the stellar surface. Assume that the pulsar is a uniformly magnetised, perfectly conducting sphere and that its interior field is $\mathbf{B}_{\text{in}} = B_0 \mathbf{e}_z \parallel \boldsymbol{\mu}$. Since the spin and magnetic axes are aligned, the external field is static. This means that charged particles on the surface of the star will experience a Lorentz force due to the rotation of the pulsar. Under the assumption that the star is a perfect conductor, charge redistribution must take place in order to establish a counterbalancing electric field $\mathbf{E}$ (due to charge separation). Since $\mathbf{E}_{\text{in}} \cdot \mathbf{B}_{\text{in}} = 0$ for a conductor, no permanent current flow exists inside the star. Therefore, the rotational form of Ohm's law applies; i.e.,

$$\mathbf{E}_{\text{in}} = -\frac{1}{c} [(\boldsymbol{\Omega} \times \mathbf{r}) \times \mathbf{B}_{\text{in}}] = -\frac{\Omega B_0 r \sin \theta}{c} (\sin \theta \mathbf{e}_r + \cos \theta \mathbf{e}_\theta), \quad (1.8)$$

where r and θ (and ϕ) are spherical coordinates centred on the NS. This interior electric field obeys the condition $\nabla \times \mathbf{E}_{\text{in}} = 0$, and solving Laplace's equation yields an external electrostatic potential (assuming that $\mathbf{E}$ is continuous at the stellar surface)

$$\Phi = -\frac{B_0 \Omega R^5}{3cr^3} P_2(\cos \theta), \quad (1.9)$$

where P_2 is the Legendre polynomial of second degree. In the case where the pulsar is surrounded by a vacuum, this potential implies an external value for the Lorentz invariant

$$\mathbf{E} \cdot \mathbf{B} = -\left(\frac{\Omega R}{c}\right) \left(\frac{R}{r}\right)^7 B_0^2 \cos^3 \theta \neq 0. \quad (1.10)$$

Since this quantity vanishes in the stellar interior, it must change continuously from its interior value to its exterior value within the surface charge layer. As the exterior value is not zero, the electric field has a non-zero component $\mathbf{E}_{\parallel}$ parallel to the magnetic field at the surface of the NS. For typical pulsar parameters the force associated with this component is many orders of magnitude larger than the gravitational binding force acting on charged particles at the surface ($\sim 10^8$ times larger for a proton). Consequently the particles at the stellar surface cannot be in dynamical equilibrium, and are pulled from the surface of the NS.

Goldreich & Julian further showed that the charge density in the co-rotating magnetosphere is

$$\rho_{\text{GJ}} = \frac{1}{4\pi} \nabla \cdot \mathbf{E} \simeq -\frac{1}{2\pi c} \boldsymbol{\Omega} \cdot \mathbf{B}; \quad (1.11)$$

also known as the Goldreich-Julian charge density. This relation implies the existence of a null-charge surface, where the local charge density is zero, which is the locus of all the points inside the corotating region where $\boldsymbol{\Omega} \cdot \mathbf{B} = 0$, and separates the magnetosphere into a positively-charged and a negatively-charged region. In the case of an aligned rotator there are two negatively charged regions (one above each polar cap) and one positively charged region (around the star's equator; see Figure 1.4). Goldreich & Julian also found that no particles can corotate beyond the light cylinder (with radius $R_{\text{LC}} = c/\Omega$; as measured from the rotation axis) as the speed of a corotating particle outside this cylinder would necessarily exceed the speed of light. There will therefore be a net outflow of particles through the surface of this cylinder along magnetic field lines that do not close inside it. The magnetic field line that crosses the null-charge surface at the light cylinder is called the critical field line. The resulting picture of the magnetosphere is shown in Figure 1.4.

Implications for Emission

In this model, $\mathbf{E} \cdot \mathbf{B} = 0$ everywhere due to the presence of the conductive plasma (the so-called force-free or ideal magnetohydrodynamic assumption), which means that no particle acceleration due to a non-zero $\mathbf{E}_{\parallel}$ will take place, and consequently no radiation will be produced. However, pulsar emission resulting from such acceleration is observed (e.g., Kanbach et al., 1980), which suggests that this model is perhaps too simplified. As will be illustrated in the following




Figure 1.4: Schematic representation of the magnetosphere of an aligned rotator. The NS is at the lower left, and the star's angular velocity vector is directed upward, parallel to its magnetic axis. The polar cap angle is denoted by θ_0 , and corresponds to the location of the footpoints of the last open field lines. The dashed vertical line indicates the light cylinder, while the slanted dashed line indicates the null-charge surface, where $\Omega \cdot \mathbf{B} = 0$. The shaded region, within which all magnetic field lines are closed, is called the corotating magnetosphere and is filled with a plasma of charged particles that corotate with the pulsar. Particles stream out through the light cylinder along the open magnetic field lines into the wind zone, with field lines above the critical field line carrying negative charges, and field lines below it carrying positive charges. The charge density inside the light cylinder is described by Eq. (1.11). From Goldreich & Julian (1969).

sections, models addressing this issue generally have an acceleration gap located somewhere in their magnetosphere, where $\mathbf{E} \cdot \mathbf{B} \neq 0$. The geometric models presented in Chapter 2 (which are used in this study) are based on these physical models.

While these geometric models are only concerned with the locations of these gaps, it is important to understand how these gaps come about in the physical models so as to better motivate the assumptions made in constructing the geometric models.

1.3.2 Emission Processes

Synchrotron Radiation

Charged particles moving relativistically through a magnetic field will emit synchrotron radiation (SR). This process is the analogue of cyclotron radiation in the relativistic limit. Neglecting radiation losses, the equations of motion for a particle moving through a uniform magnetic field $\mathbf{B}$ are (Rybicki & Lightman, 1979):

$$\frac{d}{dt}(\gamma m \mathbf{v}) = \frac{q}{c} (\mathbf{v} \times \mathbf{B}) \quad (1.12)$$

$$\frac{d}{dt}(\gamma m c^2) = q (\mathbf{v} \cdot \mathbf{B}) = 0, \quad (1.13)$$

where $\mathbf{v} = \beta c$ is the particle's velocity, q is its charge, m is its mass and $\gamma = (1 - v^2/c^2)^{-1/2}$ is the Lorentz factor associated with its motion. From the latter of these two equations it follows that the particle's γ is constant (and thus too its speed), while the former implies that the particle's acceleration is always perpendicular to its velocity. Therefore, the particle will follow a helical path with axis parallel to the magnetic field. The associated gyration frequency of

$$\omega_B = \frac{qB}{\gamma mc}, \quad (1.14)$$

is lower (by a factor of $1/\gamma$) than it would be in the non-relativistic case (Rybicki & Lightman, 1979).

Since the radiation emitted by a rapidly moving particle is strongly beamed in its direction of motion (the opening angle of the beam is $\sim 1/\gamma$), the particle's emission sweeps out a cone with angle $\alpha = \arctan(v_\perp/v_\parallel)$. Therefore, emission will only be directed towards a fixed observer once per gyration, and the observed pulses will be much smaller than the gyration period (by a factor of γ^3). The radiation spectrum emitted by a single electron in this way is

$$\frac{dP}{d\omega} \equiv P(\omega) = \frac{\sqrt{3}}{2\pi} q^2 \left(\frac{\gamma}{\rho}\right) F\left(\frac{\omega}{\omega_c}\right), \quad (1.15)$$

where $\rho = \gamma mc^2/(qB_\perp)$ is the radius of curvature of the helix, $\omega_c = 3c\gamma^3/2\rho$ is the cut-off frequency (below which this spectrum follows a power law) and $F(x)$ is the synchrotron function, which has the following asymptotic forms (Erber, 1966)

$$F(x) \sim \begin{cases} \frac{4\pi}{\sqrt{3}\Gamma(\frac{1}{3})} \left(\frac{x}{2}\right)^{1/3}, & x \ll 1 \\ \left(\frac{\pi}{2}\right)^{1/2} e^{-x} x^{1/2}, & x \gg 1. \end{cases} \quad (1.16)$$

For a single electron the corresponding photon spectrum is well known: It extends as a power law of index $-2/3$ up to ω_c , beyond which it drops off exponentially in intensity. For the case where the volume density of particles per energy unit can be described by a power law with index p (i.e., where $dN/dE \propto E^{-p}$), the integrated power spectrum is also a power law $P(\omega) \propto \omega^{-s}$, with $s = (p - 1)/2$.

Curvature Radiation

In the limit where the particles are moving along the magnetic field (e.g., in the pulsar's co-rotating frame), with no component of velocity perpendicular to it, the radiation produced is somewhat different. The radiation produced by such particles is called curvature radiation (CR), and is associated with changes in longitudinal energy with respect to the magnetic field, as opposed to changes in transverse energy in the case of SR. This process is especially important for pulsars, as the particles liberated from the stellar surface (the primary particles; see Section 1.3.1), and subsequently accelerated along the magnetic field lines, have negligible momentum perpendicular to the magnetic field (Sturrock, 1971). These particles emit CR in a cone of opening angle $\sim 1/\gamma$ centred on the magnetic field line, which means that under the extreme acceleration which occurs in NS magnetospheres, this CR is emitted close to tangent to the magnetic field line. This fact is important when constructing the geometric models (see Chapter 2).

The power spectrum for this CR is the same as that for SR for a single particle, but with the modification that ρ is no longer the gyration radius, but the radius of curvature of the particle trajectory in the lab frame (including the effects of magnetic field geometry and $\mathbf{E} \times \mathbf{B}$ drift, since the pulsar is rotating). For a power-law population the resulting spectral index is, however, not the same as for SR due to the curvature radius typically being so much larger than the gyration radius. The spectral index for CR for a power law population is $s = (p - 1)/3$ (e.g.,; Rudak & Dyks, 1999), and this spectrum again only extends up to a cut-off frequency. It is important to note, however, that this cut-off frequency is much higher for CR than it is for SR due to its dependence on ρ .

More recent models include the generalised synchro-curvature process that encapsulates both limits (perpendicular and parallel particle motion); Viganò et al. (e.g., 2015)

Magnetic (One-photon) Pair Production

The photons produced by the radiation processes discussed above may interact with the strong magnetic field of the pulsar to produce electron-positron pairs if they are sufficiently energetic. The probability that a photon with a frequency ν (i.e., having energy $h\nu \gtrsim 2m_e c^2$, where m_e is the mass of an electron) will interact with the magnetic field to produce an electron-positron pair, can be expressed as a photon attenuation coefficient $\alpha(\chi)$. This coefficient determines the actual number of pairs n_p created for a photon path length d in a magnetic field B through the relation (Erber, 1966)

$$n_p = n_\gamma \left(1 - e^{-\alpha(\chi)d}\right) \simeq n_\gamma \alpha(\chi)d, \quad (1.17)$$

where n_γ is the photon number density and the parameter $\chi = \frac{1}{2}(h\nu/m_e c^2)(B_\perp/B_{\text{cr}})$ is the Erber parameter, with $B_\perp$ the component of the magnetic field perpendicular to the particle's direction of motion and $B_{\text{cr}} = m_e^2 c^3 / e\hbar = 4.414 \times 10^{13}$ G the critical magnetic field in which the gyro energy of an electron (or positron) would be equal to its rest mass (Daugherty & Harding, 1983). This coefficient, which corresponds to the photon absorption per unit length when considering photon propagation perpendicular to a uniform, unbounded magnetic field, is given by (Erber, 1966)

$$\alpha(\chi) = \frac{1}{2} \left(\frac{\alpha_f}{\lambda_c} \right) \left(\frac{B_\perp}{B_{\text{cr}}} \right) T(\chi), \quad (1.18)$$

where $\alpha_f = e^2/\hbar c \approx 1/137$ is the fine structure constant, λ_c is the Compton wavelength and $T(\chi)$ is a dimensionless auxiliary function with the following asymptotic forms:

$$T(\chi) = \begin{cases} 0.46 e^{-\frac{4}{3\chi}}, & \chi \ll 1 \\ 0.60 (\chi)^{-1/3}, & \chi \gg 1. \end{cases} \quad (1.19)$$

Using the approximation for $T(\chi)$ where $\chi \ll 1$ (which is appropriate for pulsars), Eq. (1.18) becomes (Luo et al., 2000)

$$\alpha(\chi) = 0.46 \left(\frac{\alpha_f}{\lambda_c} \right) \left(\frac{\chi}{\epsilon_\gamma} \right) e^{-\frac{4}{3\chi}}, \quad (1.20)$$

where ϵ_γ is the photon energy. Thus, $\alpha(\chi) \sim \chi e^{-\frac{1}{\chi}}$, which means that as $B_\perp$ is increased (and therefore χ is increased), the magnetic field opacity increases for a constant ϵ_γ . This implies that sufficiently strong magnetic fields are in principle completely efficient pair producers for photons of arbitrarily high energy. In addition to the one-photon (magnetic) pair production process, the two-photon pair production process is important for outer gap models (details can be found in Cheng et al., 1986b)

1.4 The Physical Models Relevant to this Work

In this section the physical models that give rise to the geometric γ -ray models we use in this work are discussed. Specifically, the polar cap (PC), slot gap (SG) and outer gap (OG) models for the production of γ -rays are presented, with special emphasis on the details that are most relevant to the geometric models used in this study (introduced in Chapter 2).

1.4.1 The Polar Cap Model

In the polar cap (PC) model of Daugherty & Harding (1982), charged particles are liberated from the PC surface by the strong electric field component ($E_\parallel$) parallel to the local magnetic field. These particles are accelerated extremely rapidly along the curved magnetic field lines, producing CR very close to the PC surface. If the emitted photons are energetic enough, they will produce $e^\pm$ pairs via single-photon pair production (discussed in Section 1.3.2). The motion of these secondary particles will generally have a component perpendicular to the magnetic field. Therefore, these particles will produce SR, radiating the kinetic energy associated with this perpendicular component away over a very short timescale. The two oppositely charged members of each pair will then be accelerated in opposite directions along the magnetic field lines by $E_\parallel$. The stream of inwardly accelerated particles produced in this way (typically positrons for $\alpha < 90^\circ$) heats the PC and drives further thermionic emission of charged particles from the stellar surface (Harding & Muslimov, 2003). The outwardly accelerated particles also produce CR, which leads to the production of yet more $e^\pm$ pairs, and so forth. Thus, a single liberated primary particle leads to a cascade of charged particles as depicted in Figure 1.5.

The plasma of charged particles produced in this way increases in number density as one moves farther away from the stellar surface, and is sufficiently dense to screen $E_\parallel$ above the pair formation front (PFF). The resulting magnetosphere is one in which particles are accelerated close to the PC, with a pair-plasma coasting outward relativistically above the PFF. Therefore, in the PC model the emission responsible for the observed LCs originates below the PFF.




Figure 1.5: Schematic representation by Daugherty & Harding (1982) of a pair photon cascade resulting from the acceleration of a primary electron along the curved magnetic field of an NS. Four generations of photons, including both CR photons (γ_{CR}) and SR photons (γ_{SR}) are shown. The kinetic energy of the produced particles associated with motion perpendicular to the magnetic field is almost instantaneously radiated in the form of SR due to the strong magnetic field.

1.4.2 The Slot Gap Model

The slot gap (SG) model (Arons, 1983), which serves as the basis for the geometric two-pole caustic (TPC) model, is in a sense an extension of the PC model. The SG model considers the two-dimensional structure of the PFF; specifically, it considers how its height above the PC depends on colatitude. As one moves closer to the last open field line (in the SG model), the height at which the PFF is encountered increases, until it practically extends all the way to the light cylinder along the last open field line. This is because as a boundary condition, the electric potential is taken to vanish at the last open field lines bordering the co-rotating closed region.




Figure 1.6: Meridional cross-section of an axisymmetric PFF. The solid line indicates the position of the PFF (expressed in units of the height of the PFF directly above the centre of the PC) as a function of distance away from the magnetic axis (expressed as a fraction of the angular distance to the footpoint of the the last open field line). In this figure the left-most axis corresponds to the magnetic axis, and the right-most axis corresponds to the last open field line. In the region above the PFF the electric field has no vertical component (i.e., parallel to the magnetic field), as it is quenched by the charged particle plasma produced in the pair-production cascade. In the region below the PFF this is not the case, and particles can still be accelerated (Arons, 1983).

Consider a pulsar magnetosphere in which the closed field line region is filled with plasma, and particles are streaming out from the NS surface, but where no pair production cascade has occurred. In this picture there exists a region with a non-zero $E_{\parallel}$ above the PC, and a region in which $E_{\parallel}$ is zero (the closed-field-line region). As there should not be any discontinuities in the electric field, it can be inferred that the electric potential along the magnetic field lines inside the open field line region has a maximum somewhere inside this region, and decreases as one moves toward the interface between the two regions, finally becoming zero at the interface itself. This translates into the picture presented by Arons (1983) when considering the effect this boundary condition has on the magnitude of the acceleration experienced by particles in the open field line region. Weaker acceleration at larger θ (due to the weaker $E_{\parallel}$) implies that a longer acceleration distance is needed for a particle to attain an energy sufficient to produce CR photons; photons which are in turn energetic enough to




Figure 1.7: Schematic representation of the positions of the primary (I), secondary (II), and tertiary (III) regions associated with the outer gap (see text for details). The current flow along the magnetic field lines begins at (b) and extends through (a). For a pulsar spinning in the opposite direction with the same magnetic field the signs of all the charges will be reversed (Cheng et al., 1986c).

produce $e^\pm$ pairs and trigger the pair production cascade required to form the PFF. Figure 1.6 shows what this structure looks like, with the PFF closest to the stellar surface directly above the magnetic pole, and increasing in altitude as one moves closer to the last open field line.

1.4.3 The Outer Gap Model

Cheng et al. (1986a,c) postulated a mechanism for the creation and sustenance of significant outer-magnetosphere gaps by assuming a global current flow in the pulsar magnetosphere. They considered a situation where the charge density above the PC was positive (Cheng et al., 1986a), which most closely agrees with the case of an anti-aligned rotator in the context of Eq. (1.11). In their model, negative charges flow out of the inner magnetosphere along the last open field lines (crossing the light cylinder at (a) in Figure 1.7), while positive charges flow out along open field lines closer to the magnetic pole. The outflow of negative charges creates a slab-like region along the last open field line where $\mathbf{E} \cdot \mathbf{B} \neq 0$. This region is called an acceleration gap, as particles can be accelerated by the non-zero $E_\parallel$ inside it. The processes of SR, CR, pair production and inverse Compton radiation lead to an eventual steady-state configuration consisting of the three regions shown in Figure 1.7. The gap itself extends along the last open field line, from the null-charge surface up to the light cylinder (region I in Figure 1.7), and from the last open field line inward toward the magnetic axis (bounded by the ‘boundary layer,’ region II in Figure 1.7). The non-zero $E_\parallel$ inside the gap means that charges will be accelerated along the magnetic field lines. As these primary particles are accelerated along the curved field lines by the large potential in the gap (which can be as large as $\sim 10^{15}$ V; Cheng et al., 1986a), they emit CR in the form of primary γ -rays. These γ -rays subsequently convert (via the two-photon pair production mechanism; either within the gap or in the boundary layer) into $e^\pm$ pairs, which are again accelerated by the non-zero $E_\parallel$. The charged particles produced inside the gap contribute to the population of primary particles, while the particles produced farther away from the gap (i.e., in the boundary layer) are the secondary electrons and positrons. The subsequent separation of these particles (as they are accelerated in opposite directions) eventually sets up a net charge, which terminates the gap by quenching the non-zero $E_\parallel$. This quenching does not occur throughout the gap, though, but rather occurs in a fashion similar to that in which the non-zero $E_\parallel$ near the surface of the NS is quenched in the PC case. The positive curvature of the magnetic field lines, coupled with the assumed




Figure 1.8: The geometry of the current sheet in a striped wind model. The pulsar’s magnetic field axis inclined by $\alpha = 45^\circ$ with respect to its rotation axis. In panel (a) (which depicts a poloidal cross section of the magnetosphere), the thin lines are the magnetic field lines and the thicker black line is the current sheet. Notice that the direction of the magnetic field changes when one moves through the current sheet. In panel (b) (which depicts the equatorial plane cross section of the magnetosphere), only the footprint of the current sheet (the spiral black line) and the direction of the equatorial magnetic field component is shown. Notice that this direction also alternates as one moves through the current sheet (adapted from [Bogovalov, 1999](#)).

tangential emission of the primary γ -rays, means that a produced γ -ray will undergo pair production farther away from the last open field line than where it was produced.

The primary and secondary electrons and positrons produced through pair production have velocity components perpendicular to the local magnetic field, and will therefore produce SR in the form of X-rays ([Cheng et al., 1986c](#)), as well as CR in the form of X-rays and γ -rays. The weaker $\mathbf{E} \cdot \mathbf{B} \neq 0$ of the boundary layer means, however, that the separation of charges is less efficient, with only the fraction needed to establish the gap’s upper charge boundary being separated. Since the primary electrons and positrons move in opposite directions, so will the primary γ -rays and the secondary pairs they create. This means that the secondary γ -rays and X-rays are emitted in crossed beams, which leads to the production of tertiary $e^\pm$ pairs in the tertiary pair-filled magnetosphere well beyond the gap (region **III** in Figure 1.7). Most of these pairs are created with energies far less than what is necessary to produce γ -rays, or even X-rays. The radiation from this region is mainly infrared SR, which floods the entire magnetosphere and plays a very important part in the production of pairs in both regions **I** and **II** through collisions with the primary γ -rays. In addition, these infrared photons could also contribute to the primary γ -ray flux in young pulsars like Vela, through inverse Compton collisions with primary particles inside the gap ([Cheng et al., 1986c](#)). The secondary γ -rays that do not convert to tertiary $e^\pm$ pairs may escape the magnetosphere completely, and make up the bulk of the γ -ray emission that is observed from the pulsar ([Cheng et al., 1986c](#)).

1.5 Overview of Current Models

A key shortcoming of the three models discussed in Section 1.4 is that in each the global magnetosphere is decoupled from the local physics that gives rise to the acceleration gaps, and hence to the γ -ray emission; specifically, the plasma currents present in each of these models do not affect the global magnetosphere structure (contra to what is expected from plasma physics; [Li et al., 2012](#)). To overcome this, more recent γ -ray emission models typically do not assume a particular structure for the $\mathbf{B}$ -field, but attempt to solve for that structure dynamically, using sophisticated magnetohydrodynamics (MHD) or particle-in-cell (PIC) simulations of the pulsar magnetosphere.

One such model is the force-free inside, dissipative outside (or FIDO; [Kalapotharakos et al., 2014](#)) model, which

assumes that the plasma that fills the pulsar's magnetosphere is not infinitely conductive everywhere, in contrast to the Goldreich-Julian model presented in Section 1.3.1, which assumes infinite conductivity everywhere. Specifically, the FIDO model assumes that the conductivity σ' of the plasma is infinite inside the light cylinder and finite outside of it. This effectively means that particle acceleration (and therefore γ -ray emission via CR) can take place outside the light cylinder (in contrast to the Goldreich-Julian model) since $E_{\parallel}$ can be non-zero there. However, while σ' is assumed to be finite everywhere outside the light cylinder, $E_{\parallel}$ is not non-zero everywhere outside the light cylinder. To determine where in the magnetosphere this is the case, the structure of the B -field is solved using MHD. The main region where $E_{\parallel}$ is non-zero, and therefore the region where the γ -ray emission predominantly originates in this model, is along the so-called current sheet, which separates the open volume regions associated with the two magnetic poles.

The current sheet is also important in other MHD models, like the striped wind model (Coroniti, 1990; Bogovalov, 1999), wherein this sheet has a wave-like shape as depicted in Figure 1.8(a)

Another approach that is currently being explored is to implement *ab initio* relativistic PIC simulations of the dynamics of the plasma itself (e.g., Chen & Beloborodov, 2014; Philippov & Spitkovsky, 2014; Philippov et al., 2015). A typical progression in these simulations is to start with a vacuum dipole magnetic field and to then inject plasma particles into the magnetosphere (at the stellar surface, for example) while the NS spins. During each timestep the interaction between each of the plasma particles and the B - and E -fields (or, at least between comparatively small groups of particles and the fields) are modelled as fully as possible (i.e., inclusive of the mechanisms by which the particles emit γ -rays). By breaking the magnetosphere into smaller cells, the complexity of each calculation is somewhat diminished. Across the simulation the change in the B - and E -fields, as well as the motion of the plasma particles, is monitored.

In this MHD model

Chapter 2

The Geometric Models

The γ -ray *emission* models discussed in the previous chapter (which are based on the emission processes described in Section 1.3.2) serve as the bases for the *geometric* γ -ray models that will be employed throughout the rest of this study: the two-pole caustic (TPC; Dyks & Rudak, 2003; Venter et al., 2009) and geometric outer gap (OG) models. These geometric models are not concerned with the underlying processes responsible for γ -ray emission, but are rather concerned with which configuration for the magnetospheric acceleration gaps best fits with the LCs observed for γ -ray pulsars. They represented an interim step in the chronology of model development, and are being superseded by newer global and PIC models (see Section 1.5). Still, they allow us to (1) directly compare our results with those of other authors and (2) test our fitting method in a well-known context.

When coupled with an empirical hollow-cone radio model (which is also geometric in nature; Story et al., 2007; Kijak & Gil, 1998, 2003, hereafter ‘Cone’), these γ -ray models give rise to the OG+Cone and TPC+Cone dual-band models, which aim to produce phase-matched radio and γ -ray LC pairs. In either case, both of the single-band models assume a retarded vacuum dipole magnetic field (Deutsch, 1955), and the conglomeration of the two is effected by the requirement that α must be the same for both single-band models.

2.1 The Single-band Models

2.1.1 The Cone Model

The hollow cone radio model assumes that, at a given frequency ν , the radio emission occurs at a single, period- and period-derivative-dependent altitude (in units of the neutron star radius R_{NS} ; fixed at $R_{\text{NS}} = 10^6$ cm for this study)

$$R_{\text{KG}} = 40 \nu_{\text{GHz}}^{-0.26} \dot{P}_{-15}^{0.07} P^{0.3}, \quad (2.1)$$

where $\nu_{\text{GHz}} = \nu/(10^9 \text{ Hz})$ is the radio emission frequency in GHz and $\dot{P}_{-15} = \dot{P}/(10^{-15} \text{ s.s}^{-1})$ is the pulsar’s period derivative (i.e., $\dot{P}$) in units of $10^{-15} \text{ s.s}^{-1}$ (Kijak & Gil, 1998; Story et al., 2007). The assumed emission intensity profile at this altitude, as a function of magnetic colatitude θ , is an offset Gaussian (the hollow cone) that peaks at $\bar{\theta}$ and reaches 10% of its peak intensity at the last open field line. The location of the peak (i.e., $\bar{\theta}$) is determined by assuming a width for the Gaussian peak (at 10%) of $\delta = 0.18 \rho_{\text{cone}}$, where

$$\rho_{\text{cone}} = 1.24^\circ R_{\text{KG}}^{0.5} P^{-0.5} \quad (2.2)$$

is the half-opening angle of the last open field line at R_{KG} . Given this width, this location is $\bar{\theta} \approx 0.53 \rho_{\text{cone}}$.

For the purposes of this study we fix ν at 1.4 GHz, and, in the interest of simplicity, do not introduce a core beam component as others have done (e.g., Story et al., 2007; Gonthier et al., 2004). The latter of these choices is in no way intended to endorse a cone-only geometry over a core+cone geometry. Given these choices, the Cone model takes three free parameters: P , $\dot{P}$, and α ; and predicts the relative radio intensity per solid angle of the observed emission as a function of rotational phase ϕ_M and the observer angle ζ .

Figure 2.1(a) shows an example of the output this model yields (in the form of a *phaseplot*), for a pulsar with $P = 0.167$ s and $\dot{P} = 1.68 \times 10^{-14}$ s.s $^{-1}$ (the period and period derivative for PSR J0742–2822) and $\alpha = 54^\circ$. The geometric nature of the model is reflected in the fact that this phaseplot maps relative intensity, and not absolute intensity. In practice, this means that each constant- ζ cut through this phaseplot must first be appropriately scaled, and the relevant background level taken into account, before it is fit to an observed LC. Figure 2.1(c) shows an example of such an LC, produced by cutting the phaseplot in panel (a) at $\zeta = 45^\circ$, and normalising its intensity (in the interest of simplicity, the background level is assumed to be 0).

2.1.2 The γ -ray Models

The γ -ray models assume that γ -ray emission occurs inside acceleration gaps present in the typically pair-plasma-filled magnetosphere according to the outer gap and slot gap physical models (see Section 1.4); where the pulsar’s electric field is not fully screened by the pair plasma, and particles can be accelerated to produce curvature radiation. They approximate these emission regions with thin slabs located on or near the last open field line; i.e., on or near the field line that closes at the light cylinder. In the co-rotating frame, these models assume a constant, uniform γ -ray emissivity per unit arclength inside the gaps, and that the γ -ray emission occurs tangentially to the magnetic field lines (e.g., Venter et al., 2009). In the observer frame, time-of-flight delays and the magnetic field geometry lead to the formation of caustics (Morini, 1983; Dyks et al., 2004), as is evident in the phaseplot shown in Figure 2.1(b).

In both the OG and TPC models, the shape and location of the gaps themselves are parameterised using four parameters: two parameters, θ_{in} and θ_{out} , that govern the colatitudinal width and location of the slab (expressed in the open volume coordinates defined in Dyks et al., 2004, and only < 1 for open field lines), and two parameters, R_{min} and R_{max} , that govern its radial extent. The former pair of parameters (θ_{in} and θ_{out}) are typically expressed as a fraction of the azimuth-dependent polar cap angle θ_0 , where the polar cap is the region around the pulsar’s magnetic pole bounded by the locus of the footpoints of the set of last open field lines. Since only open field lines can accelerate particles in the physical γ -ray models, these two parameters are constrained to both be less than θ_0 in the geometric models. For this study we fix these two parameters at $\theta_{\text{in}} = 0.95 \theta_0$ and $\theta_{\text{out}} = 1.00 \theta_0$, and therefore assume a gap width of $0.05 \theta_0$. The latter pair of parameters (R_{min} and R_{max}) are typically expressed as a fraction of R_{LC} . Similar to θ_{in} and θ_{out} , we fix R_{max} across both γ -ray models (at $R_{\text{max}} = 1.2$), but the same cannot be done for R_{min} , since the physical OG model imposes a clear limit on the range of values that this parameter can take. Specifically, since in the OG model, stable gaps can only form where $\mathbf{\Omega} \cdot \mathbf{B} > 0$, i.e., where $\mathbf{B}$ ’s component along $\mathbf{\Omega}$ is positive, the gap can only extend as far down as the null-charge surface in the geometric model, where $\mathbf{\Omega} \cdot \mathbf{B} = 0$. For the TPC model, no such constraint on R_{min} exists other than that the gap can only extend as far down as the neutron star’s surface (i.e., $R_{\text{min}} \geq R_{\text{NS}}$). For the purposes of this study we fix the lower end of the TPC gap to be at $R_{\text{min}} = R_{\text{NS}}$ and the lower end of the OG gap to be at the null-charge surface.

These models, then, each take a total of two parameters for the purposes of this study: P (through R_{LC}) and α ; and produces a phaseplot similar the one shown in Figure 2.1(b) (for the TPC model) for each such combination. Just as for the Cone model, each constant- ζ cut through such a phaseplot produces a model LC like the one shown in Figure 2.1(d) (upon appropriately scaling its intensity and taking the relevant background level into account).

2.2 The Dual-band Models

When paired off, the single-band models constitute the OG+Cone and TPC+Cone conglomerated dual-band models. Since, in each case, both component single-band models are constrained to use the same magnetic field, α is constrained to always have the same value for both component models. Therefore, the OG+Cone and TPC+Cone models each take three parameters for the purposes of this study: P , $\dot{P}$, and α ; and produce a pair of phaseplots for each such combination. In analogue to the single-band models, a model LC pair is produced by taking a constant- ζ cut through both phaseplots (since a single observer is assumed for both wavebands).

The phaseplots and LCs shown in Figure 2.1 constitute an example of how the TPC+Cone model yields an LC pair; specifically for the P and $\dot{P}$ values associated with PSR J0742–2822 (as an example), and $(\alpha, \zeta) = (54^\circ, 45^\circ)$.




Figure 2.1: Example output for the Cone (panels (a) and (c)) and TPC (panels (b) and (d)) single-band models. The phaseplots in panels (a) and (b) plot the predicted relative intensity per solid angle for the radio and γ -ray models, respectively, for a magnetic inclination of $\alpha = 54^\circ$. The LCs in panels (c) and (d) are derived from these phaseplots by cutting through them at a constant value for the observer angle ζ (of 45° for this pair of LCs). Since the models used in this study are geometric in nature, and do not predict the absolute intensity per solid angle, these LCs are typically renormalised such that their peak intensity is 1. Since both α and ζ are the same for these two LCs, they constitute an LC pair associated with the TPC+Cone conglomerated model.




Figure 2.2: Illustration of the geometric TPC and OG models. The TPC radiation region is confined to the surface of the last closed field lines (between the thick dashed lines). It extends from the stellar surface all the way to the light cylinder. The shaded area shows the location of the OG radiation region, which extends from the null-charge surface to the light cylinder (adapted from [Dyks & Rudak, 2003](#)).

Chapter 3

Light Curve Fitting

Within the context of pulsar modelling, statistical LC fitting serves a number of purposes, the most notable of which are: the evaluation of the validity of the various pulsar models as descriptions of real-world pulsars, the extraction of pulsar parameter estimates from these models (e.g., [Johnson et al., 2014](#)) and the comparative evaluation of the various models' performance across the relevant pulsar populations (e.g., [Pierbattista et al., 2015](#)). Typically, statistical fitting works by leveraging the statistical properties of a goodness-of-fit-associated statistic. Therefore, regardless of the purpose with which statistical fitting is employed, its effectiveness is governed by how effectively that statistic characterises goodness of fit.

3.1 Pearson's χ^2 Statistic

Arguably the simplest, and most well-established such GOF statistic is Pearson's χ^2 statistic, upon which Pearson's χ^2 GOF test is built ([Pearson, 1900](#)). Within the LC fitting context, where the goal of a pulsar model is typically to reproduce the true LC D_{true} that underlies a given observed LC D (with n phase bins), the χ^2 statistic assigns a value to each model-predicted LC M , and can be defined as ([Bevington, 1969](#))

$$\chi^2(M; D, E) = \sum_{i=1}^n \left(\frac{D_i - M_i}{E_i} \right)^2, \quad (3.1)$$

where M_i , D_i and E_i are the model-predicted true LC intensity, the observed LC intensity and the measurement uncertainty (assumed to be Gaussian) in the i th phase bin; typically simply written " $\chi^2(M)$ " when there is no possibility of ambiguity. If M is identical to D_{true} , then the expected value of the absolute error $|D_i - M_i|$ in each bin is E_i , and the expected value of χ^2 is equal to n . The more M differs from D_{true} , however, the larger the expected value of $\chi^2(M)$ is.

Best Fit Identification When fitting a model that produces a predicted LC $M(\theta)$ for each parameter combination θ (comprised of n_θ parameter values) in its parameter space, the model's best-fit LC $M^{\text{pcs}} = M(\theta^{\text{pcs}})$, associated with the best-fit parameter combination θ^{pcs} , can be identified by minimising $\chi^2(M(\theta))$ over its parameter space; this yields $\chi_{\min}^2 = \chi^2(M^{\text{pcs}})$.¹ Since D is not identical to D_{true} (in general), M^{pcs} is not necessarily the model LC that is closest to D_{true} . However, if the model does, in fact, produce an LC $M_{\text{true}} = M(\theta_{\text{true}})$ that is identical to D_{true} (be it M^{pcs} or some other model LC; i.e., if the model is, in fact, valid as a description of the pulsar responsible for D), then the expected value of $\chi_{\min}^2$ is equal to the total number of degrees of freedom $\nu = n - n_\theta$ associated with the fit; subject to the conditions that the model is both continuous and linear in its parameters, and that the effective number of degrees of freedom associated with the minimisation procedure is, in fact, n_θ ([Andrae et al., 2010](#)). Notice the difference: the expected value of $\chi^2(M_{\text{true}})$ is n (as discussed above) and the expected value of $\chi_{\min}^2$ is ν .

¹Here, 'pcs' is a crude initialism for "Pearson's χ^2 " (invoked because using the " χ^2 " symbol itself would be impractical), and M^{pcs} and θ^{pcs} are labelled in the superscript to indicate that they are identified specifically via the marginalisation of a GOF statistic (the minimisation of the χ^2 statistic, in this case).

The χ^2 Goodness of Fit Test Another useful property of the χ^2 statistic (beyond the expected value of $\chi_{\min}^2$ for a valid model) is that for a valid model, $\chi_{\min}^2$ is χ^2 -distributed with ν degrees of freedom (hereafter written as “ $\chi_{\min}^2 \sim \chi_{k=\nu}^2$,” where k refers to the degrees of freedom of the χ^2 distribution²). This property means that it can serve as a test statistic with which to test the hypothesis that a given model is a valid model (i.e., that $M(\theta)$ produces an LC identical to D_{true}). This hypothesis can be rejected, at a confidence level s , if

$$\chi_{\min}^2 > \chi_{k=\nu, s}^2, \quad (3.2)$$

where $\chi_{k=\nu, s}^2$ is the corresponding confidence limit for the $\chi_{k=\nu}^2$ distribution (often read from a table). Otherwise, the hypothesis cannot be rejected, and the model is said to pass the χ^2 GOF test (with confidence s) since M^{PCS} is a good fit to D . Once a model passes the χ^2 GOF test it is reasonable to proceed under the assumption that the model produces an LC M_{true} that is identical to D_{true} .

3.1.1 Parameter Estimation

Estimates on the pulsar parameters are obtained by constraining the value of θ_{true} , and typically consist of two parts: a point estimate of θ_{true} and an uncertainty on that point estimate. The first of these two parts is already available since θ^{PCS} can serve as a point estimate of θ_{true} . The second of these parts can be constructed by administering a hypothesis test similar to that administered as part of the χ^2 GOF test. Specifically, for each LC the model produces the hypothesis that it is, in fact, M_{true} is tested.

An appropriate test statistic with which to administer this second hypothesis test (for a given model LC M) is the statistic

$$\begin{aligned} \Delta\chi^2(M) &= \chi^2(M) - \chi^2(M^{\text{PCS}}) \\ &= \chi^2(M) - \chi_{\min}^2, \end{aligned} \quad (3.3)$$

where the modified “ Δ ” symbol denotes the difference relative to $\chi_{\min}^2$. When evaluated for M_{true} , this difference is χ^2 -distributed with n_{θ} degrees of freedom (i.e., $\Delta\chi^2(M_{\text{true}}) \sim \chi_{k=n_{\theta}}^2$) since $\chi^2(M_{\text{true}}) \sim \chi_{k=n}^2$ and $\chi_{\min}^2 \sim \chi_{k=\nu}^2$ (Lampton et al., 1976). Therefore, the hypothesis that M is M_{true} can be rejected, at a given confidence level q , if

$$\Delta\chi^2(M) > \chi_{k=n_{\theta}, q}^2, \quad (3.4)$$

where $\chi_{k=n_{\theta}, q}^2$ is the corresponding confidence limit for the $\chi_{k=n_{\theta}}^2$ distribution.

To construct the desired constraints, this hypothesis test is leveraged to translate the confidence limit $\chi_{k=n_{\theta}, q}^2$ into a confidence region in the model’s parameter space (comprising all the parameter combinations that pass this hypothesis test). Together, θ^{PCS} and this confidence region constitute a full estimate of the pulsar parameters. However, since it can have a complicated shape, the latter is typically reduced to a set of uncertainties on the component values of θ^{PCS} by somehow characterising its shape. A simple way to do this (which yields conservative uncertainties) is to find the smallest n_{θ} -dimensional hyperrectangle that fully encloses the identified confidence region, and to use its extent for each of the model parameters as the uncertainty on that parameter.

Estimating only the Parameters of Interest If only a subset of the model’s parameters are of interest, the parameter estimate can be confined to only the $n_{\text{poi}} < n_{\theta}$ parameters of interest. To accomplish this, the model’s parameter space is first reduced to an n_{poi} -dimensional subspace of itself by minimising $\chi^2(M)$ over the $n_{\theta} - n_{\text{poi}}$ uninteresting parameters for each unique combination of the parameters of interest. Once this reduction is achieved, the second hypothesis test is conducted for each parameter combination in this subspace. Due to the minimisation over the uninteresting parameters, however, the $\chi_{k=n_{\theta}}^2$ distribution is not appropriate for this confined estimation; instead, the $\chi_{k=n_{\text{poi}}}^2$ distribution should be

²This choice of notation is made in the interest of clarity. Specifically, since the symbol “ χ^2 ” is used for both the distribution and the statistic, a subscript of the form “ $_{k=\nu}$ ” will be inserted throughout this thesis whenever the intent is to refer to a distribution.

used. Hence, the desired confidence region is constructed in this reduced, n_{poi} -dimensional parameter space by including only those parameter combinations (that form part of it) for which

$$\Delta\chi^2(M) \leq \chi_{k=n_{\text{poi}},q}^2, \quad (3.5)$$

where $\chi_{k=n_{\text{poi}},q}^2$ is the appropriate confidence limit for the $\chi_{k=n_{\text{poi}}}^2$ distribution (Avni, 1976). The reduction to a set of uncertainties on the individual parameters of interest is done by finding the smallest n_{poi} -dimensional hyperrectangle that encloses this region.

The Accuracy of the Estimate An important distinction to keep in mind when deriving parameter estimates using the χ^2 statistic, however, is that between the accuracy with which the model reproduces D_{true} (assessed by finding M^{pcs}) and the accuracy of the parameter estimate (i.e., how accurately θ_{true} is estimated). While the former is essentially what the χ^2 statistic measures (in the single-band case, at least), the latter cannot be determined based solely on LC fits. It might, for instance, be the case that the model in question accurately reproduces D_{true} (i.e., that it passes the χ^2 GOF test), but that θ^{pcs} differs substantially from θ_{true} ; or that the model is built on an incorrect or incomplete parameterisation altogether. While a full discussion of how a parameter estimate's accuracy might be assessed is outside the scope of this work, it is worth noting that a first step toward doing so is often to compare it to another (independent) estimate; e.g., the accuracy of an estimate derived via a radio-only LC fit might be assessed (to some extent) by comparing it to an estimate derived via a γ -ray-only or X-ray-only LC fit for the same pulsar. See Section 3.2.2 for a brief discussion of how radio dominance relates to parameter estimate accuracy within the context of dual-band LC fitting.

3.1.2 Example Synthetic Fit

Figure 3.1 shows an example single-band fit, where a synthetic γ -ray LC is extracted from the geometric two-pole caustic (TPC; see Section 2.1.2) model (by introducing Gaussian noise), and then fit using that same model. There are, in total, four free parameters involved in this fit: the pulsar's magnetic inclination angle α , the observer angle ζ , the model LC amplitude A and the position in phase of the pulsar's magnetic field axis ϕ_μ ; and the parameters of interest are α and ζ (the pulsar's so-called *viewing geometry*). Minimising the χ^2 statistic over the uninteresting parameters yields the χ^2 map plotted in panel (a).

Since $n_\theta = 4$ for this example, and the synthetic LC comprises $n = 100$ phase bins, the expected value of χ_{min}^2 is $\nu = 96$ for a valid model. The chosen confidence level for the first hypothesis test is $s = 2\sigma$, and at this confidence the model passes the test since

$$\chi_{\text{min}}^2 = 93.89 < \chi_{k=96,2\sigma}^2 \simeq 119.87.$$

The model's best-fit LC is plotted in panel (b), and the locations in (α, ζ) -space of θ_{true} and θ^{pcs} (for which $(\alpha, \zeta) = (65^\circ, 55^\circ)$ and $(67^\circ, 53^\circ)$, respectively) are marked in panel (a) with a dot and a hollow diamond (both yellow).

The white contour in panel (a) demarcates the confidence region in (α, ζ) -space that corresponds to the confidence level (of $q = 3\sigma$) that was chosen for the second hypothesis test. Since there are two parameters of interest (i.e., since $n_{\text{poi}} = 2$), this contour is constructed by including all the parameter combinations for which

$$\Delta\chi^2(M) \leq \chi_{k=2,3\sigma}^2 \simeq 11.83;$$

i.e., for which

$$\chi^2(M) \leq \chi_{\text{min}}^2 + \chi_{k=2,3\sigma}^2 = 105.72.$$

This confidence region is then reduced to uncertainties on each of the parameters of interest by constructing the yellow rectangle, which encloses it; its extent along each of the parameters' axes is used to determine the uncertainties. The full parameter estimate this fit yields (for the parameters of interest), is therefore

$$(\alpha, \zeta) = \left((67_{-3}^{+5})^\circ, (53_{-7}^{+3})^\circ \right).$$

Notice that the fit ostensibly recovers the (α, ζ) -space component of θ_{true} (which is typically unknown) since that




Figure 3.1: The results for a fit to a synthetic LC using χ_c^2 . In panel (a), χ^2 is plotted as a function of α and ζ . The dot and hollow diamond (both yellow) mark θ_{true} (the parameter combination associated with the synthetic LC) and θ^{PCS} . The irregular white contour encloses all the parameter combinations associated with model LCs that pass the second hypothesis test. The final parameter constraints are derived from the dimensions of the smallest rectangle that encloses this contour (the yellow rectangle in this example). In panel (b), the best-fit LC is plotted (in yellow) alongside the synthetic LC used (in grey with black uncertainties), both covering two rotations of the pulsar in the interest of clarity. The background emission level chosen for the synthetic LC is 50 weighted counts per bin, as indicated by the horizontal dashed line in panel (b).

component, in fact, lies inside the derived confidence region.

3.1.3 The Reduced χ^2 Statistic

The value that the χ^2 statistic assigns to a given LC fit only carries meaning when compared to the relevant value of ν . Relaying such statistic values can, therefore, become quite cumbersome. To overcome this, χ^2 values are often expressed as a fraction of ν , through the reduced χ^2 statistic

$$\chi_\nu^2(M) = \frac{1}{\nu} \chi^2(M). \quad (3.6)$$

Conveniently, this normalisation has no significant impact on the procedure described above, since

$$\chi_\nu^2(M_1) > \chi_\nu^2(M_2) \iff \chi^2(M_1) > \chi^2(M_2), \quad (3.7)$$

where M_1 and M_2 are two model LCs. This means that the same confidence regions can be constructed using the appropriate distributions for $\chi_{\nu, \min}^2$ and $\Delta\chi_\nu^2(M_{\text{true}})$.

3.1.4 Dual-band Extension

In the dual-band case, each pulsar is associated with a pair of phase-matched LCs: a radio LC D_r , with n_r phase bins and associated uncertainties E_r , and a γ -ray LC D_γ , with n_γ phase bins and associated uncertainties E_γ . Therefore, a dual-band pulsar model typically produces a pair of phase-matched LCs $M_r(\vartheta)$ and $M_\gamma(\vartheta)$ for each parameter combination ϑ (made up of n_ϑ parameters) in its parameter space, and is aimed at thereby reproducing the pair of true LCs $D_{r, \text{true}}$ and $D_{\gamma, \text{true}}$ (that underlie D_r and D_γ).

Since the two χ^2 statistics appropriate for the single-band-only fits,

$$\chi_r^2(M_r) = \chi^2(M_r; D_r, E_r) \quad \text{and} \quad (3.8)$$

$$\chi_\gamma^2(M_\gamma) = \chi^2(M_\gamma; D_\gamma, E_\gamma), \quad (3.9)$$

already assign values to M_r and M_γ individually, a natural extension of the χ^2 statistic for joint, dual-band fits would be a statistic that depends on $\chi_r^2(M_r)$ and $\chi_\gamma^2(M_\gamma)$. The simplest such statistic, which effectively treats each phase-matched pair of radio and γ -ray LCs as a single, concatenated LC (with $n_c = n_r + n_\gamma$ phase bins), is

$$\chi_c^2(M_c; D_c, E_c) = \chi_r^2(M_r) + \chi_\gamma^2(M_\gamma), \quad (3.10)$$

where $M_c(\vartheta) = (M_r, M_\gamma)$ is the model-predicted LC pair and $D_c = (D_r, D_\gamma)$ is the observed LC pair with associated errors $E_c = (E_r, E_\gamma)$; typically written “ $\chi_c^2(M_c)$ ” when there is no possibility of ambiguity. Within this framework, the overarching goal of a dual-band model is to reproduce the true LC pair $D_{c, \text{true}} = (D_{r, \text{true}}, D_{\gamma, \text{true}})$.

The primary advantage of defining χ_c^2 in this way is that it inherits essentially all of the useful properties of the χ^2 statistic (since the sum of independent χ^2 -distributed statistics is also χ^2 -distributed):

1. For a model with best-fit LC pair $M_c^{\text{PCS}} = M_c(\vartheta^{\text{PCS}})$, the expected value of $\chi_{c, \min}^2 = \chi_c^2(M_c^{\text{PCS}})$ is $\nu_c = n_c - n_\vartheta$ (the total number of degrees of freedom associated with the fit) if that model is valid (i.e., if it produces an LC pair $M_{c, \text{true}} = M_c(\vartheta_{\text{true}})$ that is identical to $D_{c, \text{true}}$; if the conditions listed in Section 3.1 are met).
2. For such a model (i.e., a valid dual-band model), $\chi_{c, \min}^2 \sim \chi_{k=\nu_c}^2$. Therefore, the hypothesis that a given dual-band model is valid can be tested via the χ^2 GOF test. Specifically, the hypothesis is rejected with confidence s if $\chi_{c, \min}^2 > \chi_{k=\nu_c, s}^2$ (the corresponding confidence limit of the $\chi_{k=\nu_c}^2$ distribution).
3. The uncertainties on the individual parameter values that make up ϑ_{true} are derived by leveraging the fact that

$$\Delta\chi_c^2(M_{c, \text{true}}) = \chi_c^2(M_{c, \text{true}}) - \chi_{c, \min}^2 \sim \chi_{k=n_\vartheta}^2; \quad (3.11)$$

specifically by characterising the shape of the appropriate confidence region (in the model’s parameter space; associated with a chosen confidence level q) that comprises all the parameter combinations for which $\Delta\chi_c^2(M_c(\vartheta)) \leq \chi_{k=n_\vartheta, q}^2$. This derivation procedure changes similarly to the one for single-band fits when aimed at estimating only n_{poi} parameters of interest.

The Accuracy of the Estimate (Dual-band)

Just as is the case for a single-band-only parameter estimate, the accuracy of the dual-band parameter estimate thus derived cannot be determined based solely on LC fits. The addition of a second waveband (as compared to a single-band-only fit) does, however, shed some light on the accuracy of the corresponding single-band-only estimates (derived by, e.g., finding the model LC pair for which χ_r^2 is a minimum), especially when the dual-band model is a conglomeration of two single-band models (as is the case for the dual-band models used in this work; see Section 2.2). Stated most succinctly: If the two single-band-only parameter estimates are not in agreement (i.e., if these estimates are not colocated in the dual-band model’s parameter space; see Section 3.2.2 for a demonstration of this), then it must be the case that at least one of the two single-band-only estimates is inaccurate. Unfortunately, LC fitting alone (be it single- or dual-band) cannot give greater clarity than this with regard to which single-band-only estimate is more accurate. Of course, if the two single-band-only estimates are in agreement (i.e., if they are colocated in the dual-band model’s parameter space), then they are equally accurate, and are just as accurate as the dual-band estimate (although this accuracy still cannot be determined based solely on LC fits).

The Dual-band Reduced χ^2 Statistic

Just as was the case for the χ^2 statistic, the χ_c^2 statistic can be written in its reduced form as

$$[\chi_c^2]_\nu(M_c) = \frac{1}{\nu_c} \chi_c^2(M_c), \quad (3.12)$$

which is equivalent to χ_c^2 with regard to the best fits it identifies. Notice that, with $\chi_{\nu, r}^2$ and $\chi_{\nu, \gamma}^2$ the reduced equivalents of χ_r^2 and χ_γ^2 ,

$$[\chi_c^2]_\nu(M_c) \neq \chi_{\nu, r}^2(M_r) + \chi_{\nu, \gamma}^2(M_\gamma). \quad (3.13)$$

In the interest of clarity (especially in later chapters, where this type of distinction will become important), this subtle difference between χ_c^2 and its reduced equivalent is encoded in the notation for the latter: The subscript that indicates summation (“ c ” in this case) precedes the subscript that indicates reduction (“ ν ” in this case). In essence, the order in which the subscripts are placed indicates the order of operations that results in the statistic in question; first summation, then division, in this case.

Example Synthetic Fit (Dual-band)

Figure 3.2 shows an example joint fit where synthetic, phase-matched radio and γ -ray LCs are fit using phase-matched LCs produced by a conglomerated dual-band model comprising the geometric TPC γ -ray model (as in the single-band example) and an empirical hollow-cone geometric radio model (see Section 2.1.1; Story et al., 2007). For this fit there are five free parameters: α , ζ , ϕ_μ , and the radio and γ -ray model LC amplitudes a_r and a_γ . The parameters of interest are again α and ζ , and minimising over the other parameters yields the χ_c^2 map plotted in panel (a).

Since $n_r = 250$ and $n_\gamma = 100$, the synthetic LC pair as a whole comprises a total of $n_c = 350$ phase bins, and the expected value for $\chi_{c, \min}^2$ for a valid model (i.e., a model that passes the dual-band χ^2 test) is $\nu_c = 350 - 5 = 345$. Since

$$\chi_{c, \min}^2 = 358.42 < \chi_{k=345, 2\sigma}^2 \simeq 389.31,$$

the model passes the first hypothesis test at the chosen confidence level of $s = 2\sigma$. The best-fit model LC pair is plotted in panels (b) and (c), and the locations in (α, ζ) -space of ϑ_{true} and ϑ^{pcs} (for which $(\alpha, \zeta) = (65^\circ, 55^\circ)$ and $(65^\circ, 56^\circ)$, respectively) are marked in panel (a) with a dot and a hollow diamond (both yellow). Note that the values for α and ζ for ϑ_{true} are the same as those for θ_{true} in the single-band example (plotted in Figure 3.1), and were used to generate both the radio and γ -ray synthetic observed LCs.




Figure 3.2: The results for a fit to a synthetic dual-band LC pair using χ_c^2 . In panel (a), χ_c^2 is plotted as a function of α and ζ . As per Eq. (3.10), this map is the sum of the radio and γ -ray single-band maps (i.e., the sum of the maps for χ_r^2 and χ_γ^2 ; the latter of which is plotted in Figure 3.1). Since emission is observed in both wavebands in this example (by construction), the best joint fit must lie in the coloured region of the map. The greyed region, where the Cone model predicts no emission, is included for reference (i.e., to aid comparison with Figure 3.1). In panels (b) and (c), the joint best-fit LC pair is plotted (consisting of a radio model LC and a γ -ray model LC; both in yellow) alongside the synthetic LC pair used (in grey with black uncertainties). The background emission level used for the synthetic radio LC (plotted in panel (c)) is 0. The effect of the additional constraint (with respect to the example single-band fit in Section 3.1.2) that the radio LC should also be reproduced is clearly visible in the morphology of the χ_c^2 map (panel (a)): In addition to the features on the map of χ_γ^2 that roughly lie along lines of constant $\zeta + \alpha$ (as is typical for γ -ray χ^2 maps), there are features on the map of χ_c^2 that roughly lie along lines of constant $\beta = \zeta - \alpha$ (as is typical for the Cone model). In a sense, the inclusion contour on this map is an intersection of the contours on the maps of χ_r^2 and χ_γ^2 for the single-band-only fits.




Figure 3.3: An example of a model LC that reproduces the overall structure of an observed LC, despite not being good enough for the model to pass the χ^2 test; i.e., of an LC that would result in a *merely favourable* fit (were it the model’s best-fit LC). While this model does not produce an LC identical to D_{true} , it is likely still a useful description of the pulsar responsible for D .

The confidence region in panel (a) (the white contour; again corresponding to a chosen confidence level of $q = 3\sigma$), is constructed similarly to the one for the single-band example, with all the parameter combinations that produce LC pairs for which

$$\chi_c^2(M_c) \leq \chi_{c,\min}^2 + \chi_{k=2,3\sigma}^2 = 370.25,$$

being included (recall that $\chi_{k=2,3\sigma}^2 \simeq 11.83$).

The estimate this fit yields for the parameters of interest is then

$$(\alpha, \zeta) = \left((65_{-1}^{+1})^\circ, (56_{-2}^{+1})^\circ \right).$$

Notice that this parameter estimate is somewhat more constrained than the estimate in the single-band example, owing to the joint nature of the fit. In essence the radio fit serves to constrain the value of $\beta = \zeta - \alpha$; i.e., it constrains the counter-diagonal extent of the contour.

3.2 Real-World Fits

Arguably the most attractive feature of the χ^2 statistic is that it can be used both to test GOF and to derive statistically robust parameter constraints. Within the pulsar astrophysics context this means that the same statistic can be used both to test how useful the available models are and to derive constraints on the parameters of the observed pulsars by using those models. The fitting procedure outlined in Section 3.1.1 demonstrates how the χ^2 statistic can be leveraged to accomplish these two goals, and the accompanying example fits can be considered representative of the ideal case: where a model passes the χ^2 test and yields constraints on the pulsar parameters.

There are, however, circumstances in which this fitting procedure proves ineffectual, specifically with regard to its ability to estimate and constrain the pulsar parameters.

3.2.1 Single-band Fits

When conducting single-band fits it often happens that a model reproduces the overall structure of the observed LC, but fails to reproduce its fine structure as fully as is required by the uncertainties on that LC. In such cases, the model typically fails the χ^2 GOF test because M^{PCS} is not a good fit to D , and can therefore not be considered valid as a description of the pulsar responsible for D . Figure 3.3 shows an example of such a case: the radio-band best fit to PSR J0631–1036 using the Cone model introduced in Section 2.1.1. As this example demonstrates, however, such models are not entirely useless since they still roughly reproduce the overall structure of the observed LC. Therefore, fits like the one shown in Figure 3.3 could still yield meaningful parameter constraints despite the model’s failure to pass the χ^2 GOF test. Within the context of pulsar parameter estimation, it is this criterion of model ‘usefulness’ that is most important, since the ultimate goal is




Figure 3.4: The effectiveness of the χ^2 statistic for real-world single-band fits. On this Euler plot, “Good” fits (which lie inside the solid boundary) are fits that pass the χ^2 GOF test while “Favourable” fits (which lie inside the dashed boundary) are fits where the model LC reproduces the observed LC to some extent, and that therefore imply that the model might be somewhat useful. While parameter estimates can be made for any of the favourable fits via the procedure described in Section 3.1.1, the χ^2 GOF test cannot be used to discriminate between favourable fits and truly bad fits (which lie outside the dashed boundary) since it rejects both merely favourable fits (which lie in the blue region) and truly bad fits. It only successfully identifies favourable fits that lie in the green region; i.e., fits that are both favourable and good.

to derive meaningful constraints on the pulsar parameters. Stated more succinctly, the χ^2 GOF test suffers from reduced statistical sensitivity when used to characterise model usefulness since it is blind to the quality of the broad-structure fit.

The Taxonomy of Single-band Fits

Figure 3.4 shows this perceived failure diagrammatically, where each region of the Euler plot represents a subset of all possible single-band fits. The region labelled “Good” (coloured green; with a solid circular boundary) includes all the fits where the model passes the χ^2 GOF test (i.e., all good fits), while the region labelled “Favourable” (with a dashed circular boundary; dashed to indicate that it is ill-defined) includes all the fits where the model reproduces the broad structure of the observed LC to some extent. Notice that the former is a subset of the latter; i.e., all ‘good’ fits are ‘favourable’ fits. The region that comprises fits that are favourable but not good (or, that are ‘merely favourable’) is coloured blue. The χ^2 GOF test, when interpreted as distinguishing between useful and useless models, rejects models that produce fits that lie in this region despite their ability to reproduce the overall structure of the observed LC being fit.

Deriving Constraints from Merely Favourable Fits

Since the procedure by which constraints on the pulsar parameters may be derived (outlined in Section 3.1.1) presupposes that the model passes the χ^2 GOF test (through its reliance on the distributions of $\Delta\chi^2(M_{\text{true}})$ and χ^2_{min}), it cannot, strictly speaking, be used to derive constraints on the pulsar parameters from fits that are merely favourable. Fortunately, as will be demonstrated in Section 7.4, $\Delta\chi^2(M)$ is still χ^2 -distributed with n_θ (or n_{poi}) degrees of freedom when evaluated for a model’s true best fit (i.e., when evaluated for the LC $M_{\text{prox}} = M(\theta_{\text{prox}})$ that fits D_{true} best). Therefore, pulsar parameter constraints can still be derived from merely favourable fits by leveraging the relevant confidence level of the $\chi^2_{k=n_\theta}$ (or $\chi^2_{k=n_{\text{poi}}}$) distribution, and in the same manner as before (i.e., by constructing the confidence region in the model’s parameter space that is associated with the relevant confidence level). Note that the resulting parameter estimate (which is constituted by the combination of θ^{PCS} with these constraints) is an estimate of θ_{prox} , not of θ_{true} (as is the case for good fits).




Figure 3.5: Results for four different dual-band fits to synthetic LCs using χ_c^2 . On each χ_c^2 map in panels (d) and (e), the green and blue dots indicate the locations of $\theta_{r,\text{true}}$ and $\theta_{\gamma,\text{true}}$, respectively; the parameter combinations used for the radio and γ -ray models in generating the synthetic observed LC pair plotted in panel (f) (in black). On the maps in panels (a) and (b), where the same parameter combination was used for both single-band models in generating the synthetic observed LC pair plotted in panel (c), only a single (white) dot is plotted. In each case, a corresponding hollow diamond is plotted in the same colour to indicate the location of θ^{PCS} (the corresponding single-band-only parameter point estimate). On each map the yellow (for balanced fits) or red (for unbalanced fits) circle indicates the location of ϑ^{PCS} (the dual-band parameter point estimate for the corresponding joint fit). The joint best-fit LC pairs are plotted against the relevant synthetic LC pair in the right-hand column (in the same colour as the circle used to mark it). The procedure described in Section 3.1.1 only fails to derive a sensible point estimate for the fit that results in the plot in panel (e), for which the single-band-only parameter estimates are non-collocated and the radio and γ -ray LC uncertainties are unbalanced (i.e., there is an uncertainty disparity; as indicated by the inscriptions regarding ε_r and ε_γ in the right-hand column).

3.2.2 Dual-band Fits

In the dual-band case, two additional, compounding effects become important: that of the non-colocation, in parameter space, of the respective single-band best fits, and that of the aforementioned uncertainty disparity between the radio and γ -ray wavebands. In contrast to the single-band case, however, these effects do not merely result in the miscategorisation of some favourable fits (specifically, in the miscategorisation of merely favourable fits as bad fits), but result in parameter estimates where one or both of the components of the estimate (comprising the point estimate and the uncertainty associated with it) are dominated by the GOF in a single waveband; i.e., these effects result in the manifestation of single-band dominance (typically radio dominance).

Point Estimates

Figure 3.5 illustrates how these two effects conspire to induce radio-dominated point estimates by comparing the results obtained when a dual-band model is fit to each of two different pairs of synthetic observed LCs, for two different sets of uniform radio and γ -ray LC uncertainties. The first pair of synthetic observed LCs, plotted in black in panel (c), was generated by adding random noise to the model LC pair associated with the model parameters indicated by the white dot in panels (a) and (b): $(\alpha, \zeta) = (45^\circ, 60^\circ)$. The second pair of synthetic observed LCs, plotted in panel (f), was generated similarly, but its component LCs are based on LCs from two different model LC pairs: the model LC pair associated with $(\alpha_r, \zeta_r) = (5^\circ, 20^\circ)$ for the radio LC and the model LC pair associated with $(\alpha_\gamma, \zeta_\gamma) = (70^\circ, 45^\circ)$ for the γ -ray LC (the green and blue dots, respectively, in panels (d) and (e)). This LC pair is generated in this way to ensure that the resulting single-band-only parameter estimates are non-colocated in the model's parameter space; i.e., to ensure that the resulting dual-band fits mimic real-world fits where one or both of the single-band models yield parameter estimates that are inaccurate to some extent (see Section 3.1.4). In essence, since the dual-band models used in real-world dual-band fits often do not produce an LC pair that is similar to the observed LC pair, the second synthetic observed LC pair is generated such that it is impossible for the dual-band model to reproduce it during the synthetic fit; thus the parameter estimates derived for this LC pair are a compromise, similar to their real-world counterparts.

As indicated by the white diamond in panels (a) and (b), the single-band-only fits to the first pair of synthetic observed LCs recover the parameters used to generate those LCs, and are therefore colocated. For the second pair of synthetic observed LCs, the single-band best fits are non-colocated, as indicated by the green and blue diamonds in panels (d) and (e)).

Each of these two pairs of synthetic observed LCs were fit twice: once with balanced radio and γ -ray uncertainties and once with unbalanced radio and γ -ray uncertainties. The orange circle in panel (a), located at $(45^\circ, 60^\circ)$, corresponds to the best fit to the first pair of synthetic observed LCs for the case where the radio and γ -ray uncertainties are balanced (as indicated by the inscriptions below the radio and γ -ray LCs in panel (c)). The model LCs associated with this fit are plotted in orange in panel (c). The red circle in panel (b), again located at $(45^\circ, 60^\circ)$, along with the red LCs in panel (c), corresponds to the best fit for disparate LC uncertainties. In both cases, since the single-band best fits are colocated, χ_c^2 successfully identifies the correct dual-band best-fit LC.

A similar, arguably correct, identification occurs for the case where the LC uncertainties are balanced but the single-band best fits are non-colocated. The dual-band fit, associated with the orange circle in panel (d) and plotted in orange in panel (f), can be regarded as successful in finding an adequate best fit in that the resulting pair of LCs roughly reproduces the overall shape of both the radio and γ -ray LCs that make up the first pair of synthetic observed LCs.

Where χ_c^2 -equipped parameter estimation procedure fails, however (with regard to the quality of the point estimate it yields), is in the case where both non-colocation and an uncertainty disparity are present. As the best-fit locations plotted in panel (e) and the pair of best-fit LCs plotted in red in panel (f) show, the radio GOF completely dominates the joint fit and the radio-only best fit is chosen as the best joint fit despite being associated with an essentially flat γ -ray LC.

Parameter Constraints

This single-band dominance also extends to the derived parameter constraints, as is evident in the morphologies of the χ_c^2 maps plotted in panels (b) and (e) in Figure 3.5, the χ_c^2 maps that are associated with the two example fits where the radio and γ -ray uncertainties are unbalanced. In contrast to the χ_c^2 maps plotted in panels (a) and (d) (associated with

the two balanced fits), where there is clear evidence that the joint GOF (i.e., χ_c^2) is sensitive to the γ -ray-only GOF, the morphologies of the χ_c^2 maps plotted in panels (b) and (e) are almost identical to those of the maps of χ_r^2 for the two radio-only fits (not shown). While this effect is somewhat hard to discern when considering only the regions where both single-band models predict emission (the non-grey coloured regions), it becomes less so when the morphology of these maps in the regions where the radio model predicts no radio emission (indicated using a greyscale colourmap) are also taken into account. In these regions (in panels (b) and (e)), the joint GOF is almost uniformly at its minimum, regardless of the γ -ray GOF.

Therefore, while the procedure described in Section 3.1.1 would yield parameter constraints for all four example fits in Figure 3.5, such constraints would only be sensitive to the GOF in both bands for the two fits where the radio and γ -ray uncertainties are balanced. For both of the fits where those uncertainties are unbalanced such constraints would be radio dominated.

Single-band Dominance and Parameter Estimate Accuracy

The plots in Figure 3.5 also serve to elucidate the impact that radio dominance (or, more generally, single-band dominance) might have on the accuracy of a dual-band parameter estimate. As discussed in Section 3.1.4, non-located single-band-only parameter estimates like those depicted in panels (d) and (e) signify that at most one of the two single-band-only parameter estimates can be accurate. Since it is not possible to know, based solely on LC fitting, how accurate (if at all) either of the two single-band-only parameter estimates are, the results plotted in panels (d) and (e) could correspond to one of (broadly) four possible scenarios:

1. **Only the radio-only estimate is accurate** (and the γ -ray-only estimate therefore inaccurate and non-located with the radio-only estimate). In this scenario the radio-dominated dual-band estimate (i.e., the dual-band estimate in panel (e)) would be accurate. The undominated estimate depicted in panel (d) would, of course, be inaccurate in this scenario. However, since it somewhat accurately recovers the true value of $\beta = \zeta - \alpha$ (i.e., the value of β associated with the accurate radio-only estimate; notice that all of the relatively good radio-only fits lie at roughly similar $|\beta|$ in panel (e)), it is somewhat more accurate than the inaccurate γ -ray-only estimate. In effect, the more accurate radio-only estimate pulls both dual-band estimates toward greater accuracy.
2. **Only the γ -ray-only estimate is accurate** (and the radio-only estimate therefore inaccurate and non-located with the γ -ray-only estimate). In this scenario the radio-dominated dual-band estimate would be just as inaccurate as the inaccurate radio estimate. In contrast to the previous scenario, however, the undominated dual-band estimate is somewhat more accurate than the radio-dominated estimate (and, by extension, the radio-only estimate), since it somewhat accurately recovers the value of $(\alpha + \zeta)$ associated with the (accurate) γ -ray-only estimate. In effect, the radio dominance prevents the more accurate γ -ray-only estimate from pulling the radio-dominated dual-band estimate toward greater accuracy.
3. **Both of the single-band-only estimates are somewhat inaccurate.** In this scenario the radio-dominated dual-band estimate would again be just as accurate as the radio-only estimate. However, if it is the case that the radio-only estimate is accurate enough that it roughly reproduces the true value of β , and the γ -ray-only estimate is accurate enough that it roughly reproduces the true value of $\alpha + \zeta$ (i.e., if it is the case for both single-band models that an increase in the quality of the broad-structure fit roughly correlates with an increase in the accuracy of the derived single-band-only estimate), then the undominated dual-band estimate is likely more accurate than either of the other three estimates.
4. **Neither the radio-only nor the γ -ray-only estimate is at all accurate.** In this scenario both of the dual-band estimates are almost certainly also not at all accurate.

In the first two of these scenarios the dual-band fit is effectively redundant, since an accurate parameter estimate can be obtained via the appropriate single-band-only fit. In the last of these scenarios the dual-band fit is ineffectual, since neither single-band-only estimate is sufficiently accurate to meaningfully mitigate the inaccuracy in the other. In the third of these scenarios, however, the accuracy of the dual-band estimate almost certainly improves upon the elimination of the radio dominance.




Figure 3.6: The effectiveness of the parameter estimation procedure described in Section 3.1.1 for dual-band fits. Just as is the case for single-band fits, dual-band fits can also be labelled “Good” and “Favourable.” Dual-band fits can, however, be further subdivided based on whether or not the single-band-only best fits they are associated with are colocated (which results in the dashed boundary labelled “Colocated”), and whether or not the radio and γ -ray LC uncertainties are balanced (which results in the boundary labelled “Balanced”). For fits that lie in the green and blue regions, both point estimates and parameter constraints can be derived. For fits that lie in the orange region, only point estimates can be derived. For fits that lie in the red region, neither sensible point estimates nor parameter constraints can be derived, and these fits are either radio- or γ -ray-dominated (i.e., single-band dominated).

Therefore, since the specific accuracy of the single-band-only estimates is typically unclear for real-world fits (and it is typically unclear which of these four scenarios one is dealing with; where the single-band-only estimates are non-colocated), an undominated dual-band estimate is generally preferable because it nominally tends to be a somewhat accurate parameter estimate more often than does a radio-dominated estimate.

The Taxonomy of Dual-band Fits

Figure 3.6 shows the perceived failure of the parameter estimation procedure for joint, dual-band fits diagrammatically. Alongside the regions present in Figure 3.4 (those labelled “Good” and “Favourable”), two new regions have been added: one labelled “Colocated,” which includes all the fits where the single-band-only best fits are colocated (e.g., the fits in panels (a) and (b) in Figure 3.5), and one labelled “Balanced,” which includes all the fits where the radio and γ -ray LC uncertainties have roughly the same relative size (e.g., the fits in panels (a) and (d) in Figure 3.5). Just as is the case for single-band fits, the procedure described in Section 3.1.1 effectively derives both parameter point estimates and parameter constraints for good joint fits (i.e., that lie in the green region in Figure 3.6). However, the procedure does not always identify the most favourable fit for fits that are not good. Specifically, it fails to do so for merely favourable fits where the single-band best fits are non-colocated and the radio and γ -ray uncertainties are unbalanced (e.g., the fit in panel (e) in Figure 3.5; the red region in Figure 3.6). In this region the parameter point estimate is single-band dominated (typically radio). Additionally, the procedure fails to yield useful parameter constraints for merely favourable fits where the radio and γ -ray uncertainties are unbalanced (i.e., for fits in the red and orange regions in Figure 3.6), yielding instead parameter constraints that are single-band dominated (typically radio). Therefore, for fits that lie in the orange region in Figure 3.6 the procedure effectively only yields parameter point estimates, while for fits that lie in the red region in Figure 3.6 the procedure effectively yields nothing.

3.2.3 Attempts at Eliminating Radio Dominance

Since the large disparity between the radio and γ -ray LC uncertainties is directly responsible for the radio dominance, a natural first step towards a solution is to minimise that disparity.

Johnson et al. (2014) attempted to do this by artificially inflating the radio LC uncertainties. They first defined an on-peak interval for the γ -ray LC, then used the average relative uncertainty $\varepsilon_{\gamma,ave}$ inside that interval to calculate a nominal absolute radio uncertainty

$$E_R = \frac{n_\gamma}{n_r} \cdot D_{r,max} \cdot \varepsilon_{\gamma,ave}, \quad (3.14)$$

where $D_{r,max}$ is the peak intensity of the observed radio LC.³ Notably, the ratio between n_γ and n_r is included as a factor in this prescription to ensure that the waveband with the greater number of bins, for which there is therefore more information, still enjoys some priority.⁴ This uncertainty was then applied uniformly across the observed radio LC.

This prescription proved effective, mitigating the effect of the large uncertainty disparity for the majority of the 40 millisecond pulsars they studied.⁴ For seven of those pulsars, however, it was necessary to adjust the calculated nominal uncertainty manually to more fully balance the radio and γ -ray contributions. In five of those seven cases the nominal uncertainty had to be reduced by a factor of between 2 and 4, while for the remaining two cases the nominal uncertainty had to be further inflated by a factor of 2.

Pierbattista et al. (2015), fitting radio and γ -ray LCs for 41 normal pulsars, instead opted to perform each joint fit for each of two possible uniform radio LC uncertainties. The first of these uncertainties was the uncertainty (Johnson et al., 2011, a predecessor to E_R)

$$E_{R,\gamma\text{-peak}} = D_{r,max} \cdot \varepsilon_{\gamma,ave}, \quad (3.15)$$

which in effect simply omits the n_γ/n_r factor present in E_R .⁵ Pierbattista et al. derived the second uncertainty from the results of the radio-only log-likelihood fit conducted using the first by requiring that the reduced χ^2 value for the radio-only best fit (i.e., $\chi^2_{\nu,r,min}$) should be 1 when it is assumed. This requirement yields the uncertainty

$$E_{R,*} = \sqrt{-\frac{2 \ln L_{max}}{\nu}} \cdot E_{R,\gamma\text{-peak}}, \quad (3.16)$$

where $\ln L_{max} \leq 0$ is the log-likelihood associated with the best fit identified using the first uncertainty (i.e., using $E_{R,\gamma\text{-peak}}$).⁶ When conducting the joint fits using these uncertainties, they only preferred the $E_{R,*}$ -associated best-fit LC pair over the $E_{R,\gamma\text{-peak}}$ -associated best-fit LC pair if the likelihood associated with it was higher.

Assessing how effectively this prescription eliminated radio dominance is difficult, and the authors themselves did not comment on this. Of note, however, is that for at least one of the pulsars they fit (PSR J1124–5916), there is evidence of significant unmitigated γ -ray dominance brought about by comparatively much smaller γ -ray LC uncertainties.

³In the interest of clarity we have opted for notation that is more consistent with the rest of this document. Johnson et al. used “ σ_R ” to denote this uncertainty.

⁴Personal communication.

⁵Pierbattista et al. used “ σ_{peak} ” to denote this uncertainty.

⁶Pierbattista et al. used “ σ_* ” to denote this uncertainty.

Part II

DEVELOPMENT

Chapter 4

Finding a More Appropriate Statistic

4.1 The Cause of the Failure

As alluded to in Section 3.2.1, a core aim when leveraging the χ^2 statistic is to determine how useful the model in question is as a description of the pulsar responsible for the LC (or LC pair) being fit. Such procedures (like the one outlined in Section 3.1.1) typically accomplish this in three steps. First, each model LC is assigned a GOF value that characterises how well the model LC fits the observed LC. Then the model's best fit (i.e., the fit with the greatest GOF) is identified. Finally, the usefulness of the model as a whole is characterised by how good its best fit is. Crucially, this procedure is only effective if the LC identified as the model's best fit is also the LC that renders the model most useful. Conversely, if an LC other than the model's best-fit LC can be identified that would render the model more useful, then the procedure failed. In other words, the procedure is only effective if the quantity used to measure GOF is a maximum for the fit that renders the model most useful.

Mathematically, this means that the procedure is only effective if the model's usefulness is strictly increasing with respect to the quantity τ used to measure the GOF of the LCs it produces. In terms of a comparison between two candidate best-fit model LCs M_1 and M_2 , this monotonic relationship can be expressed as

$$\pi(M_1) > \pi(M_2) \iff \tau(M_1) > \tau(M_2), \quad (4.1)$$

where π is some quantity that characterises the model usefulness each LC would imply if it were the model's best fit.

The χ^2 statistic does not, however, measure GOF directly, but indirectly. Specifically, the χ^2 statistic measures deviation from observation (DFO), and can act as a proxy for a direct measurement of GOF under the assumption that GOF increases as DFO decreases. This is why *lower* χ^2 values are associated with *better* fits in Pearson's χ^2 test. Taking the indirect nature of the χ^2 statistic into account, the procedure is only effective if the usefulness of the model strictly increases as the value of the χ^2 statistic decreases; i.e., in the context of Eq. (4.1), if

$$\pi(M_1) > \pi(M_2) \iff \chi^2(M_1) < \chi^2(M_2). \quad (4.2)$$

As described in Section 3.2.2, real-world dual-band fits conducted using χ_c^2 are often radio dominated, and it is often the case that an LC pair other than M_c^{PCS} can be identified that more fully reproduces the observed LC pair, and hence corresponds to greater implied model usefulness. Therefore, χ_c^2 arguably only satisfies this condition for a comparatively small subset of real-world dual-band fits.

4.1.1 Scrutinising the χ_c^2 Statistic using a By-eye Fit

Stated plainly, the more fully a model can reproduce what is observed, the more useful the model is. A model that fully reproduces what is observed, is maximally useful. Figure 4.1 shows an example dual-band fit, for PSR J1048–5832, using an OG+Cone conglomerated dual-band model. The observed LCs, as reported in the 2PC, are plotted in grey, and the LC




Figure 4.1: The joint best-fit LC pair M_c^{pcs} (plotted in yellow; identified using χ_c^2) and an LC pair (M_c^{eye} , the best fit identified by eye; plotted in brown) that is arguably a better joint fit to the observed LC pair, for PSR J1048–5832 (using an OG+Cone conglomerated dual-band model).

Table 4.1: The statistic values associated with the candidate joint best-fit LC pairs plotted in Figure 4.1.

LC pair	χ_r^2	$\chi_{\nu,r}^2$	χ_γ^2	$\chi_{\nu,\gamma}^2$	χ_c^2	$[\chi_c^2]_\nu$
M_c^{pcs}	175508.0	697	17959.4	187	193467.4	551
M_c^{eye}	295656.9	1173	523.9	5	296180.8	844

uncertainties are shown as black vertical lines. The radio LC uncertainties (taken to be uniformly equal to the standard deviation in the off-pulse region) are very small compared to the γ -ray uncertainties (so small that the associated error bars in Figure 4.1 almost disappear), and the single-band-only best-fit parameter combinations are non-colocated. As a consequence, the fit in Figure 4.1 is radio-dominated, similarly to the synthetic fit shown in Figure 3.5(e) (with best-fit LC pair plotted in red in panel (f)). The χ_c^2 statistic identifies the LC pair M_c^{pcs} (plotted in yellow in Figure 4.1) as the model's best joint fit to the observed LC pair. Alongside M_c^{pcs} , another model-predicted LC pair M_c^{eye} is also plotted (in brown), which is the best-fit LC pair identified by eye.

The Single-band Fits

The radio components of both M_c^{pcs} and M_c^{eye} roughly reproduce the observed radio LC. While neither passes the χ^2 GOF test, both can be considered favourable fits (similarly to the fit shown in Figure 3.3), reproducing both the observed peak multiplicity and pulse width. As the values for χ_r^2 and $\chi_{\nu,r}^2$ in Table 4.1 reveal, the radio component of M_c^{pcs} outperforms that of M_c^{eye} , almost by a factor of 2.

For the γ -ray fits, the situation is dramatically different. The γ -ray component of M_c^{pcs} fails to reproduce the observed γ -ray LC entirely, while that of M_c^{eye} reproduces it reasonably well. Just as for the radio components, neither component LC passes the χ^2 GOF test. Unlike the radio components, however, only one of the two γ -ray components is a favourable fit to the observed γ -ray LC (the γ -ray component of M_c^{eye}); a flat LC would fare better than the γ -ray component of M_c^{pcs} , since it would deviate less from the observed LC. This stark apparent contrast in γ -ray GOF is reflected in the values for χ_γ^2 and $\chi_{\nu,\gamma}^2$ in Table 4.1; the γ -ray component of M_c^{eye} outperforms that of M_c^{pcs} almost by a factor of 40.

The Joint Fits

Evaluating the joint fits similarly reveals that, as suspected, a decrease in χ_c^2 does not necessarily translate to an increase in implied model usefulness (i.e., the degree to which the observed LC pair is reproduced does not necessarily increase as

the value of χ_c^2 decreases). With the two component LCs taken together, M_c^{pcs} only succeeds in roughly reproducing the observed radio LC (failing entirely to reproduce the observed γ -ray LC. In essence, the χ_c^2 fit seems to opt for an LC pair that reproduces the observed radio LC marginally better (as compared to M_c^{eye}), even though that marginal improvement carries with it the cost of completely failing to reproduce the γ -ray LC. In contrast, M_c^{eye} succeeds in roughly reproducing both of the observed LCs.

In aggregate, then, M_c^{eye} seems to imply that the model in question is capable of reproducing the observed LC pair reasonably well, while M_c^{pcs} seems to imply that it is only capable of reproducing the observed radio LC (as is typical for radio-dominated fits). Since a dual-band model that roughly reproduces both of the observed LCs is more useful (as a dual-band model) than one that only roughly reproduces one of the two observed LCs, M_c^{eye} implies greater model usefulness than does M_c^{pcs} . Therefore, a statistic that satisfies the condition stipulated in Eq. (4.2) would arguably rather identify M_c^{eye} as the model's best-fit LC pair, not M_c^{pcs} (i.e., would assign greater joint GOF to M_c^{eye} than to M_c^{pcs}). Since χ_c^2 does not prefer M_c^{eye} over M_c^{pcs} , it apparently does not satisfy the condition in Eq. (4.2), and is arguably not an appropriate statistic with which to gauge the implied model usefulness a given joint fit corresponds to.

4.2 The Scaled Flux

Comparing the values of χ_r^2 and χ_γ^2 for M_c^{pcs} and M_c^{eye} (in Table 4.1) reveals that equal decreases in the DFO in the radio and γ -ray bands do not necessarily correspond to equal increases in the joint GOF (as demonstrated in Figure 4.1). The reason for this is that the single-band χ^2 statistic effectively measures DFO in units of the LC uncertainties (since the deviation in each bin is divided by the corresponding LC uncertainty). Consequently, the definition of χ_c^2 (in Eq. (3.10)) essentially calls for two quantities that express GOF in different units to be added together without taking their units into account.

To define a more appropriate statistic, a more detailed understanding of the interplay between DFO, GOF, and implied model usefulness, is required. To gain such an understanding, it is first necessary to convert the indirect single- and dual-band χ^2 GOF statistics to direct GOF statistics. While such a conversion is problematic in general (i.e., beyond the LC fitting context), since for general curves it is impossible to coherently identify a class of model curves that have zero goodness, it is less so within the LC fitting context.

The key insight that allows this conversion to be done, is that the observed LCs are the result of a perturbation to the observed flux, effected by the pulsar. Had the pulsar not been present, the observed LC in each of the two bands would simply have been the LC constituted by the background emission in the that band. Each of the single-band observed LCs, D_r and D_γ , can therefore be viewed as a deviation from the background LC associated with it (B_r and B_γ , respectively), and the dual-band observed LC pair D_c can be viewed as a deviation from the background LC pair $B_c = (B_r, B_\gamma)$.

Since deviation is inherently symmetrical (i.e., since D_r deviates as much from B_r as B_r deviates from D_r , for instance), these deviations can be characterised by calculating

$$\Phi^2 = \chi^2(B), \quad (4.3)$$

where B is the relevant background LC; this yields Φ_r^2 and Φ_γ^2 for D_r and D_γ , and

$$\begin{aligned} \Phi_c^2 &= \chi_c^2(B_c) \\ &= \Phi_r^2 + \Phi_\gamma^2 \end{aligned} \quad (4.4)$$

for D_c . Understood physically, Φ_r and Φ_γ are the total scaled pulsar-associated fluxes contained in D_r and D_γ , and Φ_c is the total scaled pulsar-associated flux contained in D_c .

Relation to Measurement Precision An immediate consequence of the identification of Φ is that the precision with which the observed LCs capture the radio and γ -ray components of the perturbation, can be quantified. Upon recalling that the expected value of χ^2 for M_{true} (and by extension for D_{true}) is n (the number of phase bins), the estimated measurement

precision can be written as

$$\rho = \frac{\Phi^2}{n}, \quad (4.5)$$

which yields ρ_r and ρ_γ for D_r and D_γ , and

$$\rho_c = \frac{\Phi_c^2}{n_c} = \frac{n_r}{n_c} \rho_r + \frac{n_\gamma}{n_c} \rho_\gamma \quad (4.6)$$

for D_c .

4.3 Scaled-flux Normalised GOF

The identification of Φ (i.e., of Φ_r , Φ_γ , and Φ_c) also allows for the unambiguous identification of an interval in each of the relevant χ^2 -spaces (i.e., those associated with χ_r^2 , χ_γ^2 , and χ_c^2) that corresponds to a change in implied model usefulness from effectively useless to maximally useful: the intervals for which $\chi^2 \in [0, \Phi^2]$. In accordance with the indirect nature of the χ^2 statistic (as discussed in Section 4.1), the sense of these intervals is reversed, with the lower χ^2 value corresponding to the higher implied model usefulness in each case (and vice versa).

Stated in terms of GOF, this quantity allows for the identification of a class of model LCs that can sensibly be said to have zero GOF: all those model LCs for which $\chi^2(M) = \Phi^2$ (the set of all background-equivalent LCs); and hence allows for the definition of a direct GOF statistic,

$$\Psi^2(M) = \Phi^2 - \chi^2(M), \quad (4.7)$$

which is 0 for a fit to a background-equivalent LC (a “null” fit) and 1 for a fit to a model LC that perfectly reproduces D (a “perfect” fit). This statistic yields Ψ_r^2 and Ψ_γ^2 for the radio and γ -ray single-band fits, and

$$\begin{aligned} \Psi_c^2(M_c) &= \Phi_c^2 - \chi_c^2(M_c) \\ &= \Psi_r^2(M_r) + \Psi_\gamma^2(M_\gamma) \end{aligned} \quad (4.8)$$

for the joint fit, and can be viewed as the GOF complement to the χ^2 statistic (which measures DFO), since $\chi^2(M) + \Psi^2(M) = \Phi^2$.

The identification of Φ can be leveraged even further, though, to define the scaled-flux normalised GOF statistic

$$\begin{aligned} \Psi_\Phi^2(M) &= \frac{1}{\Phi^2} \Psi^2(M), \\ &= 1 - \frac{1}{\Phi^2} \chi^2(M), \end{aligned} \quad (4.9)$$

which not only measures GOF directly, but also expresses it in units that more naturally relate to model usefulness (i.e., in units of the scaled flux).¹ This statistic yields $\Psi_{\Phi,r}^2(M_r)$ and $\Psi_{\Phi,\gamma}^2(M_\gamma)$ for the radio and γ -ray single-band fits, and

$$\begin{aligned} [\Psi_c^2]_\Phi(M_c) &= 1 - \frac{1}{\Phi_c^2} \chi_c^2(M_c) \\ &= \frac{\Phi_r^2}{\Phi_c^2} \Psi_{\Phi,r}^2(M_r) + \frac{\Phi_\gamma^2}{\Phi_c^2} \Psi_{\Phi,\gamma}^2(M_\gamma) \end{aligned} \quad (4.10)$$

for the joint fit; the scaled-flux normalised, direct equivalent of the χ_c^2 statistic.² Each of these three statistics ($\Psi_{\Phi,r}^2$, $\Psi_{\Phi,\gamma}^2$, and $[\Psi_c^2]_\Phi$) yields a value of 0 for a null fit and a value of 1 for a perfect fit, and the interval in the respective spaces that corresponds to a change in implied model usefulness from effectively useless to maximally useful is the interval $[0, 1]$.

¹Notice that $\Psi_\Phi^2(M)$ effectively takes B as a parameter alongside D and E (which are the parameters that $\chi^2(M)$ takes; see Eq. (4.9)) since it depends on Φ (which appears both in the denominator in Eq. (4.9) and the definition of $\Psi^2(M)$)

²Here the order of operations is again encoded in the notation used since $[\Psi_c^2]_\Phi(M_c) \neq \Psi_{\Phi,r}^2(M_r) + \Psi_{\Phi,\gamma}^2(M_\gamma)$ (similarly to $[\chi_c^2]_\nu$, see Section 3.1.4)

Since the definitions of Ψ^2 and Ψ_Φ^2 reverse the usefulness correspondence associated with the relevant intervals ($[0, \Phi^2]$ and $[0, 1]$, respectively), best-fit identification using either statistic is done by maximising its value, instead of minimising it. Similarly to χ_ν^2 (see Section 3.1.3), since both Ψ^2 and Ψ_Φ^2 are linear in χ^2 (in both the single- and dual-band case), they will always yield the same best fits as the χ^2 statistic (and, by extension, the same parameter estimates). Mathematically (for Ψ_Φ^2), this equivalence can be expressed as

$$\Psi_\Phi^2(M)|_{\max} \equiv \chi^2(M)|_{\min}, \quad (4.11)$$

where the $\cdot|_{\max}$ and $\cdot|_{\min}$ operators indicate that Ψ_Φ^2 is equivalent *under maximisation* to χ^2 *under minimisation* (over the set of model LCs to which M belongs).³

An interesting property of Ψ_Φ^2 (and Ψ^2) is that

$$\Psi_\Phi^2(M) > 0 \iff \chi^2(M) < \Phi^2, \quad (4.12)$$

i.e., $\Psi_\Phi^2(M) > 0$ only when M is a better fit to D than B . A powerful corollary to this is that an LC for which $\Psi_\Phi^2(M) < 0$, corresponds to a model that reproduces the observed perturbation worse than does a model that posits the absence of a pulsar. Given this property, the Ψ_Φ^2 statistic naturally divides the set of possible model LCs into three distinct categories: *favourable fits* (or “better-than-background” fits; for which $\Psi_\Phi^2 > 0$), *null fits* (for which $\Psi_\Phi^2 = 0$), and *bad fits* (for which $\Psi_\Phi^2 < 0$).

Scaled-flux Normalised DFO The considerations that give rise to Ψ_Φ^2 also give rise to the scaled-flux normalised DFO statistic

$$\begin{aligned} \chi_\Phi^2(M) &= \frac{1}{\Phi^2} \chi^2(M) \\ &= 1 - \Psi_\Phi^2(M), \end{aligned} \quad (4.13)$$

which also effectively normalises the interval $[0, \Phi^2]$ in the relevant spaces, but does not reverse the sense of the associated usefulness correspondence. In essence, χ_Φ^2 expresses deviation in GOF units, due to the effective normalisation of the scaled flux. This statistic yields $\chi_{\Phi,r}^2$ and $\chi_{\Phi,\gamma}^2$ for the single-band fits, and

$$\begin{aligned} [\chi_c^2]_\Phi(M_c) &= \frac{1}{\Phi_c^2} \chi_c^2(M_c) \\ &= \frac{\Phi_r^2}{\Phi_c^2} \chi_{\Phi,r}^2(M_r) + \frac{\Phi_\gamma^2}{\Phi_c^2} \chi_{\Phi,\gamma}^2(M_\gamma) \\ &= 1 - [\Psi_c^2]_\Phi(M_c) \end{aligned} \quad (4.14)$$

for the joint fit; the scaled-flux normalised, indirect equivalent of χ_c^2 . This statistic yields a value of 1 for a null fit and a value of 0 for a perfect fit, and is also equivalent to the χ^2 statistic (both its single-band and dual-band incarnations) within the LC fitting context (upon minimisation).

4.4 Revisiting the χ_c^2 Statistic

The form of Eq. (4.10) reveals that the use of χ_c^2 , as a GOF statistic, carries with it an implicitly assumed weighting for the single-band GOF contributions. Specifically, it reveals 1) that (in general) equal changes in radio and γ -ray GOF do not translate to the same change in joint GOF, and 2) that the weights

$$w_r^{\text{PCS}} = \frac{\Phi_r^2}{\Phi_c^2} \quad \text{and} \quad w_\gamma^{\text{PCS}} = \frac{\Phi_\gamma^2}{\Phi_c^2}$$

³Note that this is an *equivalence*, not an *equality*.

carried by the radio and γ -ray GOF contributions, depend on the relative precision of the radio and γ -ray observations. Defining the scaled flux ratio $\lambda_{r\gamma} = \Phi_r/\Phi_\gamma > 0$, the implicitly assumed relative radio weight is

$$\omega_r^{\text{PCS}} = \frac{w_r^{\text{PCS}}}{w_\gamma^{\text{PCS}}} = \lambda_{r\gamma}^2 = \frac{n_r \rho_r}{n_\gamma \rho_\gamma}. \quad (4.15)$$

4.4.1 Single-band Dominance

Together, these two attributes of the χ_c^2 statistic cause single-band dominance in joint, dual-band fits where there is a large uncertainty disparity between the two bands. Radio dominance, specifically, occurs for pulsars where the radio observation is much more precise than the γ -ray observation; i.e., for which $\Phi_r^2 \gg \Phi_\gamma^2$, and hence $\rho_r \gg \rho_\gamma$ (assuming $n_r \sim n_\gamma$). For such pulsars,

$$w_r^{\text{PCS}} \simeq 1 \quad \text{and} \quad w_\gamma^{\text{PCS}} \simeq 0 \quad (4.16)$$

since $\Phi_c^2 = \Phi_r^2 + \Phi_\gamma^2 \simeq \Phi_r^2$, which means that $\omega_r^{\text{PCS}} \gg 1$ and

$$[\Psi_c^2]_\Phi(M_c) \simeq \Psi_{\Phi,r}^2(M_r), \quad (4.17)$$

and χ_c^2 effectively disregards the γ -ray-band GOF. Similarly, χ_c^2 effectively disregards the radio-band GOF for pulsars where the inverse is true, since for those pulsars $\Phi_\gamma^2 \gg \Phi_r^2$.

4.4.2 Implied Model Usefulness

Beyond the effect of its precision dependence, the weighting in Eq. (4.10) also determines how implied model usefulness depends on the single-band GOF contributions (as per the Eq. (4.1)).

Consider, for instance, a comparison between two candidate, dual-band models, $F_c(\vartheta)$ and $G_c(\vartheta)$, that yield the best-fit LC pairs $F_c^{\text{PCS}} = (F_r, F_\gamma)$ and $G_c^{\text{PCS}} = (G_r, G_\gamma)$, respectively, when fit to an observed LC pair $D_c = (D_r, D_\gamma)$, for which $\Phi_r > \Phi_\gamma \Rightarrow w_r^{\text{PCS}} > w_\gamma^{\text{PCS}}$. Suppose that F_c^{PCS} fits D_r perfectly but fits D_γ no better than B_γ , i.e., that

$$\Psi_{\Phi,r}^2(F_r) = 1 \quad \text{and} \quad \Psi_{\Phi,\gamma}^2(F_\gamma) = 0;$$

and that the reverse is true for G_c^{PCS} , i.e., that

$$\Psi_{\Phi,r}^2(G_r) = 0 \quad \text{and} \quad \Psi_{\Phi,\gamma}^2(G_\gamma) = 1.$$

Then, under Eq. (4.10),

$$[\Psi_c^2]_\Phi(F_c^{\text{PCS}}) = w_r^{\text{PCS}} \quad \text{and} \quad (4.18)$$

$$[\Psi_c^2]_\Phi(G_c^{\text{PCS}}) = w_\gamma^{\text{PCS}}, \quad (4.19)$$

which means that F_c^{PCS} is a better joint fit to D_c than G_c^{PCS} (according to χ_c^2 , at least), by a factor of $w_r^{\text{PCS}}/w_\gamma^{\text{PCS}} = \lambda_{r\gamma}^2$; i.e., with respect to D_c , the LC pair F_c^{PCS} is associated with greater joint GOF than the LC pair G_c^{PCS} (greater by a factor of $\lambda_{r\gamma}^2$).

Extending this statement about GOF to a statement about model usefulness, however, presents a problem: Both models perfectly reproduce one of the two observed LCs, yet F_c is deemed more useful because the radio observation is more precise. This problem is even more pronounced when $\lambda_{r\gamma}$ is large, with F_c (in this case) essentially deemed infinitely more useful than G_c when $\Phi_r \gg \Phi_\gamma$. Put another way, since F_c and G_c are effectively single-band models, the use of χ_c^2 essentially carries with it the assertion that (in this case) a perfect radio model is more useful than a perfect γ -ray model because the radio observation is more precise. While it may be the case that one of these models is more useful than the other, discriminating between them based on the relative precision of the observations involved, does not hold water.

4.5 The Scaled-flux Standardised $\Psi_{\Phi,c}^2$ Statistic

The key to finding a sensible alternative to the χ_c^2 statistic is to realise that its apparent shortcomings are not inherent to it, but a result of its effective misappropriation as a GOF statistic (where parameter estimation is concerned). The discussion in Section 4.3 reveals that χ_r^2 and χ_γ^2 carry units appropriate for measuring deviation (i.e., DFO), not GOF (since Φ_r and Φ_γ typically differ), which essentially renders the addition in Eq. (3.10) only proper when χ_c^2 is interpreted as a measurement of deviation. When interpreted as a measurement of GOF, as in Eq. (4.10) (upon defining $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$, which carry units more appropriate for measuring GOF), χ_c^2 is revealed to disproportionately favour the GOF contribution of the waveband associated with the more precise observation.

A more appropriate extension of the single-band χ^2 statistic, with regard to GOF, would be a statistic that grants the single-band GOF contributions equal weight, just as χ_c^2 grants the single-band DFO contributions equal weight. Adding to this the requirement that the resulting statistic should, itself, be scaled-flux normalised, yields the scaled-flux standardised statistic

$$\begin{aligned}\Psi_{\Phi,c}^2(M_c) &= \frac{1}{2} \Psi_{\Phi,r}^2(M_r) + \frac{1}{2} \Psi_{\Phi,\gamma}^2(M_\gamma) \\ &= 1 - \frac{1}{2} \left[\frac{1}{\Phi_r^2} \chi_r^2(M_r) + \frac{1}{\Phi_\gamma^2} \chi_\gamma^2(M_\gamma) \right].\end{aligned}\tag{4.20}$$

Just as is the case for $[\Psi_c^2]_\Phi$, since both are scaled-flux normalised, $\Psi_{\Phi,c}^2 = 0$ for a null fit and $\Psi_{\Phi,c}^2 = 1$ for a perfect fit. Unlike $[\Psi_c^2]_\Phi$, however, $\Psi_{\Phi,c}^2 = 0.5$ for both LC pairs considered in Section 4.4.2.

Scaled-flux Standardised DFO This statistic, of course, has a DFO counterpart: the scaled-flux standardised DFO statistic

$$\begin{aligned}\chi_{\Phi,c}^2(M_c) &= \frac{1}{2} \chi_{\Phi,r}^2(M_r) + \frac{1}{2} \chi_{\Phi,\gamma}^2(M_\gamma) \\ &= \frac{1}{2} \left[\frac{1}{\Phi_r^2} \chi_r^2(M_r) + \frac{1}{\Phi_\gamma^2} \chi_\gamma^2(M_\gamma) \right] \\ &= 1 - \Psi_{\Phi,c}^2(M_c).\end{aligned}\tag{4.21}$$

Interestingly, just as the identification of a GOF counterpart to χ_c^2 (in Eq. (4.7)) revealed that χ_c^2 is a weighted sum of the single-band GOF statistics ($\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$; through Eq. (4.10)), this identification of a DFO counterpart to $\Psi_{\Phi,c}^2$ reveals that it, in turn, is a weighted sum of the single-band DFO statistics (χ_r^2 and χ_γ^2). Specifically, upon defining the equivalent statistic (equivalent to $\chi_{\Phi,c}^2$)

$$\begin{aligned}[\chi_{\Phi,c}^2]_{E_c} &= \frac{2\Phi_r^2\Phi_\gamma^2}{\Phi_c^2} \chi_{\Phi,c}^2 \\ &= \frac{\Phi_\gamma^2}{\Phi_c^2} \chi_r^2 + \frac{\Phi_r^2}{\Phi_c^2} \chi_\gamma^2,\end{aligned}\tag{4.22}$$

which normalises the weights of χ_r^2 and χ_γ^2 in Eq. (4.21), it becomes clear that $\Psi_{\Phi,c}^2$ effectively weights the single-band deviations inversely to how χ_c^2 effectively weights the single-band GOF contributions (compare Eq. (4.10)).

4.5.1 Scrutinising $\Psi_{\Phi,c}^2$ Using a By-eye Fit

An immediate first check of whether this new test statistic constitutes a more effective measurement of joint GOF is to compare its performance to that of χ_c^2 for the radio-dominated fit shown in Figure 4.1. The values for the relevant single-band and joint GOF statistics, some of which take the scaled fluxes $\Phi_r \simeq 1914$ and $\Phi_\gamma \simeq 61$ into account (these scaled fluxes yield $\lambda_{r\gamma} \simeq 31.38$), for both best-fit LC pairs M_c^{pcs} and M_c^{eye} , are listed in Table 4.2.

In accordance with Table 4.1, the radio-band GOF (encoded in $\Psi_{\Phi,r}^2$) for M_c^{pcs} is greater than that for M_c^{eye} , and vice versa for the γ -ray-band GOF (encoded in $\Psi_{\Phi,\gamma}^2$). Additionally, the value of $\Psi_{\Phi,\gamma}^2$ for M_c^{pcs} is negative, which confirms

Table 4.2: Single- and dual-band GOF values for the radio-biased fit shown in Figure 4.1. For this observed LC pair $\Phi_r \simeq 1914$ and $\Phi_\gamma \simeq 61$.

LC pair	$\Psi_{\Phi,r}^2$	$\Psi_{\Phi,\gamma}^2$	Ψ_c^2	$[\Psi_c^2]_\Phi$	$\Psi_{\Phi,c}^2$
M_c^{pcs}	0.95	-3.83	3473937.4	0.947	-1.44
M_c^{eye}	0.92	0.86	3371224.1	0.919	0.89

that $M_{c,\gamma}^{\text{pcs}}$ (its γ -ray component) is, indeed, a bad fit to D_γ in the sense introduced in Section 4.1; i.e., it fits D_γ worse than B_γ does. Similarly, the joint GOF implied by χ_c^2 (encoded in both Ψ_c^2 and $[\Psi_c^2]_\Phi$) is greater for M_c^{pcs} than for M_c^{eye} .

As argued in Section 4.1.1, an effective joint GOF statistic would assign greater GOF to M_c^{eye} than to M_c^{pcs} , since M_c^{eye} implies greater model usefulness than does M_c^{pcs} . This means that, since neither Ψ_c^2 nor $[\Psi_c^2]_\Phi$ do so, the condition in Eq. (4.1) is not satisfied by either, just as the condition in Eq. (4.2) is not satisfied by χ_c^2 , and that both prove ineffectual as statistics with which to characterise model usefulness.

The joint GOF values assigned by $\Psi_{\Phi,c}^2$, however, do satisfy the condition in Eq. (4.1), and $\Psi_{\Phi,c}^2$ apparently effectively characterises model usefulness in this case. This does not, of course, constitute conclusive evidence that $\Psi_{\Phi,c}^2$ does indeed eliminate single-band dominance. In Chapter 8 a more comprehensive comparison of the performance of χ_c^2 and $\Psi_{\Phi,c}^2$ (with regard to single-band dominance) is conducted.

Chapter 5

The Vector Space Framework

More formally, $\Psi_{\Phi,c}^2$ can be derived by identifying precisely what it means to say that the single-band χ^2 statistic expresses DFO in units of the LC uncertainties and considering what this implies in the context of the definition of χ_c^2 in Eq. (3.10). To do this, a more robust mathematical framework for working with LCs is necessary, which this Chapter introduces.

5.1 Single-band Fits

5.1.1 Associating LCs with Vectors

An observed (single-band) LC with n phase bins, which is a sequence of intensities¹

$$D = (d_i \mathfrak{u})_n \quad (5.1)$$

in some measurement unit $\mathfrak{u}$ (e.g., “mJy” in the radio case),² can be viewed as a point in an n -dimensional space $\mathbb{K}^n$, and mapped onto a vector

$$\mathbf{d} = \sum_{i=1}^n d_i \mathbf{u}_i \quad (5.2)$$

in the real vector space of translations $\mathcal{L}$ associated with $\mathbb{K}^n$, with ordered, orthogonal basis

$$\mathcal{U} = \{\mathbf{u}_i\}_n, \quad (5.3)$$

composed of vectors that each have magnitude $|\mathbf{u}_i| = \sqrt{\mathbf{u}_i \cdot \mathbf{u}_i} = \mathfrak{u}$. This vector encodes the location (in $\mathbb{K}^n$) of D relative to the zero-flux LC $O = (0\mathfrak{u})_n$ with respect to $\mathcal{U}$ through the coordinate vector

$$\mathbf{d} = (d_i)_n \in \mathbb{R}^n. \quad (5.4)$$

A model can be thought of as a function that produces an LC

$$M(\theta) = (m_i \mathfrak{u})_n \in \mathbb{K}^n \quad (5.5)$$

for any given parameter combination θ in its parameter space $\mathcal{P}$, where each m_i is a function of θ . Since each θ does not necessarily correspond to a unique LC, $M(\theta)$ is, in general, non-injective with respect to the set of all possible LCs; i.e., a single point in $\mathbb{K}^n$ potentially corresponds to multiple points in $\mathcal{P}$ under $M(\theta)$. An example of this would be a model that predicts no pulsar-associated emission for multiple parameter combinations, effectively predicting the background-only LC $B = (b\mathfrak{u})_n$ for each of those parameter combinations (where $b\mathfrak{u}$ is the background level; assumed to be constant in phase).

¹A subscript to a sequence or set defined using set-builder notation indicates that the element indices range from 1 to the integer in the subscript (inclusive). Additionally, all sets containing i -subscripted elements are assumed to be ordered. Note that, unlike sequences, sets cannot contain repeated elements.

²A so-called “upright italic” font is used here to signify that $\mathfrak{u}$ is not a variable that can take a value, but is still a mathematical object.




Figure 5.1: A diagrammatical representation of the n -dimensional geometry associated with LC fitting in $\mathbb{K}^n$. The LCs D and $M \in \mathcal{M}$ (as well as O and B) are each associated with a point in this space, and a vector in the corresponding space of translations $\mathcal{L}$ (the vector associated with B is omitted). Also shown is the deviation vector s and the orthogonal basis $\mathcal{U}$ for $\mathcal{L}$. Note that the arrows drawn for the latter (which is a set of vectors) are merely representative; none of the vectors in $\mathcal{U}$ are necessarily coplanar with d , m and s . Similarly, the points O , B , D and M are not necessarily coplanar. The grey dashed line, which passes through B and is the locus of all the flat LCs in $\mathbb{K}^n$, is included as a reference.

Therefore, given a set of parameter combinations $\mathcal{P} \subset \mathcal{P}$, a model produces the set of candidate LCs

$$\mathcal{M} = \{M(\theta) : \theta \in \mathcal{P}'\} \subset \mathbb{K}^n, \quad (5.6)$$

where $\mathcal{P}'$ is any of the largest subsets of $\mathcal{P}$ for which $M(\theta_1) \neq M(\theta_2) \forall \theta_1, \theta_2 \in \mathcal{P}'$ (thereby ensuring that each $M \in \mathcal{M}$ is unique). Just as is the case for d , each $M \in \mathcal{M}$ can be associated with a vector

$$m = \sum_{i=1}^n m_i u_i \in \mathcal{L}, \quad (5.7)$$

which encodes its location relative to O with respect to $\mathcal{U}$ through the coordinate vector $m = (m_i)_n \in \mathbb{R}^n$.

Figure 5.1 depicts these associations diagrammatically, and reveals that, with respect to D , each $M \in \mathcal{M}$ can additionally be associated with a deviation vector

$$\begin{aligned} s &= m - d \\ &= \sum_{i=1}^n (m_i - d_i) u_i \in \mathcal{L}. \end{aligned} \quad (5.8)$$

Notice that this diagram attempts to depict the n -dimensional geometry of $\mathbb{K}^n$ directly.

The Difference between $\mathbb{K}^n$ and $\mathbb{R}^n$ Notice that this abstraction explicitly preserves the fact that the intensities that make up D are *physical* scalars, as opposed to *mathematical* scalars; i.e., that an observed LC is a sequence of *intensities*, not a sequence of *numbers*. The advantage of this, in light of the important role units play when evaluating joint GOF (as argued in Section 4.1), is that no *a priori* choice is made for the units in which these intensities are expressed (since $\mathcal{U}$ is not taken to be a *unit* basis), which means that any choices regarding units must be made explicitly.

5.1.2 Identifying the Best-fit LC

To determine which of the LCs in $\mathcal{M}$ is the best fit to D , an appropriate ordering must be imposed onto $\mathcal{M}$, which is generally unordered. The most natural way to achieve this, is to characterise each $M \in \mathcal{M}$ based on the deviation it corresponds to; specifically by defining a norm on $\mathcal{L}$ (a function that maps every vector in $\mathcal{L}$ onto $\mathbb{R}^+$; the non-negative real number) that can serve as a distance metric on $\mathbb{K}^n$.

Since every vector $k \in \mathcal{L}$ can be associated with a coordinate $k = (k_i)_n \in \mathbb{R}^n$ with respect to $\mathcal{U}$ (as demonstrated in

Eqs. (5.2) and (5.7)), the simplest option in this regard is to define the norm $\|\cdot\|_{\mathcal{U}}$ on $\mathcal{L}$ such that

$$\|\mathbf{k}\|_{\mathcal{U}} = \|\mathbf{k}\|, \quad (5.9)$$

where $\|\cdot\|$ is the Euclidean norm on $\mathbb{R}^n$. This norm gives rise to the distance metric

$$\begin{aligned} \delta_{\mathcal{U}}(K_1, K_2) &= \|\mathbf{k}_2 - \mathbf{k}_1\|_{\mathcal{U}} \\ &= \|k_2 - k_1\| \end{aligned} \quad (5.10)$$

on $\mathbb{K}^n$, where $\mathbf{k}_1, \mathbf{k}_2 \in \mathcal{L}$ are the vectors associated with the LCs $K_1, K_2 \in \mathbb{K}^n$, with coordinates $k_1, k_2 \in \mathbb{R}^n$ relative to O with respect to $\mathcal{U}$.

Applied to the deviations associated with $\mathcal{M}$, this metric enables the statistic³

$$\begin{aligned} S^2(M; D) &= \delta_{\mathcal{U}}^2(M, D) \\ &= \|\mathbf{s}\|_{\mathcal{U}}^2 \\ &= \sum_{i=1}^n (m_i - d_i)^2 \end{aligned} \quad (5.11)$$

with which to characterise each $M \in \mathcal{M}$. This statistic allows for the definition of a total order over $\mathcal{M}$, in the form of the binary relation $\geq_{S^2}$, defined such that

$$M_1 \geq_{S^2} M_2 \iff S^2(M_1) \leq S^2(M_2) \quad (5.12)$$

for $M_1, M_2 \in \mathcal{M}$, which encodes the intuitive idea that M_1 is a better fit to D than M_2 if it is closer to D than M_2 , and M_1 and M_2 are equally good fits to D if they are equidistant from D . Finding the best-fit LC (or LCs), then, reduces to minimising the value of $S^2(M)$ over $\mathcal{M}$. Specifically, any maximal element of the totally ordered set $(\mathcal{M}, \geq_{S^2})$, formed by equipping $\mathcal{M}$ with $\geq_{S^2}$, can be regarded as a best-fit LC with respect to D . If $(\mathcal{M}, \geq_{S^2})$ has a greatest element (i.e., only a single maximal element), a unique best-fit LC can be identified. Within the context of the geometrical representation of $\mathbb{K}^n$ in Figure 5.1, this is equivalent to determining which $M \in \mathcal{M}$ is closest to D under $\delta_{\mathcal{U}}$.

Accounting for the Uncertainties

This statistic, of course, does not take into account the uncertainties $E = (\varepsilon_i \mathbf{u})_n$ associated with the intensities that comprise D .⁴ Since each of these uncertainties expressly aims to characterise the deviation in its corresponding bin, a more appropriate statistic with which to characterise the overall deviation (of M from D) would be one that takes E into account. An example of such a statistic is the χ^2 statistic defined in Eq. (3.1), which scales the deviation in each bin by the uncertainty in that bin.

Within the vector-space framework, the definition of such a statistic is most elegantly realised by equipping $\mathbb{K}^n$ with a distance metric that somehow depends on E . The key to doing this effectively lies in understanding precisely how $\|\cdot\|_{\mathcal{U}}$, which gives rise to $\delta_{\mathcal{U}}$, maps vectors in $\mathcal{L}$ onto numbers in $\mathbb{R}^+$; specifically in realising that $\|\cdot\|_{\mathcal{U}}$, as defined in Eq. (5.9), effectively maps $\mathcal{U}$ onto $\{\mathbf{e}_i\}_n$, the natural basis for $\mathbb{R}^n$. This observation reveals that the use of $\delta_{\mathcal{U}}$ as the distance metric for $\mathbb{K}^n$, carries with it the implicit assumption that $\mathcal{U}$ is a good choice as the natural basis for $\mathcal{L}$. Since $\mathcal{U}$ does not depend on E , this is not the case when characterising deviation.

The definition of a more appropriate statistic, then, boils down to the identification of a more appropriate natural basis for $\mathcal{L}$ (with regard to deviation). The simplest such basis, derived from E , is the set of vectors

$$\mathcal{E} = \{\mathbf{e}_i = \varepsilon_i \mathbf{u}_i\}_n,$$

³Defined as a squared quantity in line with the definition of χ^2 in Eq. (3.1).

⁴Note that E is not an LC.

with respect to which each $\mathbf{k} \in \mathcal{L}$ can be rewritten

$$\mathbf{k} = \sum_{i=1}^n \frac{k_i}{\varepsilon_i} \boldsymbol{\epsilon}_i, \quad (5.13)$$

and the corresponding LC $K \in \mathbb{K}^n$ has coordinates

$$[k]_{\mathcal{E}} = \left(\frac{k_i}{\varepsilon_i} \right)_n \quad (5.14)$$

relative to O with respect to $\mathcal{E}$. This basis, therefore, gives rise to the norm $\|\cdot\|_{\mathcal{E}}$ on $\mathcal{L}$, defined such that

$$\|\mathbf{k}\|_{\mathcal{E}} = \|[k]_{\mathcal{E}}\|, \quad (5.15)$$

and the distance metric

$$\begin{aligned} \delta_{\mathcal{E}}(K_1, K_2) &= \|\mathbf{k}_2 - \mathbf{k}_1\|_{\mathcal{E}} \\ &= \|[k_2]_{\mathcal{E}} - [k_1]_{\mathcal{E}}\|, \end{aligned} \quad (5.16)$$

which, essentially, measures the distance between two LCs K_1 and K_2 in units of the LC uncertainties.

Applied to the deviations associated with $\mathcal{M}$, this metric, of course, enables the familiar statistic

$$\begin{aligned} \chi^2(M; D, E) &= \delta_{\mathcal{E}}^2(M, D) \\ &= \|\mathbf{s}\|_{\mathcal{E}}^2 \\ &= \sum_{i=1}^n \left(\frac{m_i - d_i}{\varepsilon_i} \right)^2, \end{aligned} \quad (5.17)$$

which results in the totally ordered set $(\mathcal{M}, \geq_{\chi^2})$, where $\geq_{\chi^2}$ is defined (analogously to $\geq_{S^2}$) such that

$$M_1 \geq_{\chi^2} M_2 \iff \chi^2(M_1) \leq \chi^2(M_2). \quad (5.18)$$

The best-fit LC (or LCs) in $\mathcal{M}$ can again be identified by finding the greatest element (or the set of maximal elements, should it not have a greatest element) of $(\mathcal{M}, \geq_{\chi^2})$. Within the context of Figure 5.1, this is equivalent to determining which $M \in \mathcal{M}$ is closest to D under $\delta_{\mathcal{E}}$.

5.1.3 The Relationship between S^2 and χ^2

The definition of $S^2(M)$ and $\chi^2(M)$ in terms of distance metrics on $\mathbb{K}^n$ ($\delta_{\mathcal{U}}$ and $\delta_{\mathcal{E}}$, respectively; Eqs. (5.11) and (5.17)), enables a useful geometrical representation of the relationship between these two statistics. Taken as functions that map $\mathbb{K}^n$ onto $\mathbb{R}^n$, the coordinate representations for the various LC-associated vectors, under the operators $[\cdot]_{\mathcal{U}}$ (implicit; e.g., Eq. (5.2)) and $[\cdot]_{\mathcal{E}}$ (e.g., Eq. (5.13)), reveal that the definition of $\chi^2(M)$ leverages a mapping that is an anisotropic scaling (centred on the LC O , with scale factors $\{\varepsilon_i\}_n$) of that leveraged in the definition of $S^2(M)$. The key, definitional quality of this scaling is that it renders the n -dimensional normal distribution that E defines (centred on D_{true}), spherically symmetric.

This relationship is depicted diagrammatically in Figure 5.2: In the left-hand mapping (which results in S^2), $\mathcal{U}$ maps onto the semi-axes of a unit hypersphere (centred on O) and $\mathcal{E}$ defines a hyperellipse, while in the right-hand mapping (an anisotropic scaling of the left-hand mapping; which results in χ^2), $\mathcal{E}$ itself maps onto the semi-axes of a unit hypersphere (centred on D_{true}). In this sense, $S^2(M)$ measures DFO in units of $\mathcal{U}$ and $\chi^2(M)$ measures it in units of $\mathcal{E}$. Notice that, similarly to Figure 5.1, these diagrams attempt to depict the n -dimensional geometry of these mappings directly.




Figure 5.2: The effective mappings of $\mathbb{K}^n$ onto $\mathbb{R}^n$ under the operators $[\cdot]_{\mathcal{U}}$ (left) and $[\cdot]_{\mathcal{E}}$ (right). In the former mapping, $\mathcal{U}$ maps onto the semi-axes of a unit hypersphere (the black, dashed circle centred on O), and $\mathcal{E}$ onto the semi-axes of a hyperellipsoid (the grey, dashed ellipse centred on D_{true}); therefore the distance metric $\delta_{\mathcal{U}}$ effectively carries units of $\mathcal{U}$. In the latter mapping, which is an anisotropic scaling of the former mapping, the reverse is true: $\mathcal{E}$ maps onto the semi-axes of a unit hypersphere and $\mathcal{U}$ onto the semi-axes of a hyperellipsoid. Therefore, the χ^2 statistic (which results from the latter mapping) carries units of the LC errors through $\mathcal{E}$. In each mapping, the grey dashed line (again included as a reference) is the locus of all the LCs in $\mathbb{K}^n$ that are rendered flat by the relevant mapping operator (i.e., by $[\cdot]_{\mathcal{U}}$ in the mapping on the left and by $[\cdot]_{\mathcal{E}}$ in the mapping on the right).




Figure 5.3: The effective mappings of $\mathbb{K}^n$ onto $\mathbb{R}^n$ under the operators $[\cdot]_{\mathcal{E}}$ (left) and $[\cdot]_{\mathcal{E}_{\Phi}}$ (right). In contrast to the mappings shown in Figure 5.2, D and M are each additionally associated with a displacement vector relative to B : the perturbation $\mathbf{p}$ (with magnitude Φ under $\|\cdot\|_{\mathcal{E}}$) and $\mathbf{m}^*$. In the former mapping, just as in Figure 5.2, $\mathcal{E}$ maps onto the semi-axes of a unit hypersphere (the black, dashed circle centred on D_{true}). In this mapping, the locus of all possible background-equivalent LCs (for which $\chi^2 = \Phi^2$) is also a hypersphere, centred on D (the grey, dashed circle), albeit a somewhat larger hypersphere. In the latter mapping, it is instead this locus that is a unit hypersphere, thereby effectively normalising the magnitude of $\mathbf{p}$ (i.e., the scaled flux), and granting χ_{Φ}^2 and Ψ_{Φ}^2 explicit GOF units (as per Section 4.3). Notice that the equivalence between χ^2 and Ψ_{Φ}^2 (with regard to best-fit identification; under the appropriate marginalisations) is reflected in the fact that this locus is a hypersphere in both mappings.

Equivalence of Statistics

The choice made in Eq. (5.17) for the unit basis of $\mathcal{L}$ is not the only choice that results in a totally ordered set where the elements of $\mathcal{M}$ are ordered as in $(\mathcal{M}, \geq_{\chi^2})$. Since the definition of the total order $\geq_{\chi^2}$ relies on the $\leq$ relation (a total order on $\mathbb{R}$), any choice that results in a statistic that is proportional to χ^2 is equivalent to the choice made in Eq. (5.17). Understood geometrically, any choice for the natural basis of $\mathcal{L}$ that results in a mapping wherein $\mathcal{E}$ maps onto the semi-axes of a hypersphere (other than a unit hypersphere), will result in a statistic that is equivalent to χ^2 with regard to the ordering it imposes on $\mathcal{M}$, and therefore with regard to the best fit (or fits) it would identify.

An example of such an equivalent statistic is χ_ν^2 , which results when the basis

$$\mathcal{E}_\nu = \sqrt{\nu}\mathcal{E} = \{\sqrt{\nu}\epsilon_i = \sqrt{\nu}\epsilon_i \mathbf{u}_i\}_n \quad (5.19)$$

is chosen to serve as natural basis for $\mathcal{L}$. Under this choice, $\mathcal{E}$ maps onto the semi-axes of the hypersphere in $\mathbb{R}^n$ with radius $(1/\sqrt{\nu})$. Just as was the case for $\mathcal{U}$ and $\mathcal{E}$, this choice gives rise to a distance metric on $\mathbb{K}^n$: the distance metric

$$\delta_{\mathcal{E}_\nu}(K_1, K_2) = \|[k_2]_{\mathcal{E}_\nu} - [k_1]_{\mathcal{E}_\nu}\|, \quad (5.20)$$

where

$$\begin{aligned} [k]_{\mathcal{E}_\nu} &= \left(\frac{k_i}{\sqrt{\nu}\epsilon_i} \right)_n \\ &= \frac{1}{\sqrt{\nu}} [k]_{\mathcal{E}} \in \mathbb{R}^n \end{aligned} \quad (5.21)$$

is the coordinate vector, with respect to $\mathcal{E}_\nu$, associated with an LC $K \in \mathbb{K}^n$. Written in terms of this metric,

$$\begin{aligned} \chi_\nu^2(M) &= \delta_{\mathcal{E}_\nu}^2(M, D) \\ &= \sum_{i=1}^n \left(\frac{m_i - d_i}{\sqrt{\nu}\epsilon_i} \right)^2 \\ &= \frac{1}{\nu} \chi^2(M), \end{aligned} \quad (5.22)$$

just as in Eq. (3.6).

5.1.4 Scaled-flux Normalised GOF

The Perturbation and the Scaled Flux

The notion introduced in Section 4.2, that the pulsar's emission essentially perturbs the background LC B (constituted by the background level b), to produce D , comes into sharp focus within the vector space framework. Since both B and D are points in $\mathbb{K}^n$, this perturbation can be associated with the vector

$$\begin{aligned} \mathbf{p} &= \mathbf{d} - \mathbf{b} \\ &= \sum_{i=1}^n (d_i - b_i) \mathbf{u}_i \in \mathcal{L}, \end{aligned} \quad (5.23)$$

which connects B and D , as shown in Figure 5.3. The magnitude of this vector under $\|\cdot\|_{\mathcal{E}}$ (i.e., its magnitude in units of $\mathcal{E}$), is

$$\Phi = \|\mathbf{p}\|_{\mathcal{E}} = \chi(B); \quad (5.24)$$

the scaled flux identified in Section 4.2. Since a model that yields $\chi_{\min}^2 = \Phi^2$ is 'useless' as a description for the pulsar (as discussed in Section 4.1), and a model that yields $\chi_{\min}^2 = 0$ is 'maximally useful,' Φ^2 can be viewed as a GOF standard against which the values that $\chi^2(M)$ yields, should be compared, especially with regard to implied model usefulness.

The Scaled-flux Normalised Ψ_Φ^2 Statistic

As discussed in Section 4.3, the identification of $\mathbf{p}$ allows for the definition of the scaled-flux normalised GOF statistic $\Psi_\Phi^2(M)$. Within the vector space framework, upon choosing

$$\mathcal{E}_\Phi = \Phi \mathcal{E} = \{\Phi \varepsilon_i \mathbf{u}_i\}_n \quad (5.25)$$

as the natural basis for $\mathcal{L}$ (which gives rise to the norm $\|\cdot\|_{\mathcal{E}_\Phi}$ on $\mathcal{L}$ and the distance metric $\delta_{\mathcal{E}_\Phi}$ on $\mathbb{K}^n$),

$$\begin{aligned} \Psi_\Phi^2(M; D, E, B) &= \|\mathbf{p}\|_{\mathcal{E}_\Phi}^2 - \delta_{\mathcal{E}_\Phi}^2(M, D) \\ &= 1 - \frac{1}{\Phi^2} \chi^2(M; D, E). \end{aligned} \quad (5.26)$$

Notice that upon choosing $\mathcal{E}_\Phi$ as the natural basis for $\mathcal{L}$, the magnitude of the perturbation (and hence the scaled flux) is effectively normalised since $\|\mathbf{p}\|_{\mathcal{E}_\Phi} = 1$. In fact, $\mathcal{E}_\Phi$ arises naturally if this property is taken as a constraint alongside the mandate that the resulting statistic should be equivalent to the χ^2 statistic (under maximisation; i.e., that it should impose the same ordering on $\mathcal{M}$ under maximisation as the χ^2 statistic does under minimisation).

As also discussed in Section 4.3, this choice for the natural basis of $\mathcal{L}$ also gives rise to the scaled-flux normalised DFO statistic

$$\begin{aligned} \chi_\Phi^2(M; D, E, B) &= \delta_{\mathcal{E}_\Phi}^2(M, D) \\ &= \frac{1}{\Phi^2} \chi^2(M; D, E), \end{aligned} \quad (5.27)$$

which is the scaled-flux normalised equivalent of $\chi^2(M; D, E)$. Considered together, Eqs. (5.26) and (5.27) encode the notion that DFO and GOF are complementary (which underpins the use of the χ^2 statistic as a GOF statistic in Pearson's χ^2 test, as discussed in Section 4.1), since

$$\chi_\Phi^2(M) + \Psi_\Phi^2(M) = 1. \quad (5.28)$$

Geometrically, these normalised statistics can be understood by recognising that under $[\cdot]_{\mathcal{E}}$ the locus of all the possible LCs that would result in a null fit to D (i.e., for which $\chi^2(K) = \Phi^2$) is a hypersphere in $\mathbb{R}^n$ (centred on D), with radius Φ . Since these LCs also lie on a hypersphere under $[\cdot]_{\mathcal{E}_\Phi}$ (a unit hypersphere, in fact), $\chi_\Phi^2(M)$ is equivalent to $\chi^2(M)$. This equivalence also extends to $\Psi_\Phi^2(M)$ (under maximisation) since it can also be associated with a distance in $\mathbb{K}^n$; specifically, the distance from M to the null hypersphere along a line that is perpendicular to $\mathbf{s}$ (the relevant deviation vector; as shown in Figure 5.3). Notice that this distance is not the distance from M to any specific point in $\mathbb{K}^n$, but rather the distance from M to a surface in $\mathbb{K}^n$ (the null hypersphere). This distinction is the reason for the form of Eq. (5.26), where Ψ_Φ^2 is defined in terms of the magnitude of $\mathbf{p}$ instead of the distance between B and D (i.e., in terms of $\|\mathbf{p}\|_{\mathcal{E}_\Phi}$ instead of $\delta_{\mathcal{E}_\Phi}(B, D)$).

It is worth noting here that it is due to these equivalences that $\chi^2(M)$ can be said to satisfy the criterion in Eq. (4.2) in the single-band case, since χ_Φ^2 and Ψ_Φ^2 are more directly connected to usefulness, through their connection to the Φ .

5.2 Extended to Dual-band Fits

An important, yet subtle, distinction between Eqs. 3.1 and 5.17 is that the former defines χ^2 as a novel quantity (derived from D , M , and E), while the latter identifies χ^2 as the square of the distance between D and M in units of E (through $\mathcal{E}$). This identification is not only accompanied by that of an entire family of χ^2 -equivalent statistics (which includes χ_ν^2), but also by that of the scaled-flux normalised Ψ_Φ^2 statistic, which expresses GOF in units that normalise the magnitude of the perturbation $\mathbf{p}$.

While this distinction is trivial for single-band fits, since it in no way impacts the ordering imposed on $\mathcal{M}$, it becomes important for dual-band fits, where joint GOF is typically characterised by taking the sum of the single-band-only statistics. To understand precisely how this distinction manifests for dual-band fits, and what its effect is, it is useful to understand how χ_c^2 , as defined in Eq. (3.10), comes about within the vector space framework.

5.2.1 Associating LC Pairs with Vectors

As discussed in Section 3.1.4, the utility of Pearson's χ^2 statistic can be extended to the evaluation of joint, dual-band fits by effectively treating the pair of observed radio and γ -ray LCs D_r and D_γ , with n_r and n_γ phase bins and uncertainties E_r and E_γ , as a single, concatenated LC

$$\begin{aligned} D_c &= D_r \# D_\gamma \\ &= (d_{r,i}u_r)_{n_r} \# (d_{\gamma,i}u_\gamma)_{n_\gamma} \\ &= (d_{r,1}u_r, \dots, d_{r,n_r}u_r, d_{\gamma,1}u_\gamma, \dots, d_{\gamma,n_\gamma}u_\gamma), \end{aligned} \quad (5.29)$$

with $n_c = n_r + n_\gamma$ phase bins and associated uncertainties $E_c = E_r \# E_\gamma$ (where the “ $\#$ ” symbol denotes concatenation). In analogue to the single-band case, this LC can be thought of as a point in the n_c -dimensional space

$$\mathbb{K}^{n_c} = \mathbb{K}^{n_r} \oplus \mathbb{K}^{n_\gamma}, \quad (5.30)$$

which is the direct sum of the n_r - and n_γ -dimensional single-band spaces, and associated with a vector

$$\begin{aligned} \mathbf{d}_c &= \mathbf{d}_r + \mathbf{d}_\gamma \\ &= \sum_{i=1}^{n_r} d_{r,i} \mathbf{u}_{r,i} + \sum_{i=1}^{n_\gamma} d_{\gamma,i} \mathbf{u}_{\gamma,i}, \end{aligned} \quad (5.31)$$

where $\mathbf{d}_r \in \mathcal{L}_r$ and $\mathbf{d}_\gamma \in \mathcal{L}_\gamma$, in the real vector space of translations $\mathcal{L}_c = \mathcal{L}_r \oplus \mathcal{L}_\gamma$ associated with $\mathbb{K}^{n_c}$, with ordered, orthogonal basis

$$\mathcal{U}_c = \mathcal{U}_r \# \mathcal{U}_\gamma, \quad (5.32)$$

where $\mathcal{U}_r = \{\mathbf{u}_{r,i}\}_{n_r}$ and $\mathcal{U}_\gamma = \{\mathbf{u}_{\gamma,i}\}_{n_\gamma}$ (composed of vectors of magnitude u_r and u_γ , respectively). The vector $\mathbf{d}_c$ encodes the position of D_c relative to the zero-flux (concatenated) LC $O_c = O_r \# O_\gamma$, and the sequence of coefficients

$$\mathbf{d}_c = (d_{r,i})_{n_r} \# (d_{\gamma,i})_{n_\gamma} \in \mathbb{R}^{n_c} \quad (5.33)$$

can be thought of as the coordinates of D_c relative to O_c with respect to $\mathcal{U}_c$.

A dual-band model can be thought of as a (non-injective) function $M_c(\vartheta)$ that produces a concatenated LC

$$\begin{aligned} M_c(\vartheta) &= M_r(\vartheta) \# M_\gamma(\vartheta) \\ &= (m_{r,i}u_r)_{n_r} \# (m_{\gamma,i}u_\gamma)_{n_\gamma}, \end{aligned} \quad (5.34)$$

where every $m_{r,i}, m_{\gamma,i} \in \mathbb{R}$ is a function of ϑ , for each parameter combination ϑ in its parameter space $\mathcal{Q}$. Given a set of parameter combinations $\mathcal{Q}' \in \mathcal{Q}$, such a dual-band model produces a set of candidate (concatenated) LCs

$$\mathcal{M}_c = \{M_c(\vartheta) : \vartheta \in \mathcal{Q}'\}, \quad (5.35)$$

where $\mathcal{Q}'$ is any of the largest subsets of $\mathcal{Q}$ for which $M_c(\vartheta_1) \neq M_c(\vartheta_2) \forall \vartheta_1, \vartheta_2 \in \mathcal{Q}'$. Just as is the case for $\mathbf{d}_c$, each $M_c \in \mathcal{M}_c$ can be associated with a vector

$$\begin{aligned} \mathbf{m}_c &= \mathbf{m}_r + \mathbf{m}_\gamma \\ &= \sum_{i=1}^{n_r} m_{r,i} \mathbf{u}_{r,i} + \sum_{i=1}^{n_\gamma} m_{\gamma,i} \mathbf{u}_{\gamma,i}, \end{aligned} \quad (5.36)$$

where $\mathbf{m}_r \in \mathcal{L}_r$ and $\mathbf{m}_\gamma \in \mathcal{L}_\gamma$ are the vectors associated with M_r and M_γ .

The geometry of $\mathbb{K}^{n_c}$ is (essentially by construction) exactly analogous to the geometry of $\mathbb{K}^n$ in the single-band case,

and hence each $M_c \in \mathcal{M}_c$ can be associated with a deviation vector

$$\begin{aligned} \mathbf{s}_c &= \mathbf{m}_c - \mathbf{d}_c \\ &= \mathbf{s}_r + \mathbf{s}_\gamma \in \mathcal{L}_c, \end{aligned} \quad (5.37)$$

where $\mathbf{s}_r \in \mathcal{L}_r$ and $\mathbf{s}_\gamma \in \mathcal{L}_\gamma$ are the deviation vectors associated with $\mathbf{m}_r$ and $\mathbf{m}_\gamma$.

5.2.2 Identifying the Joint Best-fit LC Pair using χ_c^2

The first step toward defining χ_c^2 within the vector space framework, is to choose

$$\begin{aligned} \mathcal{E}_c &= \mathcal{E}_r \uplus \mathcal{E}_\gamma \\ &= \{\epsilon_{r,i}\}_{n_r} \uplus \{\epsilon_{\gamma,i}\}_{n_\gamma}, \end{aligned} \quad (5.38)$$

where each $\epsilon_{r,i} = \varepsilon_{r,i} \mathbf{u}_{r,i}$ and each $\epsilon_{\gamma,i} = \varepsilon_{\gamma,i} \mathbf{u}_{\gamma,i}$ (derived from E_c), as the natural basis for $\mathcal{L}_c$. This choice induces the norm $\|\cdot\|_{\mathcal{E}_c}$ on $\mathcal{L}_c$, defined such that, for a given vector $\mathbf{k}_c \in \mathcal{L}_c$,

$$\begin{aligned} \|\mathbf{k}_c\|_{\mathcal{E}_c} &= \|[k_c]_{\mathcal{E}_c}\| \\ &= \sqrt{\|[k_r]_{\mathcal{E}_r}\|^2 + \|[k_\gamma]_{\mathcal{E}_\gamma}\|^2}, \end{aligned} \quad (5.39)$$

where $[k_c]_{\mathcal{E}_c}$ is that vector's coordinates relative to O_c with respect to $\mathcal{E}_c$, and results in the distance metric (given a pair of concatenated LCs $K_{c,1}, K_{c,2} \in \mathbb{K}^{n_c}$)

$$\begin{aligned} \delta_{\mathcal{E}_c}(K_{c,1}, K_{c,2}) &= \|\mathbf{k}_{c,2} - \mathbf{k}_{c,1}\|_{\mathcal{E}_c} \\ &= \|[k_{c,2}]_{\mathcal{E}_c} - [k_{c,1}]_{\mathcal{E}_c}\| \\ &= \sqrt{\delta_{\mathcal{E}_r}^2(K_{r,1}, K_{r,2}) + \delta_{\mathcal{E}_\gamma}^2(K_{\gamma,1}, K_{\gamma,2})} \end{aligned} \quad (5.40)$$

on $\mathbb{K}^{n_c}$.

Applied to the deviations the LCs in $\mathcal{M}_c$ correspond to, this metric enables the statistic

$$\begin{aligned} \chi_c^2(M_c; D_c, E_c) &= \delta_{\mathcal{E}_c}^2(M_c, D_c) \\ &= \delta_{\mathcal{E}_r}^2(M_r, D_r) + \delta_{\mathcal{E}_\gamma}^2(M_\gamma, D_\gamma) \\ &= \chi_r^2(M_r; D_r, E_r) + \chi_\gamma^2(M_\gamma; D_\gamma, E_\gamma), \end{aligned} \quad (5.41)$$

where χ_r^2 and χ_γ^2 are the single-band statistics that arise for the radio and γ -ray cases. This expression for χ_c^2 is, of course, equivalent to that in Eq. (3.10) under the definitions in this Section for the various quantities.

5.2.3 Component Perturbations and Dominance

As argued in Section 4.4, when evaluated with regard to GOF this statistic is effectively weighted in favour of the waveband with the more precise observed LC. This weighting is reflected in the relative magnitudes of the radio and γ -ray components of the perturbation

$$\begin{aligned} \mathbf{p}_c &= \mathbf{d}_c - \mathbf{b}_c \\ &= \mathbf{p}_r + \mathbf{p}_\gamma \in \mathcal{L}_c \end{aligned} \quad (5.42)$$

the observed LC pair corresponds to under $\|\cdot\|_{\mathcal{E}_c}$, where $\mathbf{p}_r \in \mathcal{L}_r$ and $\mathbf{p}_\gamma \in \mathcal{L}_\gamma$; i.e., in the fact that

$$\frac{\|\mathbf{p}_r\|_{\mathcal{E}_r}}{\|\mathbf{p}_\gamma\|_{\mathcal{E}_\gamma}} = \frac{\Phi_r}{\Phi_\gamma} = \lambda_{r\gamma} \quad (5.43)$$




Figure 5.4: The effective mappings of $\mathbb{K}^{n_c}$ onto $\mathbb{R}^{n_c}$ under the operators $[\cdot]_{E_c}$ (left) and $[\cdot]_{E_c^*}$ (right). In the former mapping, just as is the case for $\mathcal{E}$ under $[\cdot]_{\mathcal{E}}$ in Figure 5.3, $\mathcal{E}_c$ maps onto the semi-axes of a unit hypersphere (in $\mathbb{R}^{n_c}$; the black, dashed circle centred on $D_{c,\text{true}}$). The locus of all possible LC pairs for which $\Psi_{\Phi,r}^2 = 1$ and $\Psi_{\Phi,\gamma}^2 = 0$ (e.g., F_c^{pcs}) is an n_γ -dimensional hypersphere centred on D_c ; in effect, this hypersphere is the γ -ray instantiation of the locus plotted in grey in the left-hand panel of Figure 5.3. Here, this locus presents as a one-dimensional sphere with radius Φ_γ (in the left-hand pane), constituted by the hollow grey diamonds directly above and below D_c (and labelled “ F_c^{pcs} ”). The n_r -dimensional hypersphere that is the locus of all possible LC pairs for which $\Psi_{\Phi,r}^2 = 0$ and $\Psi_{\Phi,\gamma}^2 = 1$ (e.g., G_c^{pcs}) similarly presents as a one-dimensional sphere with radius Φ_r in the left-hand panel (labelled “ G_c^{pcs} ”, and depicted using hollow grey squares directly to the left and right of D_c). Together, these two loci constitute a hyperellipse in $\mathbb{R}^{n_c}$ (the smaller grey dashed ellipse in the left-hand mapping), which is the locus of all possible LC pairs that have the same GOF as F_c^{pcs} and G_c^{pcs} (under the assumption that GOF is characterised by a ξ -equivalent statistic; as per the argument in Section 4.4.2). In the right-hand mapping, both this locus and the locus of all possible LC pairs that are similarly equivalent to B_c , are hyperspheres (as a result of the constraint in Eq. (5.45)); the latter rendered a unit hypersphere by the constraint in Eq. (5.46). Notice that the eccentricity of the smaller ellipse in the left-hand panel is effectively determined by the relative lengths of p_r and p_γ , and its aspect ratio is simply $\Phi_r/\Phi_\gamma = \lambda_{r\gamma}$ (the aspect ratio of its analogue in the right-hand panel is, of course, 1).

is not necessarily equal to 1.

Geometrically, as shown in the left-hand pane of Figure 5.4, this weighting is reflected in the fact that the best-fit LC pairs associated with the two equally useful hypothetical dual-band models in Section 4.4.2, F_c^{PCS} and G_c^{PCS} , map onto points in $\mathbb{R}^{n_c}$ that lie on a hyperellipse centred on D_c under $[\cdot]_{\mathcal{E}_c}$. Since these two LC pairs are equivalent with regard to implied model usefulness, they would instead map onto points that lie on a hypersphere (centred on D_c) under a mapping that is unweighted with regard to GOF (as would be the case if $\lambda_{r\gamma}$ were 1). More generally, the locus of all the possible concatenated LCs that are equivalent to B_c under $\Psi_{\Phi,c}^2$ (i.e., for which $\Psi_{\Phi,c}^2 = 0$), is a hyperellipse in $\mathbb{R}^{n_c}$ under $[\cdot]_{\mathcal{E}_c}$, with hyperspherical images in $\mathbb{R}^{n_r}$ and $\mathbb{R}^{n_\gamma}$ (with radii Φ_r and Φ_γ , respectively). Therefore, $\lambda_{r\gamma}$ effectively characterises the uncertainty disparity by quantifying how much these hyperellipses deviate from being hyperspherical as a result thereof.

5.2.4 The Scaled-flux Standardised $\Psi_{\Phi,c}^2$ Statistic

As argued in Section 4.5, a more appropriate extension of the single-band χ^2 statistic with regard to parameter estimation is one that grants equal weight to the radio and γ -ray GOF components; i.e., to $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$. Within the vector space framework, this assertion translates to the assertion that a more appropriate choice for the natural basis for $\mathcal{L}_c$, is a basis

$$\begin{aligned}\mathcal{E}_c^* &= \mathcal{E}_r^* \# \mathcal{E}_\gamma^* \\ &= \ell_r \mathcal{E}_r + \ell_\gamma \mathcal{E}_\gamma,\end{aligned}\tag{5.44}$$

where $\ell_r, \ell_\gamma > 0$ are constants (which ensures that the resulting dual-band statistic is, indeed, linear in χ_r^2 and χ_γ^2 ; or, equivalently, in $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$), chosen such that the magnitudes of $\mathbf{p}_r$ and $\mathbf{p}_\gamma$, as measured under this choice, are equal (see also Section 6.1); i.e., chosen such that

$$\frac{\|\mathbf{p}_r\|_{\mathcal{E}_r^*}}{\|\mathbf{p}_\gamma\|_{\mathcal{E}_\gamma^*}} = \frac{\Phi_r}{\ell_r} \bigg/ \frac{\Phi_\gamma}{\ell_\gamma} = 1\tag{5.45}$$

where $\|\cdot\|_{\mathcal{E}_r^*}$ and $\|\cdot\|_{\mathcal{E}_\gamma^*}$ are the norms on $\mathcal{L}_r$ and $\mathcal{L}_\gamma$ that result from this choice. This single constraint is not, however, enough to fully determine ℓ_r and ℓ_γ ; a second constraint must be brought to bear on the choice. The simplest, and arguably most natural, constraint to introduce, is to mandate that the choice be such that the magnitude of $\mathbf{p}_c$ is 1 under this choice; i.e., to mandate that

$$\|\mathbf{p}_c\|_{\mathcal{E}_c^*} = \sqrt{\left(\frac{\Phi_r}{\ell_r}\right)^2 + \left(\frac{\Phi_\gamma}{\ell_\gamma}\right)^2} = 1,\tag{5.46}$$

where $\|\cdot\|_{\mathcal{E}_c^*}$ is the norm on $\mathcal{L}_c$ that results from this choice. Within the context of parameter estimation, this constraint is, of course, superfluous, since all possible χ_c^2 -like statistics that satisfy the constraint in Eq. (5.45) are essentially equivalent by construction (i.e., result in the same ordering for $\mathcal{M}_c$, and yield the same best fit).

Solving these constraints, $\ell_r = \sqrt{2} \Phi_r$ and $\ell_\gamma = \sqrt{2} \Phi_\gamma$, and therefore

$$\mathcal{E}_c^* = \left(\sqrt{2}\Phi_r\right) \mathcal{E}_r + \left(\sqrt{2}\Phi_\gamma\right) \mathcal{E}_\gamma.\tag{5.47}$$

The GOF statistic that results upon choosing this basis as the natural basis for $\mathcal{L}_c$ is the scaled-flux standardised statistic

$$\begin{aligned}\Psi_{\Phi,c}^2(M_c; D_c, E_c, B_c) &= \|\mathbf{p}_c\|_{\mathcal{E}_c^*}^2 - \delta_{\mathcal{E}_c^*}^2(M_c, D_c) \\ &= \|\mathbf{p}_c\|_{\mathcal{E}_c^*}^2 - \left[\delta_{\mathcal{E}_r^*}^2(M_r, D_r) + \delta_{\mathcal{E}_\gamma^*}^2(M_\gamma, D_\gamma)\right] \\ &= 1 - \frac{1}{2} \left[\frac{1}{\Phi_r^2} \chi_r^2(M_r; D_r, E_r) + \frac{1}{\Phi_\gamma^2} \chi_\gamma^2(M_\gamma; D_\gamma, E_\gamma) \right],\end{aligned}\tag{5.48}$$

where $\|\cdot\|_{\mathcal{E}_c^*}$ and $\delta_{\mathcal{E}_c^*}$ are the norm on $\mathcal{L}_c$ and distance metric on $\mathbb{K}^{n_c}$ associated with $\mathcal{E}_c^*$, and $\delta_{\mathcal{E}_r^*}$ and $\delta_{\mathcal{E}_\gamma^*}$ are the distance metrics on $\mathbb{K}^{n_r}$ and $\mathbb{K}^{n_\gamma}$ associated with $\mathcal{E}_r^*$ and $\mathcal{E}_\gamma^*$; this is precisely the statistic defined in Eq. (4.20). Note that, since the normalisation constraint is invoked in deriving it, $\Psi_{\Phi,c}^2$ is also a scaled-flux normalised statistic.

Just as was the case for Ψ_{Φ}^2 , the constraints that give rise to $\Psi_{\Phi,c}^2$ also give rise to a DFO statistic:

$$\begin{aligned}\chi_{\Phi,c}^2(M_c; D_c, E_c, B_c) &= \delta_{\mathcal{E}_c^*}^2(M_c, D_c) \\ &= \frac{1}{2} \left[\frac{1}{\Phi_r^2} \chi_r^2(M_r; D_c, E_c) + \frac{1}{\Phi_\gamma^2} \chi_\gamma^2(M_\gamma; D_c, E_c) \right],\end{aligned}\tag{5.49}$$

which can be considered the complement to $\Psi_{\Phi,c}^2$ since $\chi_{\Phi,c}^2(M_c) + \Psi_{\Phi,c}^2(M_c) = 1$.

Effective Radio-only Uncertainty Scaling An interesting alternative choice for ℓ_r and ℓ_γ exists when the radio LC uncertainty in each bin is the same, as is typically assumed for pulsar fits. Specifically, choosing $\ell_r = \lambda_{r\gamma}$ and $\ell_\gamma = 1$ reveals that using $\Psi_{\Phi,c}^2$ to identify the model's best fit is effectively equivalent to inflating the radio uncertainties (or deflating them, if $\Phi_r < \Phi_\gamma$) by a factor $\lambda_{r\gamma}$ and using χ_c^2 to identify the model's best fit; i.e., equivalent to conducting the fit using χ_c^2 but assuming a uniform radio uncertainty of

$$\varepsilon_{r,\text{sfs}} = \lambda_{r\gamma} \varepsilon_r.\tag{5.50}$$

Understood Geometrically

Geometrically, as shown in the right-hand panel in Figure 5.4, $\mathcal{E}_c^*$ is simply the natural basis for $\mathcal{L}_c$ under which the locus of all possible background-equivalent LCs in $\mathbb{K}^{n_c}$ maps onto a unit hypersphere in $\mathbb{R}^{n_c}$. In so doing, this choice ensures that GOF is standardised across both wavebands, and that the joint GOF statistic is normalised. The alternative choice for ℓ_r and ℓ_γ described above effectively relaxes the requirement that the joint GOF statistic be normalised, and results in a mapping onto $\mathbb{R}^{n_c}$ wherein the locus of all possible background-equivalent fits maps onto a hypersphere that is not a unit hypersphere.

Chapter 6

Gauging a Statistic's Effectiveness

Before the effectiveness of $\Psi_{\Phi,c}^2$ can be evaluated, specifically with regard to its supposed elimination of single-band dominance in joint LC fits, it is first necessary to establish precisely what it means for such a joint fit to be radio- or γ -ray-dominated within the context of the discussion in Section 3.2.2. Ideally, this should be done in such a manner that the effectiveness of both $\Psi_{\Phi,c}^2$ and χ_c^2 can be evaluated. To facilitate this, it is useful to more closely examine the relationship between these two statistics, specifically with regard to how they identify a dual-band model's best fit.

6.1 The Relationship between χ_c^2 and $\Psi_{\Phi,c}^2$

The expressions for $[\Psi_c^2]_{\Phi}$ and $\Psi_{\Phi,c}^2$ in terms of $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$ (in Eqs. (4.10) and (4.20)), suggest that, for a given pulsar, they can be thought of as members of a family of statistics, parameterised by the relative radio weight $\omega_r = w_r/w_\gamma > 0$, where $w_r, w_\gamma \in (0, 1)$ are the normalised radio and γ -ray weights (i.e., by definition $w_r + w_\gamma = 1$). The simplest such parameterisation is the generalised statistic

$$\begin{aligned}\xi(M_c; \omega_r) &= w_r \Psi_{\Phi,r}^2(M_r) + w_\gamma \Psi_{\Phi,\gamma}^2(M_\gamma) \\ &= \frac{1}{1 + \omega_r} [\omega_r \Psi_{\Phi,r}^2(M_r) + \Psi_{\Phi,\gamma}^2(M_\gamma)],\end{aligned}\tag{6.1}$$

which yields the expressions

$$[\Psi_c^2]_{\Phi}(M_c) = \xi(M_c; \lambda_{r\gamma}^2) \quad \text{and} \tag{6.2}$$

$$\Psi_{\Phi,c}^2(M_c) = \xi(M_c; 1) \tag{6.3}$$

for these two statistics (since $\omega_r^{\text{PCS}} = \lambda_{r\gamma}^2$).¹ Since the former of these, under maximisation, is equivalent to χ_c^2 under minimisation, both χ_c^2 and $\Psi_{\Phi,c}^2$ can be related to instances of ξ through the equivalence relations

$$\chi_c^2(M_c)|_{\min} \equiv \xi(M_c; \lambda_{r\gamma}^2)|_{\max} \quad \text{and} \tag{6.4}$$

$$\Psi_{\Phi,c}^2(M_c)|_{\max} \equiv \xi(M_c; 1)|_{\max}. \quad \text{and} \tag{6.5}$$

An immediate consequence of this is that $\chi_c^2(M_c)|_{\min}$ and $\Psi_{\Phi,c}^2(M_c)|_{\max}$ are only equivalent if $\lambda_{r\gamma} = 1$, and are not equivalent otherwise; stated more succinctly,

$$\chi_c^2(M_c)|_{\min} \equiv \Psi_{\Phi,c}^2(M_c)|_{\max} \iff \lambda_{r\gamma} = 1. \tag{6.6}$$

Beyond this singular point of equivalence, the parameter estimates yielded by χ_c^2 converge to those yielded by $\Psi_{\Phi,c}^2$ as $\lambda_{r\gamma} \rightarrow 1$; as will be demonstrated in Chapter 8.

¹While ξ technically also takes D_c , E_c and B_c as parameters, they are omitted here in the interest of clarity.




Figure 6.1: A scatter plot of $\Psi_{\Phi,r}^2$ (radio-band GOF) vs. $\Psi_{\Phi,\gamma}^2$ (gamma-ray-band GOF) for PSR J2241–5236. Each grey dot represents an LC pair produced by an OG+Cone conglomerated model. The red and blue (vertical and horizontal) dashed lines identify the LC pairs M_c^r and M_c^γ that contain the single-band best-fit radio and gamma-ray LCs (as components), respectively. The yellow line (with slope -15.1) identifies the M_c^{pcs} , which is the best joint fit according to χ_c^2 ; while the (darker) green line (with slope -1) identifies M_c^{sfs} , the best joint fit according to $\Psi_{\Phi,c}^2$. Notice that since both $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$ are less than 1 for all possible fits (by construction), no LC pairs can lie in the grey region.

6.1.1 Understood Geometrically

In line with the geometric interpretation presented in Chapter 5, the relationship between χ_c^2 and $\Psi_{\Phi,c}^2$ can be understood geometrically by reframing best-fit identification as a geometric operation in $(\Psi_{\Phi,r}^2, \Psi_{\Phi,\gamma}^2)$ -space; leveraging the fact that each $M_c \in \mathcal{M}_c$ can be mapped onto a point in this space via its radio and γ -ray GOF components $\Psi_{\Phi,r}^2(M_{c,r})$ and $\Psi_{\Phi,\gamma}^2(M_{c,\gamma})$. Figure 6.1 shows an example of such a mapping, which results when the observed LCs for PSR J2241–5236 are fit using a dual-band model that is a conglomeration of an empirical, single-altitude Gaussian hollow-cone model (Story et al., 2007; Kijak & Gil, 1998, 2003, see Section 2.1.1) for the radio emission and a geometric outer gap model (Cheng et al., 1986b, see Section 2.1.2) for the γ -ray emission (the details of these models, and of their conglomeration, are not important here). The filled black circles at $(0, 0)$ and $(1, 1)$ correspond to the background and observed LC pairs B_c and D_c , since (by construction)

$$\Psi_{\Phi,r}^2(B_r) = \Psi_{\Phi,\gamma}^2(B_\gamma) = 0 \quad \text{and} \quad (6.7)$$

$$\Psi_{\Phi,r}^2(D_r) = \Psi_{\Phi,\gamma}^2(D_\gamma) = 1, \quad (6.8)$$

and each grey dot corresponds to a different $M_c \in \mathcal{M}_c$.

A useful consequence of this mapping is that any statistic that is equivalent to an instance of $\xi(M_c; \omega_r)$ (and therefore is a linear combination of $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$), is implicitly associated with a family of parallel straight lines in this space (specifically, with slope $-\omega_r$), with each line in the family corresponding to a different value for the statistic. The χ_c^2 statistic, therefore, is associated with the family of lines with slope $-\lambda_{r\gamma}^2$, and the $\Psi_{\Phi,c}^2$ statistic with the family of lines with slope -1 (by virtue of the equivalences in Eqs. (6.4) and (6.5)). Identifying the model LC pair that best fits D_c according to χ_c^2 , for example, then reduces to finding the line with slope $-\lambda_{r\gamma}^2$ that passes closest to D_c whilst also passing through at least one member of $\mathcal{M}_c$.

In this manner, the yellow dashed line in Figure 6.1 (with slope of -15.1 ; since $\Phi_r \simeq 105.1$ and $\Phi_\gamma \simeq 27.0$ for PSR J2241–5236), identifies the best-fit LC pair M_c^{pcs} associated with χ_c^2 (for which $\Psi_{\Phi,r}^2 = 0.95$ and $\Psi_{\Phi,\gamma}^2 = 0.16$), and the green dashed line (with slope -1) the best-fit LC pair M_c^{sfs} associated with $\Psi_{\Phi,c}^2$ (for which $\Psi_{\Phi,r}^2 = 0.87$ and $\Psi_{\Phi,\gamma}^2 = 0.54$).

Analogously, the radio-only and γ -ray-only statistics $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$ (and equivalently χ_r^2 and χ_γ^2) can be associated with the family of vertical and family of horizontal lines, respectively. Hence, the red (vertical) and blue (horizontal) dashed lines in Figure 6.1 identify the LC pairs M_c^r and M_c^γ , which contain the radio- and γ -ray-only best-fit LCs $M_{c,r}^r$ and $M_{c,\gamma}^\gamma$, respectively.

6.2 Characterising Dual-band Fits

This geometric interpretation also allows for a more robust characterisation of a dual-band fit in terms of uncertainty disparity and single-band best-fit non-colocation, the two factors that conspire to cause single-band dominance (as discussed in Section 3.2.2). Beyond this, it also allows for the evaluation of the magnitude of the effect these two factors have on a given fit, and, by extension, how sensitive a given statistic is to these factors.

6.2.1 The Causative Factors

The example mapping presented in Figure 6.1 demonstrates why radio dominance only occurs for pulsars where both an uncertainty disparity and single-band best-fit non-colocation is present. If, for instance, a single LC pair contained both the radio and γ -ray single-band-only best-fit LCs, i.e., if it were the case that $K_c^{\text{col}} \in \mathcal{M}_c$ and there was, therefore, no single-band best-fit non-colocation, then any ξ -equivalent statistic would identify that LC pair as the joint best fit, regardless of whether or not there were an uncertainty disparity. Similarly, if there were no significant uncertainty disparity, the slope of the lines associated with χ_c^2 would be $\simeq -1$ since it would be the case that $\lambda_{r\gamma} \simeq 1$ (and therefore $\omega_r^{\text{pcs}} \simeq 1$), and both the radio and γ -ray components of the resulting best-fit LC pair (which would lie closer to M_c^{sfs}) would have comparatively high GOF scores. It is only when both of these factors are present that radio dominance can occur.

This example mapping also demonstrates why radio dominance is more severe for fits where the non-colocation is greater. Stated most simply, the greater the separation between M_c^r and M_c^γ along the perimeter of $\mathcal{M}_c$ (i.e., the more LC pairs lie between M_c^r and M_c^γ on this perimeter), the greater the degeneracy in the identity of the dual-band model's best joint fit, and the greater the effect an uncertainty disparity can have. Conversely, if $K_c^{\text{col}} \in \mathcal{M}$ then the identity of the model's best joint fit is unambiguous, and any ξ -equivalent statistic would identify K_c^{col} as such regardless of the slope it corresponds to.

Uncertainty Disparity

A simple first option when trying to characterise the magnitude of the *uncertainty disparity* is to do so using $\lambda_{r\gamma}$. However, since either band can dominate, a more sensible characterisation is achieved by using

$$\log(\lambda_{r\gamma}) = \log(\Phi_r) - \log(\Phi_\gamma), \quad (6.9)$$

which encodes the severity of the uncertainty disparity in its magnitude (i.e., larger values for $|\log(\lambda_{r\gamma})|$ correlate with larger uncertainty disparities), and the sense of the disparity in its sign (i.e., if $\log(\lambda_{r\gamma}) > 0$, then the average relative radio uncertainty is likely smaller than the average relative γ -ray uncertainty, and vice versa).² For the example mapping shown in Figure 6.1, $\log(\lambda_{r\gamma}) \simeq 0.59$.

Single-band Best-fit Non-colocation

A simple, and natural way to characterise *single-band best-fit non-colocation* is to consider how much the respective single-band GOF statistics differ for M_c^r and M_c^γ . The amount by which $\Psi_{\Phi,r}^2$ differs for $M_{c,r}^r$ and $M_{c,r}^\gamma$ (the radio components of these LC pairs), is the *radio compromise* associated with the γ -ray-only best fit (i.e., with M_c^γ); this difference can be written

$$\begin{aligned} \triangle \Psi_{\Phi,r}^2(M_{c,r}^\gamma) &= \Psi_{\Phi,r}^2(M_{c,r}^r) - \Psi_{\Phi,r}^2(M_{c,r}^\gamma) \\ &= \Psi_{\Phi,r,\max}^2(\mathcal{M}_c) - \Psi_{\Phi,r}^2(M_{c,r}^\gamma), \end{aligned} \quad (6.10)$$

where the modified “ $\triangle$ ” symbol specifically denotes the difference relative to the maximum value for the statistic it precedes (for a given set of model LC pairs; in this case the maximum value of $\Psi_{\Phi,r}^2$ for $\mathcal{M}_c$). The *γ -ray compromise* associated with the radio-only best fit (i.e., with M_c^r) is, by analogy, the amount by which $\Psi_{\Phi,\gamma}^2$ differs for $M_{c,\gamma}^r$ and $M_{c,\gamma}^\gamma$ (the γ -ray components of M_c^r and M_c^γ), and can be written

$$\begin{aligned} \triangle \Psi_{\Phi,\gamma}^2(M_{c,\gamma}^r) &= \Psi_{\Phi,\gamma}^2(M_{c,\gamma}^\gamma) - \Psi_{\Phi,\gamma}^2(M_{c,\gamma}^r) \\ &= \Psi_{\Phi,\gamma,\max}^2(\mathcal{M}_c) - \Psi_{\Phi,\gamma}^2(M_{c,\gamma}^r). \end{aligned} \quad (6.11)$$

Given these two quantities, the *single-band best fit non-colocation* can be characterised using the product

$$\eta_s(M_c^r, M_c^\gamma) = \triangle \Psi_{\Phi,\gamma}^2(M_{c,\gamma}^r) \cdot \triangle \Psi_{\Phi,r}^2(M_{c,r}^\gamma). \quad (6.12)$$

Since $\triangle \Psi_{\Phi,r}^2 \geq 0$ and $\triangle \Psi_{\Phi,\gamma}^2 \geq 0$ for all $M_c \in \mathcal{M}_c$ (by construction), it follows that $\eta_s \geq 0$. The definition of η_s as the product of the two differences (as opposed to their sum) accounts for the possibility that there may be multiple best fits in either band. Geometrically, η_s is twice the surface area of the triangle defined by M_c^r , M_c^γ , and $K_c^{\text{col}} = (M_{c,r}^r, M_{c,\gamma}^\gamma)$, the LC pair formed by combining the radio- and γ -ray-only best-fit LCs (which corresponds to the ideal case, wherein the single-band best fits are colocated).

To put the value of η_s into context: $\eta_s = 0$ for the case where the single-band best fits are colocated, and $\eta_s = 1$ for the case where $M_{c,r}^r$ and $M_{c,\gamma}^\gamma$ are perfect fits to $D_{c,r}$ and $D_{c,\gamma}$, respectively (i.e., $\Psi_{\Phi,r,\max}^2 = \Psi_{\Phi,\gamma,\max}^2 = 1$), and their respective counterparts are null fits (i.e., $\Psi_{\Phi,\gamma}^2(M_{c,\gamma}^r) = \Psi_{\Phi,r}^2(M_{c,r}^\gamma) = 0$). Since the separation between M_c^r and M_c^γ

²Note that the logarithms in Eq. (6.9) are base 10.

Table 6.1: Characterising statistic values associated with the single- and dual-band best-fit LC pairs circled in Figure 6.1; for these fits: $\log(\lambda_{r\gamma}) \simeq 0.59$ (and therefore $\omega_{\text{F}}^{\text{PCS}} \simeq 15.1$) and $\eta_s \simeq 0.74$.

LC pair	$\Psi_{\Phi,r}^2$	$\Psi_{\Phi,\gamma}^2$	$\Psi_{\Phi,c}^2$	κ	Θ_r ($^\circ$)	f_r
M_c^r	0.96	-0.04	0.500	0.345	90.0	1.00
M_c^γ	0.00	0.73	0.365	0.480	0.0	-1.00
M_c^{PCS}	0.95	0.16	0.555	0.290	89.1	0.98
M_c^{sfs}	0.87	0.54	0.705	0.140	64.8	0.44
K_c^{col}	...	...	0.845	0.000	...	...

on the perimeter of $\mathcal{M}_c$ will tend to grow as η_s grows (and vice versa), η_s serves as a crude measure of the amount of degeneracy brought about by the non-colocation. For the example mapping in Figure 6.1, $\eta_s \simeq 0.74$.

An interesting property of η_s , as defined in Eq. (6.12), is that $\eta_s > 1$ for any fit where either $M_{c,\gamma}^r$ or $M_{c,r}^\gamma$ (the counterparts to the radio- and γ -ray-only best-fit LCs, respectively) is a bad fit to the relevant single-band component of D_c ; i.e., if $\eta_s > 1$, then either $\Psi_{\Phi,\gamma}^2(M_{c,\gamma}^r) < 0$ or $\Psi_{\Phi,r}^2(M_{c,r}^\gamma) < 0$.

It is important to reiterate that single-band best fit non-colocation does not cause radio dominance; it merely enables radio dominance by inducing a degeneracy in the identity of the model's joint best fit. The value of η_s , therefore, does not reflect the degree to which the radio GOF dominates. Instead, it reflects the degree to which the dual-band model achieves the potential implied by its component single-band models, which impacts the extent to which radio dominance can manifest.

6.2.2 The Magnitude of their Effect

Evaluating whether or not a given fit was significantly impacted by non-colocation and uncertainty disparity proves somewhat difficult, however, because the magnitude of the effect these two factors have depends on more than simply how severe they are. A natural first option might be to characterise the amount of priority the radio-band GOF enjoys, for a given statistic ξ , by comparing the GOF components associated with the best-fit LC pair M_c^ξ it identifies, directly; e.g., by comparing $\Psi_{\Phi,r}^2(M_{c,r}^{\text{PCS}})$ and $\Psi_{\Phi,\gamma}^2(M_{c,r}^{\text{PCS}})$ directly for a fit conducted using χ_c^2 . Such a characterisation proves ineffectual in general, however, since the component single-band models do not necessarily produce perfect, or even equally good, single-band fits. It may, for instance, be the case that the dual-band model produces only LC pairs for which the γ -ray GOF is greater than the radio GOF (i.e., that $\Psi_{\Phi,\gamma}^2(M_{c,\gamma}) > \Psi_{\Phi,r}^2(M_{c,r})$ for all $M_c \in \mathcal{M}_c$). For such model, a simple comparison of $\Psi_{\Phi,r}^2(M_{c,r}^{\text{PCS}})$ and $\Psi_{\Phi,\gamma}^2(M_{c,r}^{\text{PCS}})$ would invariably (and erroneously) conclude that the γ -ray GOF enjoys priority over the radio GOF, even for a radio-dominated dual-band statistic.

A better option is to, instead, consider where M_c^ξ lies relative to the ideal LC pair $K_c^{\text{col}} = (M_{c,r}^r, M_{c,\gamma}^\gamma) \notin \mathcal{M}_c$, made up of the single-band-only best-fit LCs, in $(\Psi_{\Phi,r}^2, \Psi_{\Phi,\gamma}^2)$ -space. In doing so it is possible to characterise a given joint fit, conducted using a statistic ξ , in terms of the apparent loss of joint GOF incurred as a result of the non-colocation (the *non-colocation penalty*) and the extent to which the radio-band GOF enjoys priority over the γ -ray-band GOF (the *radio priority*).

Non-colocation Penalty

The former of these can be comparatively easily quantified, by calculating the average loss in single-band-only GOF that is associated with the best-fit LC pair, since these losses are incurred due to the non-colocation of the single-band best-fit parameter point estimates. This yields the *non-colocation penalty*

$$\begin{aligned}\kappa(M_c^\xi) &= \frac{1}{2} \Delta \Psi_{\Phi,r}^2(M_{c,r}^\xi) + \frac{1}{2} \Delta \Psi_{\Phi,\gamma}^2(M_{c,\gamma}^\xi) \\ &= \Psi_{\Phi,c}^2(K_c^{\text{col}}) - \Psi_{\Phi,c}^2(M_c^\xi).\end{aligned}\tag{6.13}$$

Note that, while the definition of κ reduces to a difference of $\Psi_{\Phi,c}^2$ values, it is defined purely in terms of the single-band GOF statistics $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$. Interestingly, the $\Psi_{\Phi,c}^2$ statistic could, in principle, be derived quite simply from the $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$ statistics by finding the ξ -like dual-band statistic that minimises κ .

For M_c^{pcs} and M_c^{sfs} in Figure 6.1, this yields $\kappa_{\text{pcs}} = 0.290$ and $\kappa_{\text{sfs}} = 0.140$; as listed in Table 6.1.

Radio Priority

The latter of these can be quantified by comparing the radio and γ -ray compromises M_c^ξ necessitates, thereby gauging the extent to which the radio-band GOF is prioritised over the γ -ray-band GOF. Since either of these compromises can be 0 (as the radio compromise is when the radio band dominates), however, simply using their ratio cannot suffice, since it can sometimes be undefined. Fortunately, this can be easily circumvented by, instead, leveraging the angle

$$\Theta_r(M_c^\xi) = \begin{cases} \tan^{-1} \left(\frac{\Delta \Psi_{\Phi,\gamma}^2(M_{c,\gamma}^\xi)}{\Delta \Psi_{\Phi,r}^2(M_{c,r}^\xi)} \right), & \Delta \Psi_{\Phi,r}^2(M_{c,r}^\xi) > 0 \\ \frac{\pi}{2}, & \text{otherwise,} \end{cases}\tag{6.14}$$

between the horizontal and the line that passes through both M_c^ξ and K_c^{col} . With this angle in hand, the *radio priority* associated with a fit is

$$f_r(M_c^\xi) = \begin{cases} \frac{\Theta_r(M_c^\xi) - \pi/4}{\pi/4}, & \eta_s > 0 \\ 0, & \text{otherwise,} \end{cases}\tag{6.15}$$

which effectively maps $\Theta_r(M_c^\xi)$ onto the interval $(-1, 1)$. Therefore, $f_r(M_c^\xi) = 0$ indicates that neither band enjoys priority (this includes colocated fits), $f_r(M_c^\xi) > 0$ indicates that the radio band enjoys priority, and $f_r(M_c^\xi) < 0$ indicates that the γ -ray band enjoys priority.

For M_c^{pcs} and M_c^{sfs} in Figure 6.1, $\Theta_r = 89.1^\circ$ and 64.8° , and $f_r = 0.98$ and 0.44 ; as listed in Table 6.1.

6.3 Characterising Dual-band Statistics

As mentioned in Section 6.2.2, the magnitude of the effect that uncertainty disparity and non-colocation has on any given fit depends on more than just the magnitude of these two factors. It may, for instance, be the case that $f_r \neq 0$ for every LC pair on the perimeter of $\mathcal{M}_c$, or, that $f_r \neq 0$ for the model's true, undominated best fit. Therefore, a high radio priority for a single fit cannot, by itself, establish that the statistic used is radio dominated. Instead, the behaviour of a given statistic across a number of fits should be considered.

Single-band Dominance A statistic for which $f_r \simeq 1$ for the fits to all the pulsars in a population $\mathcal{D}_c$ (comprising N pulsars, with associated backgrounds $\mathcal{B}_c$), can be said to be radio dominated for that population. Similarly, a statistic for which $f_r \simeq -1$ across such a population can be said to be γ -ray-dominated for that population. Collectively, such statistics can be characterised as being single-band dominated, and are easily contrasted with undominated statistics, for which the expected value of f_r is 0 (for a given population of pulsars).

Figure 6.2 (a) shows an example of what the radio priority distribution (as a function of $\lambda_{r\gamma}$) looks like for each of these types of dual-band statistics; the red dots correspond to the results for a radio-dominated statistic, the blue dots




Figure 6.2: Example plots of f_r as function of $\log(\lambda_{r\gamma})$ for four types of dual-band statistic. The red plot in panel (a) is for a radio-dominated statistic, with the fits all clustering close to $f_r = 1$. Conversely, the blue plot (with the fits clustering close to $f_r = -1$) is for a γ -ray-dominated statistic. The green plot, with the fits scattered (somewhat broadly) around $f_r = 0$, is for an undominated statistic. The yellow plot, in panel (b), is for a statistic that is prone to precision-determined dominance. It is, essentially, a mixture of the three plots in panel (a): γ -ray-dominated to the left, undominated in the centre (around $\log(\lambda_{r\gamma}) = 0$), and radio-dominated to the right.

correspond to the results for a γ -ray-dominated statistic and the green dots correspond to the results for an undominated statistic. Notice that these three example statistics are not sensitive to $\log(\lambda_{r\gamma})$, in contrast to χ_c^2 .

6.3.1 Precision-determined Priority

As demonstrated in Section 4.3, the relative weight of the radio component of $[\Psi_c^2]_\Phi$, the direct, normalised equivalent to χ_c^2 , is determined by the value of $\lambda_{r\gamma}$. This means that for χ_c^2 one might expect a radio priority distribution more similar to that shown in Figure 6.2 (b), which resembles a mixture of the three profiles in Figure 6.2 (a). In the left- and right-most regions (delimited by the dashed vertical lines), χ_c^2 resembles a γ -ray-dominated and radio-dominated statistic, respectively, while in the central region it resembles an undominated statistic (albeit with a very wide range of possible f_r values). As discussed in Section 3.2, however, χ_c^2 is typically either apparently undominated or apparently radio-dominated, since for most pulsars $\log(\lambda_{r\gamma}) \geq 0$.

Taking this apparent $\lambda_{r\gamma}$ -dependent mixing into account, a good statistic by which to characterise the *precision-determined dominance* observed for χ_c^2 , is the weighted sum of the radio priorities across $\mathcal{D}_c$; i.e., via

$$F(\mathcal{D}_c, \mathcal{B}_c, \mathcal{M}_c; \xi) = \frac{\sum_{k=1}^N \log(\lambda_{r\gamma,k}) \cdot \tilde{\eta}_{s,k} \cdot f_{r,k}}{\sum_{k=1}^N |\log(\lambda_{r\gamma,k})| \cdot \tilde{\eta}_{s,k}}, \quad (6.16)$$

where, for every $D_{c,k} \in \mathcal{D}_c$ (i.e., for the k th pulsar in the population), $\log(\lambda_{r\gamma,k})$ characterises the uncertainty disparity, $\tilde{\eta}_{s,k} = \min(1, \eta_{s,k})$ characterises the single-band best-fit non-colocation (constrained to be ≤ 1 to avoid outliers), and $f_{r,k}$ characterises the resulting radio priority. The weighting in the sum in the numerator of F (of the radio priority associated with each $D_{c,k}$), is intended to render it sensitive to the degree to which the distribution of the fits as a function of $\log(\lambda_{r\gamma,k})$ and $f_{r,k}$ resembles an “S” shape, similar to that of the distribution plotted in Figure 6.2(b).

Of specific note is the interaction between the $\log(\lambda_{r\gamma,k})$ and $f_{r,k}$ factors in this weighting. If the precision-determined dominance is such that the *more precise* band tends to enjoy priority (as is the case for χ_c^2), then these two factors will tend to carry the same sign, and the terms of the sum in the numerator will tend to be positive. This, in turn, means that the magnitude of F will tend to be closer to the maximum value allowed by the magnitudes of the terms of the sum in

the numerator (ignoring their signs), since comparatively little of this potential magnitude will be lost to the effective cancellation of terms; i.e., for such a distribution, $F \simeq 1$. Similarly, if the dominance is such that the *less precise* band tends to enjoy priority, then these two factors will tend to carry opposite signs, and the terms of the sum will tend to be negative. This, however, still means that comparatively little of F 's potential magnitude is lost to cancellation; all that changes for such a distribution is the sign of F (being negative for this latter case). Therefore, for such a distribution, $F \simeq -1$.

For radio priority distributions associated with undominated or purely radio- or γ -ray-dominated statistics, like the ones shown in Figure 6.2 (a), $F \simeq 0$ since the terms for which $\log(\lambda_{r\gamma}) > 0$ (i.e., associated with dots to the right of the dotted vertical line) will tend to roughly cancel out with the terms for which $\log(\lambda_{r\gamma}) < 0$ (associated with the dots to the left of the dotted vertical line). Interestingly, F can only effectively discriminate between pure single-band dominance and precision-determined dominance if the marginal distribution of the LCs in $\mathcal{D}$ with respect to $\log(\lambda_{r\gamma})$ is roughly symmetrical about $\log(\lambda_{r\gamma}) = 0$. If, for instance, $\mathcal{D}$ consists mainly of pulsar for which $\log(\lambda_{r\gamma}) > 0$, then F would be $\simeq 1$ both when a radio-dominated statistic is used and when a statistic is used that suffers from precision-determined dominance, since no cancellation can occur in the absence of pulsars for which $\log(\lambda_{r\gamma}) < 0$.

6.4 The Convergence of χ_c^2 and $\Psi_{\Phi,c}^2$

As stated in Section 3.2.2, beyond their equivalence (with regard to parameter estimation) when $\lambda_{r\gamma} = 1$, the χ_c^2 and $\Psi_{\Phi,c}^2$ statistics also tend to yield more similar parameter estimates as $\lambda_{r\gamma} \rightarrow 1$. Since it is possible, however, that very similar model LCs can correspond to parameter combinations that are widely separated in the model's parameter space, characterising this convergence by comparing the best-fit parameter combinations themselves is likely generally ineffectual.

Fortunately, since a disagreement between χ_c^2 and $\Psi_{\Phi,c}^2$ simply means that M_c^{pcs} and M_c^{sfis} are non-colocated, the definition of η_s (which characterises the non-colocation of M_c^{r} and M_c^{y}) provides an ideal blueprint for the definition of a statistic that can serve this purpose: the *dual-band best-fit non-colocation*

$$\eta_c(M_c^{\text{pcs}}, M_c^{\text{sfis}}) = \underline{\Delta} \Psi_{\Phi,c}^2(M_c^{\text{pcs}}) \cdot \underline{\Delta} [\Psi_c^2]_{\Phi}(M_c^{\text{sfis}}), \quad (6.17)$$

where

$$\begin{aligned} \underline{\Delta} \Psi_{\Phi,c}^2(M_c^{\text{pcs}}) &= \Psi_{\Phi,c}^2(M_c^{\text{sfis}}) - \Psi_{\Phi,c}^2(M_c^{\text{pcs}}) \\ &= \Psi_{\Phi,c,\text{max}}^2(\mathcal{M}_c) - \Psi_{\Phi,c}^2(M_c^{\text{pcs}}), \end{aligned} \quad (6.18)$$

is the apparent joint GOF compromise associated M_c^{pcs} , as determined using $\Psi_{\Phi,c}^2$, and

$$\begin{aligned} \underline{\Delta} [\Psi_c^2]_{\Phi}(M_c^{\text{sfis}}) &= [\Psi_c^2]_{\Phi}(M_c^{\text{pcs}}) - [\Psi_c^2]_{\Phi}(M_c^{\text{sfis}}) \\ &= [\Psi_c^2]_{\Phi,\text{max}}(M_c) - [\Psi_c^2]_{\Phi}(M_c^{\text{sfis}}) \end{aligned} \quad (6.19)$$

is the apparent joint GOF compromise associated M_c^{sfis} , as determined using $[\Psi_c^2]_{\Phi}$. Stated in terms of η_c , then, Chapter 8 will demonstrate that, in general, parameter estimates derived using χ_c^2 converge on those derived using $\Psi_{\Phi,c}^2$ as $\lambda_{r\gamma} \rightarrow 1$ since $\eta_c \rightarrow 0$ as $\lambda_{r\gamma} \rightarrow 1$ (or, equivalently, as $\log(\lambda_{r\gamma}) \rightarrow 0$).

Part III

IMPLEMENTATION

Chapter 7

Dual-band Fitting: Conceptual Design

Before the newly-developed $\Psi_{\Phi,c}^2$ statistic can be leveraged to derive pulsar parameter estimates from real-world pulsar data (in Chapter 10; using the geometric models introduced in Chapter 2), a few preparatory steps must be taken. The first of these steps is to prepare the observed radio and γ -ray LCs (from the 2PC) for dual-band fitting using χ_c^2 and $\Psi_{\Phi,c}^2$. The second step is to prepare the raw model LCs that the geometric models produce, to render them compatible with the observed LCs under χ_c^2 and $\Psi_{\Phi,c}^2$ (i.e., to convert the raw model LCs into proper LCs that take into account the measurement units, background level and phase definition used for the 2PC LCs). The last step is to construct a Monte Carlo simulation algorithm by which constraints on the point estimates yielded by χ_c^2 and $\Psi_{\Phi,c}^2$ can be derived.

7.1 The Observed LC Pairs

As described in Section 3.1, an observed LC consists of two parts (for the purposes of LC fitting): a sequence of intensities (that covers one rotation of the pulsar) and a sequence of uncertainties on those intensities. Unfortunately, only one of these components is present for the observed radio LCs in the datasets we use; specifically, the observed radio LCs are not furnished with an accompanying sequence of uncertainties. To remedy this, we derive an estimate for the measurement uncertainty associated with each of these LCs by calculating the standard deviation of the relative intensity in the off-peak region. The uncertainty on the intensity in each radio LC bin (i.e., the uniform radio uncertainty ε_r) is then assumed to be equal to the estimate thusly derived.

For the 23 pulsars in Chapter 10, the relevant off-peak regions were identified by eye (and are indicated on the LC plots for those pulsars in Appendix A), and the estimates they yield are listed in Table 10.1.

7.2 Extracting LCs from Geometric Models

As shown in Figure 2.1 (for both a geometric radio and geometric γ -ray model), each constant- ζ cut through a model-produced phaseplot yields a raw, normalised model LC (with n phase bins; in the notation introduced in Chapter 5)

$$C(\theta) = (c_i)_n, \quad (7.1)$$

where $c_i \in \mathbb{R}$ is the relative LC intensity in the i th LC bin (i.e., relative to the peak intensity $C_{\max} := 1$; note that we assume that C is not flat) and $\theta \in \mathcal{P}$ (the model's parameter space) is the parameter combination associated with that cut. To translate C into a prediction for the true intensity of the pulsar-associated emission (i.e., a proper LC), we introduce an additional model parameter: the relative amplitude $a \in \mathbb{R}^+$ of the model-predicted pulsar-associated emission (relative to that of the observed pulsar-associated emission). In essence, this parameter governs the ratio between the predicted above-background emission (in the eventual model-predicted LC) and the observed above-background emission.

In terms of this parameter (alongside the background level bu ; where u is the relevant measurement unit), the general form for the LC a geometric model predicts, is

$$M(a, C; D, B) = ([a \cdot c_i \cdot (d_{\max} - b)]u + bu)_n \in \mathbb{K}^n, \quad (7.2)$$

where $D_{\max} = d_{\max}u$ is the peak intensity in D . Upon recognising that $(d_{\max} - b)u$ is, in fact, the magnitude of largest component of the perturbation vector $\mathbf{p} = \mathbf{d} - \mathbf{b}$ (see Eq. (5.23)), this expression can be rewritten as

$$M = ([a \cdot c_i \cdot p_{\max}]u + bu)_n, \quad (7.3)$$

where $p_{\max} = d_{\max} - b$ and $B = (bu)_n$ is the background LC constituted by the flat background level. Notice that, with $m_i^* = a \cdot c_i \cdot p_{\max}$, the quantity $[a \cdot c_i \cdot p_{\max}]u$ is, in fact, the magnitude of the i th component of the model perturbation vector $\mathbf{m}^*$ that connects B and M (as depicted in Figure 5.3). This identification makes it quite clear that the geometric models are truly aimed at reproducing the observed perturbation $\mathbf{p}$.

As an example of what the value of a means: For an observed LC for which $D_{\max} = 80u$ and $b = 50$ (i.e., with background level $50u$ and $p_{\max} = 80 - 50 = 30$), an amplitude of $a = 1$ would result in a model LC for which $M_{\max} = 80u$ (where $M_{\max}$ is the peak intensity in M), while an amplitude of $a = 0.5$ would result in a model LC for which $M_{\max} = [0.5 \cdot 1 \cdot 30]u + 50u = 65u$ (recall that $C_{\max} := 1$). In both cases the model LC intensities would uniformly be greater than the background level; i.e., in both cases $M_i \geq 50u$ for all $i \in [1, n]$.

7.2.1 Finding the Optimal Amplitude

A convenient side effect of this formulation is that, given a raw LC C and an observed LC D (with LC uncertainties E), the amplitude a^* that results in the greatest GOF for the eventual model-predicted LC (i.e., the optimal amplitude), can be found analytically. Since all of the single-band statistics defined in Chapters 4 and 5 (i.e., the χ^2 and Ψ_{Φ}^2 statistics) are equivalent with regard to best-fit identification, this can be accomplished by taking the partial derivative (with respect to a) of the expression for $\chi^2(M)$ that arises when the model LC is assumed to have the form specified in Eq. (7.2), setting it equal to zero and solving for a ; i.e., by solving for a in

$$\begin{aligned} \frac{\partial}{\partial a} \chi^2(M; D, E) &= \frac{\partial}{\partial a} \sum_{i=1}^n \left[\frac{(a \cdot c_i \cdot p_{\max} + b) - d_i}{\varepsilon_i} \right]^2 \\ &= \frac{\partial}{\partial a} \sum_{i=1}^n \left[\frac{a \cdot c_i \cdot p_{\max}}{\varepsilon_i} - \frac{d_i - b}{\varepsilon_i} \right]^2 \\ &= \frac{\partial}{\partial a} \left[a^2 \sum_{i=1}^n \left(c_i \cdot \frac{p_{\max}}{\varepsilon_i} \right)^2 - 2a \sum_{i=1}^n \left(c_i \cdot \frac{p_{\max}}{\varepsilon_i} \cdot \frac{p_i}{\varepsilon_i} \right) + \sum_{i=1}^n \left(\frac{p_i}{\varepsilon_i} \right)^2 \right] \\ &= 2a \sum_{i=1}^n \left(c_i \cdot \frac{p_{\max}}{\varepsilon_i} \right)^2 - 2 \sum_{i=1}^n \left(c_i \cdot \frac{p_{\max}}{\varepsilon_i} \cdot \frac{p_i}{\varepsilon_i} \right) \\ &= 0, \end{aligned} \quad (7.4)$$

where $E_i = \varepsilon_i u$ is the uncertainty in the i th LC bin and $p_i = p_i u$ is the component of $\mathbf{p}$ associated with that bin. This yields

$$a^*(C; D, E, B) = \frac{1}{p_{\max}} \cdot \sum_{i=1}^n \left(\frac{c_i}{\varepsilon_i} \cdot \frac{p_i}{\varepsilon_i} \right) / \sum_{i=1}^n \left(\frac{c_i}{\varepsilon_i} \right)^2, \quad (7.5)$$

which minimises $\chi^2(M; D, E)$ since

$$\frac{\partial^2}{\partial a^2} \chi^2(M; D, E) = 2 \sum_{i=1}^n \left(c_i \cdot \frac{p_{\max}}{\varepsilon_i} \right)^2 \geq 0. \quad (7.6)$$

For observed LCs with uniform LC uncertainties (i.e., that have the same uncertainty in each LC bin), as is typically

assumed for observed radio LCs, this expression for the optimal amplitude reduces to

$$a^*(C; D, B) = \frac{1}{p_{\max}} \cdot \frac{\sum_{i=1}^n c_i \cdot p_i}{\sum_{i=1}^n c_i^2}, \quad (7.7)$$

which is independent of the LC uncertainties.

Since the single-band models used for the fits in this study are geometric in nature, the optimal amplitudes a_r^* and a_γ^* associated with the radio and γ -ray components of the LC pairs they produce (for a given observed LC pair D_c), are independent of each other. Therefore, each of the conglomerated dual-band models effectively take two additional parameters: the radio and γ -ray amplitudes a_r and a_γ .

7.2.2 Accounting for Phase Ambiguity

The inherent degeneracy brought about by the different conventions for the definition of rotational phase used in the phase-matched *Fermi* LAT observations (i.e., in the 2PC) and in the geometric models, also needs to be taken into account. The former of these derives its definition (of the observed rotational phase ϕ) from the shape of the radio LC, with $\phi = 0$ at the radio LC peak for single-peak profiles and roughly at the centre of the main radio pulse for more complicated profiles. The latter, on the other hand, defines a model-specific rotational phase ϕ_M , such that $\phi_M = 0$ at the point in the pulsar's rotation where its magnetic field axis passes between its rotation axis and the observer. In the interest of simplicity (and consistency with the available literature; e.g., [Johnson et al., 2014](#)), the convention used in the observations takes precedence throughout this document (hence its greater notational primacy; unless otherwise specified).

Unfortunately, it is (in general) unknown *a priori* at which point in observed rotational phase (i.e., at which value for ϕ) the pulsar's magnetic field axis passes between its rotation axis and the observer. Therefore, this ambiguity cannot simply be accounted for by shifting the model LCs in phase by a set (likely pulsar-specific) amount. Instead, we introduce an additional free parameter to each of the geometric single-band models: the point ϕ_μ in observed rotational phase where this passage occurs. In effect, ϕ_μ (which is cyclical) is the temporal equivalent of α and ζ , in that it specifies the pulsar's orientation in rotational phase. The best-fit value for ϕ_μ , therefore, constitutes an estimate of the point in rotational phase at which the pulsar's magnetic field axis passes between its rotation axis and the observer. This means that the fits we conduct also constrain the value of ϕ_μ for each pulsar.

Through this additional parameter, the two definitions for rotational phase can be related via the identity

$$\phi_M = \phi - \phi_\mu. \quad (7.8)$$

Therefore, when conducting a fit, each value for ϕ_μ (which we vary from 0 to 1 in steps equivalent to two radio LC bins) corresponds to a model LC that is shifted in phase by ϕ_μ relative to the model LC for which $\phi_\mu = 0$ (since, by the above identity, $\phi = \phi_M + \phi_\mu$).

When conducting dual-band fits, this parameter (just like α and ζ) is shared between the component single-band models, because allowing it to take different values for the two would effectively assume a different magnetic field for each (i.e., the location of the pulsar's magnetic field axis in rotational phase would not be the same for both). In the dual-band case (for a given joint fit), then, each value for ϕ_μ corresponds to a pair of shifted LCs, each shifted by the same amount relative to its counterpart LC in the LC pair for which $\phi_\mu = 0$.

Implications for the Models' Parameter Space

The introduction of ϕ_μ allows the symmetry of the retarded vacuum dipole field to be leveraged to reduce the range of values necessary for α and ζ when conducting a fit; from $\alpha, \zeta \in [0^\circ, 180^\circ]$ to $\alpha, \zeta \in [0^\circ, 90^\circ]$. This reduction is possible because, under the symmetry inherent to this field, the LC associated with a parameter combination $(\alpha, \zeta, \phi_\mu)$ is identical to the LCs associated with the parameter combinations $(\alpha, 180^\circ - \zeta, \phi_\mu + 0.5)$, $(180^\circ - \alpha, \zeta, \phi_\mu + 0.5)$, and $(180^\circ - \alpha, 180^\circ - \zeta, \phi_\mu)$. Notice that, by the cyclical nature of rotational phase, $\phi_\mu + 0.5 = \phi_\mu - 0.5$; i.e., even if $\phi_\mu > 0.5$, it is still the case that $\phi_\mu + 0.5 \in [0, 1]$.




Figure 7.1: Phaseplots for the TPC γ -ray model for $\alpha = 54^\circ$ and $\alpha = 180^\circ - 54^\circ = 126^\circ$. All four constant- ζ cuts produce the same LC due to symmetry.

Figure 7.1 demonstrates this symmetry for a parameter combination of $(\alpha = 54^\circ, \zeta = 45^\circ, \phi_\mu = 0)$ using the appropriate phaseplots for the TPC γ -ray model; i.e., using the phaseplots for which $\alpha = 54^\circ$ and $\alpha = 180^\circ - 54^\circ = 126^\circ$, with $\zeta \in [0^\circ, 180^\circ]$. All four constant- ζ cuts (which each correspond to one of the four equivalent cases noted above) produce the same LC: the example γ -ray LC plotted in Figure 2.1. The green, blue and red cuts are, in effect, the symmetric counterparts to the yellow cut (which is the same cut as in Figure 2.1).

7.3 The Sets of Model LC Pairs

Upon taking the steps described above, each LC pair that a dual-band model produces (for a given pulsar fit) is effectively associated with a parameter combination of the form

$$\vartheta = (\alpha, \zeta, \phi_\mu, a_r, a_\gamma). \quad (7.9)$$

(Note that the same α and ζ value is effectively used for both wavebands.) To construct the set of model LC pairs involved in a given dual-band fit, after finding the optimal radio and γ -ray amplitudes as described above, we vary the remaining three free parameters as follows (taking into account the symmetries discussed in Section 7.2.2): α and ζ are each (independently) varied from 0° to 90° at a 1° resolution, and ϕ_μ is varied at a resolution equal to two radio bin widths (which may differ between pulsars; as stated in Section 7.2.2). Given these resolutions, each set of model LC pairs (i.e., both the set for the TPC+Cone dual-band model and the set for the OG+Cone dual-band model) contains at most $(90 \times 90 \times n_r/2) = (4050 \times n_r)$ distinct LC pairs (assuming that n_r is even, as is fortunately the case for all the observed radio LCs that appear in this study). Since a dual-band pulsar observation expressly implies that there is pulsar-associated emission present in both bands, only model LC pairs where both component LCs contained some pulsar-associated emission (i.e., where the pulsar was visible in both bands) were included in the final sets of LC pairs.

As an additional precaution against spurious best-fit identifications (which primarily occurred for fits conducted using χ^2_c during the development of this approach), the range of a_γ was confined such that $|D_{\gamma, \max} - M_{\gamma, \max}|$ is no greater than $1.5E_{\gamma, \max}$ (where $E_{\gamma, \max}$ is the uncertainty on $D_{\gamma, \max}$).

7.4 Deriving Parameter Constraints

As discussed in Section 3.2.1, the procedure by which constraints on the pulsar parameters may be derived (by leveraging the properties of the χ^2 statistic; described in Section 3.1.1) can only, strictly speaking, yield parameter constraints for valid models, since it presupposes that the model in question passes the χ^2 GOF test (i.e., is a valid model). This means that, in principle, an alternative approach is necessary for fits that are merely favourable (i.e., that are favourable but not good). A similar situation exists for fits conducted using the $\Psi_{\Phi,c}^2$ statistic, since it is, in effect, a weighted sum of two χ^2 statistics, and is therefore (in general) not χ^2 -distributed. This means that an alternative approach is necessary for both merely favourable fits conducted using χ_c^2 and fits conducted using $\Psi_{\Phi,c}^2$.

Fortunately, the procedure described in Section 3.1.1 provides an ideal template for a more generalised procedure. Specifically, if the distributions of the appropriate analogues to $\Delta\chi^2(M_{\text{true}})$ (and $\Delta\chi_c^2(M_{c,\text{true}})$) could somehow be determined, or at least approximated, they could be leveraged to derive the desired parameter constraints just as the distribution of $\Delta\chi^2(M_{\text{true}})$ is leveraged in the ideal case. Interestingly, as alluded to in Section 3.2.1, the analogues for merely favourable fits conducted using χ^2 and χ_c^2 turn out to also be χ^2 -distributed with n_θ and n_θ degrees of freedom, respectively.

7.4.1 A Generalised Procedure

To effectively leverage the template provided by the procedure described in Section 3.1.1, however, a somewhat more nuanced understanding of why parameter constraints are necessary in the first place, is required. Stated most simply (as alluded to in Section 3.1), they are necessary because D is not the true pulsar LC D_{true} . This means that the χ^2 statistic, which takes D as a parameter in order to calculate DFO, is merely an estimator for the true deviation, with respect to D_{true} , that the model LC in question corresponds to. More formally: the statistic $\chi^2(M; D, E)$ is an estimator for the true deviation

$$\chi_{\text{true}}^2(M) = \chi^2(M; D_{\text{true}}, E) \quad (7.10)$$

associated with a given model LC M (under the assumption that the LC uncertainties that comprise E accurately capture the variance associated with D_{true}). Understood in these terms, the hypothesis that the χ^2 GOF test aims to test is that “there exists an LC $M_{\text{true}} \in \mathcal{M}$ for which $\chi_{\text{true}}^2(M_{\text{true}}) = 0$,” and it leverages the distribution of an estimator for $\chi_{\text{true}}^2(M_{\text{true}})$ (i.e., that of $\chi^2(M_{\text{true}})$) to do so (compare the discussion in Section 3.1). For the valid-model case, then, parameter constraints are necessary because it is uncertain which LC in $\mathcal{M}$ is M_{true} since D_{true} is unknown (i.e., which LC in $\mathcal{M}$ is identical to D_{true}). The procedure by which these constraints are derived (for valid models) leverages the distribution of $\Delta\chi^2(M_{\text{true}})$, which is an estimator for $\Delta\chi_{\text{true}}^2(M_{\text{true}})$ (which is always 0 for a valid model).

More generally, parameter constraints are necessary because it is uncertain which LC

$$M_{\text{prox}} = M(\theta_{\text{prox}}) \in \mathcal{M}$$

is closest to D_{true} since D_{true} is unknown; i.e., for which model LC the value of $\chi_{\text{true}}^2(M)$ is a minimum.¹ In the valid model case, the closest model LC to D_{true} is M_{true} , since $\chi_{\text{true}}^2(M_{\text{true}}) = 0$; i.e., in the valid-model case $M_{\text{prox}} = M_{\text{true}}$. Since M_{prox} is, in essence, the generalised analogue of M_{true} , constraints on the true best-fit parameter combination θ_{prox} can, in principle, be derived in the invalid-model case by leveraging the distribution that is the generalised analogue of that of $\Delta\chi^2(M_{\text{true}})$: the distribution of $\Delta\chi^2(M_{\text{prox}})$. Unfortunately, since M_{prox} is both unknown (since D_{true} is unknown) and not identical to D_{true} (since the model failed the χ^2 GOF test), the exact distribution of this quantity is unknown (compared to the valid-model case, where the exact distribution of $\Delta\chi^2(M_{\text{true}})$ is known).

Fortunately, a closely related distribution is available (in principle): that of $\Delta\chi^2(M^{\text{PCS}})$. Since the expected value of D is D_{true} (under the assumption that D is a stochastic variable taken from a multivariate normal distribution with mean D_{true} and standard deviations E ; the core assumption of the χ^2 GOF test), the expected value of M^{PCS} is M_{prox} ; i.e., since it is most likely that D is identical to D_{true} , it is most likely that M^{PCS} is identical to M_{prox} . Therefore, to derive parameter constraints for fits where the model is not valid, it may be sufficient to construct the distribution of $\Delta\chi^2(M^{\text{PCS}})$, and use

¹Here we assume that the model is linear in its parameters, and that only one model LC is therefore closest to D_{true} .




Figure 7.2: Mappings of $\mathbb{K}^n$ onto $\mathbb{R}^n$ under the operator $[\cdot]_{\mathcal{E}}$ (which results in the χ^2 statistic) for two fits to an observed LC D with uncertainties E , one using a valid model $F(\theta)$ (left) and one using an invalid model $G(\theta)$ (right). As discussed in Chapter 5, E gives rise to the basis $\mathcal{E}$ for $\mathcal{L}$ (the vector space of translations associated with $\mathbb{K}^n$), which maps onto the unit hypersphere centred on D_{true} in $\mathbb{R}^n$ in these mappings (the black dashed circles). All the LCs that map onto the outer hyperspheres (the grey dashed circles) are a distance $\sqrt{n}$ away from D_{true} ; the expected value of $\chi(F_{\text{true}})$. Minimising χ^2 over the model's parameters effectively reduces the dimensionality of the multivariate distribution, by n_{θ} , which means that the analogous hypersphere for M^{pcs} has radius $\sqrt{\nu}$. The two grey diamonds (in each panel; that lie on the outer grey circles) indicate one such hypersphere. As discussed, $\chi(F^{\text{pcs}})$ and $\chi(G^{\text{pcs}})$ are proxies for the true deviations $\chi_{\text{true}}(F_{\text{true}})$ and $\chi_{\text{true}}(G_{\text{prox}})$; used because the D_{true} is unknown.

it as a proxy for that of $\Delta\chi^2(M_{\text{prox}})$. In lieu of a better alternative, then, we opt to derive parameter constraints for fits that fail the χ^2 GOF test by leveraging the distribution of $\Delta\chi^2(M^{\text{pcs}})$. Fortunately, since M^{pcs} has direct analogues for χ_c^2 and $\Psi_{\Phi,c}^2$, the approach is easily extended for use with these two statistics; specifically, for fits using the former we leverage the distribution of $\Delta\chi_c^2(M_c^{\text{pcs}})$ and for the latter the distribution of $\Delta\Psi_{\Phi,c}^2(M_c^{\text{sfbs}})$ (both constructed numerically; see Section 7.4.2 for a more detailed discussion).

In summary, the key difference between the two procedures for deriving parameter constraints is that for the valid-model procedure, parameter constraints are necessary because it is uncertain which LC in $\mathcal{M}$ is *identical* to D_{true} , while for the more generalised procedure, they are necessary because it is uncertain which LC in $\mathcal{M}$ is *closest* to D_{true} .

Figure 7.2 demonstrates this within the vector space framework (introduced in Chapter 5), for two fits to the same observed LC. One fit (the one depicted in the left-hand pane) uses a valid model $F(\theta)$ (which we assume, for argument's sake, does, in fact, produce an LC F_{true} that is identical to the true LC D_{true}), while the other fit (depicted in the right-hand pane) uses an invalid model $G(\theta)$. The χ^2 statistic identifies F^{pcs} as the former model's best fit, and G^{pcs} as the latter model's best fit. These LCs are not, however, the model LCs that are closest to D_{true} ; the LCs F_{true} and G_{prox} are the model LCs that are closest to D_{true} . Since $\chi^2(F^{\text{pcs}}) \simeq \nu = n - n_{\theta}$, the former model would pass the χ^2 GOF test, and the hypothesis that it produces an LC F_{true} that is identical to D_{true} cannot confidently be rejected. Note that, in this case, χ_{true}^2 would (if it were treated as a GOF statistic) identify F_{true} as the LC in $F(\theta)$ that is closest to D_{true} (since it is identical to D_{true}); hence, in this case F_{true} is F_{prox} . The latter model, on the other hand, would likely not pass the χ^2 GOF test, since $\chi^2(G^{\text{pcs}})$ is significantly larger than ν .

Upon rejection of the hypothesis that $\chi_{\text{true}}^2(G_{\text{prox}}) = 0$ (i.e., upon establishing that $G(\theta)$ is an invalid model), the question becomes ‘‘How big is the true deviation associated with this model's true best fit?’’ or ‘‘How big is $\chi_{\text{true}}^2(G_{\text{prox}})$?’’ Since D_{true} is unknown, however, both this quantity and the most immediate estimator for it, $\chi^2(G_{\text{prox}})$, cannot be calculated. The most immediate estimator for $\chi_{\text{true}}^2(G_{\text{prox}})$ that is available, is $\chi^2(G^{\text{pcs}})$.

A Geometric Proof that $\Delta\chi^2(M_{\text{prox}}) \sim \chi_{k=n_{\theta}}^2$

The diagram in the right-hand panel in Figure 7.2 allows for a quite simple geometric proof that $\Delta\chi^2(M_{\text{prox}}) \sim \chi_{k=n_{\theta}}^2$, and hence that the procedure described in Section 3.1.1 remains essentially unchanged for merely favourable fits. Since the distance between D and G_{prox} is $\chi(G_{\text{prox}})$ and the distance between D and G^{pcs} is $\chi(G^{\text{pcs}})$, the distance between G_{prox}

and G^{PCS} is (by Pythagoras's Theorem)

$$\sqrt{\chi^2(G_{\text{prox}}) - \chi^2(G^{\text{PCS}})} = \sqrt{\Delta\chi^2(G_{\text{prox}})}.$$

This distance is, however, the same as the perpendicular distance between D and the line that runs through G_{prox} and D_{true} (of which only a segment is shown). Therefore, since this perpendicular distance is χ^2 -distributed with n_θ degrees of freedom, $\Delta\chi^2(M_{\text{prox}}) \sim \chi_{k=n_\theta}^2$.

7.4.2 The Monte Carlo Algorithm

In analogue to the procedure outlined by Avni (1976), the Monte Carlo algorithm we implement aims to numerically construct the distributions of $\Delta\Psi_{\Phi,c}^2(M_c^{\text{sfs}})$ and $\Delta\chi_c^2(M_c^{\text{PCS}})$, and estimate the critical values for these quantities that correspond to the 2σ confidence levels for these distributions: $\delta_{2\sigma,\text{sfs}}$ and $\delta_{2\sigma,\text{PCS}}$. A confidence level of 2σ was chosen based on the requirements for the first real-world test of the effectiveness of $\Psi_{\Phi,c}^2$, conducted in Chapter 10. In the interest of clarity, the description that follows focuses on the procedure for χ_c^2 , with a short discussion of the details involved in its extension to $\Psi_{\Phi,c}^2$ afterwards.

The basic assumption that the Monte Carlo algorithm makes is that $D_c = D_{c,\text{true}}$ and (by extension) $M_c^{\text{PCS}} = M_{c,\text{prox}}$. Upon making this assumption, it proceeds to construct the distribution of $\Delta\chi_c^2(M_c^{\text{PCS}})$ by iterating over an ensemble of alternative observed LC pairs, constructed via repeated sampling of a multivariate normal distribution with mean D and standard deviations E . In practice this sampling means that, for the k th alternative observed LC pair $D_{c,k}$ in the ensemble (i.e., the LC pair involved in the k th iteration), the intensity in the i th bin of, for instance, its radio component $D_{r,k}$, is taken from a univariate normal distribution with mean $D_{r,i}$ and standard deviation $E_{r,i}$; i.e., with mean and standard deviation equal to the observed intensity in the i th radio LC bin and the uncertainty on that intensity. In effect, each LC pair in the ensemble is constituted by perturbing D_c based on the LC uncertainties E_c .

Since each simulated observed LC pair (i.e., each $D_{c,k}$) is different, the specific χ_c^2 statistic used is effectively different for each iteration; for the k th iteration the relevant statistic is (as per Eq. (3.10))

$$\chi_{c,k}^2 = \chi_c^2(M_c; D_{c,k}, E_c), \quad (7.11)$$

which identifies the LC pair $M_{c,k}^{\text{PCS}}$ as the model's best fit (to $D_{c,k}$) in the k th iteration (for which $\chi_{c,k}^2(M_{c,k}^{\text{PCS}}) = \chi_{c,k,\text{min}}^2$). The quantity of interest for each iteration is the difference

$$\begin{aligned} \Delta\chi_{c,k}^2(M_c^{\text{PCS}}) &= \chi_{c,k}^2(M_c^{\text{PCS}}) - \chi_{c,k}^2(M_{c,k}^{\text{PCS}}) \\ &= \chi_{c,k}^2(M_c^{\text{PCS}}) - \chi_{c,k,\text{min}}^2, \end{aligned} \quad (7.12)$$

which is the k th-iteration analogue of the difference $\Delta\chi_c^2(M_{\text{true}})$ (the ensemble distribution of which yields constraints in the valid-model case).

The diagram in the left-hand panel of Figure 7.3 is a representation of the procedure at the core of the Monte Carlo algorithm. The cloud of grey dots centred on D_c (which is denser toward its centre) is the ensemble that this procedure leverages. The best-fit LC pair for the k th iteration (i.e., to $D_{c,k}$) is $M_{c,k}^{\text{PCS}}$, and the locus of all possible LC pairs that would fit $D_{c,k}$ equally well (according to $\chi_{c,k}^2$) is the n_c -dimensional hypersphere centred on $D_{c,k}$ with radius $\chi_{c,k}(M_{c,k}^{\text{PCS}})$ (the right-most red arc). As might be expected, the initial best-fit LC pair M_c^{PCS} (which fits D_c best, according to χ_c^2) is a worse fit to $D_{c,k}$ than $M_{c,k}^{\text{PCS}}$ is. This is reflected in the fact that the former lies on a larger $D_{c,k}$ -centred hypersphere (the left-most red arc) than does the latter, since $\chi_{c,k}(M_c^{\text{PCS}}) > \chi_{c,k}(M_{c,k}^{\text{PCS}})$ (in this example). Within this framework, $\Delta\chi_{c,k}^2(M_c^{\text{PCS}})$ effectively characterises the difference of these two hyperspheres. Specifically, this difference is the square of the distance between $M_{c,k}^{\text{PCS}}$ and the n_c -dimensional hypersphere with radius $\chi(M_c^{\text{PCS}})$ (i.e., the $D_{c,k}$ -centred hypersphere that contains M_c^{PCS}) along a line tangent to the $D_{c,k}$ -centred hypersphere that contains $M_{c,k}^{\text{PCS}}$.

Termination via Convergence In the interest of reliability, the algorithm terminates iteration based on a convergence criterion for $\delta_{2\sigma,\text{PCS}}$ (for χ_c^2 ; which is calculated at the end of each iteration). After the first 400 Monte Carlo iterations, and




Figure 7.3: The quantities associated with $D_{c,k}$, the k th alternative observed LC pair in the ensemble; i.e., associated with the LC pair involved in the k th iteration of the Monte Carlo simulation.

at 200 iteration intervals thereafter, the standard deviation in $\delta_{2\sigma, \text{pcs}}$ across the preceding 200 iterations is calculated. At any of these checks, the algorithm is taken to have converged if this deviation is less than 5% of $\delta_{2\sigma, \text{pcs}}$ itself. To mitigate the possibility of a false positive, such a convergent estimate is not immediately accepted, however. An estimate is only accepted once two consecutive checks for convergence are positive.

The final value for $\delta_{2\sigma, \text{pcs}}$ is then converted into an acceptance contour in the model's parameter space by identifying all the LC pairs for which $\chi_c^2 \leq \chi_{c, \text{min}}^2 + \delta_{2\sigma, \text{pcs}}$. The extent of this contour for each parameter then translates into the desired constraint on that parameter.

7.4.3 Applied to Fits using $\Psi_{\Phi, c}^2$

As alluded to earlier, the extension of this algorithm to $\Psi_{\Phi, c}^2$ is comparatively simple: For the k th iteration, the relevant GOF statistic is (as per Eq. (5.48))

$$\Psi_{\Phi, c, k}^2(M_c) = \Psi_{\Phi, c}^2(M_c; D_{c, k}, E_c, B_c), \quad (7.13)$$

where B_c is background LC pair implied by the background levels b_r and b_γ , and the quantity of interest is

$$\begin{aligned} \Delta \Psi_{\Phi, c, k}^2(M_c^{\text{sfs}}) &= \Psi_{\Phi, c, k}^2(M_{c, k}^{\text{sfs}}) - \Psi_{\Phi, c, k}^2(M_c^{\text{sfs}}) \\ &= [1 - \chi_{\Phi, c, k}^2(M_{c, k}^{\text{sfs}})] - [1 - \chi_{\Phi, c, k}^2(M_c^{\text{sfs}})] \\ &= \chi_{\Phi, c, k}^2(M_c^{\text{sfs}}) - \chi_{\Phi, c, k}^2(M_{c, k}^{\text{sfs}}) \\ &= \Delta \chi_{\Phi, c, k}^2(M_c^{\text{sfs}}), \end{aligned} \quad (7.14)$$

where M_c^{sfs} is the best fit to D_c and $M_{c, k}^{\text{sfs}}$ is the best fit to $D_{c, k}$; i.e., $\Psi_{\Phi, c, k}^2(M_{c, k}^{\text{sfs}}) = \Psi_{\Phi, c, k}^2(M_{c, k}^{\text{sfs}})$ and $\chi_{\Phi, c, k}^2(M_{c, k}^{\text{sfs}}) = \chi_{\Phi, c, k}^2(M_{c, k}^{\text{sfs}})$.

The diagram in the right-hand panel of Figure 7.3 is a representation of the geometry at play for the $\Psi_{\Phi, c}^2$ statistic. In both the left- and right-hand panes the same synthetic observed LC pair $D_{c, k}$ is being fit with the same model $M_c(\vartheta)$, but using two different statistics ($\chi_{c, k}^2$ and $\Psi_{\Phi, c, k}^2$). Notice that the mapping in the right-hand panel (under $[\cdot]_{\mathcal{E}_c^*}$) is essentially an anisotropic scaling of the mapping in the left-hand panel (under $[\cdot]_{\mathcal{E}_c}$), just as in Figure 5.4. As a consequence, the image (in $\mathbb{R}^{n_c}$) of the ensemble of synthetic observed LCs also appears scaled; specifically, in the left-hand mapping its image is roughly an n_c -dimensional hypersphere while in the right-hand mapping its image is roughly an n_c -dimensional hyperellipse (see Section 5.2.4). The image the model produces in $\mathbb{R}^{n_c}$ is similarly affected, which causes a different LC pair to be identified as the model's best fit (as compared to the fit in the left-hand pane; M_c^{sfs} instead of M_c^{pcs}). Similarly, $M_{c, \text{prox}}^{\text{sfs}}$ is now the model LC pair that is closest to $D_{c, \text{true}}$ (instead of $M_{c, \text{prox}}^{\text{pcs}}$), and $M_{c, k}^{\text{sfs}}$ is the best fit to $D_{c, k}$ (instead of $M_{c, k}^{\text{pcs}}$). For comparison, the locations of $M_{c, \text{prox}}^{\text{pcs}}$, M_c^{pcs} and $M_{c, k}^{\text{pcs}}$ are shown in the right-hand panel of Figure 7.3 using half-sized symbols. Note that in both cases the difference being calculated characterizes the difference of two hyperspheres centred on $D_{c, k}$.

Leveraging this extension, the Monte Carlo algorithm yields the approximate distribution of $\Delta \Psi_{\Phi, c}^2(M_{c, \text{prox}})$, and

hence the requisite critical value estimate $\delta_{2\sigma,\text{sfs}}$. Since the joint best fits are identified by maximisation (rather than minimization, as is the case for χ_c^2), the 2σ confidence contours are constructed in the model's parameter space by identifying all the model LC pairs for which

$$\Psi_{\Phi,c}^2(M_c) \geq \Psi_{\Phi,c,\text{max}}^2 - \delta_{2\sigma,\text{sfs}}. \quad (7.15)$$

These contours are again translated into constraints on the model parameters by finding the smallest n_θ -dimensional hyperrectangle that encloses them. See Appendix B for examples of such contours.

The Distribution of $\underline{\Delta}\Psi_{\Phi,c}^2(M_{\text{prox}})$

The mapping in the right-hand panel of Figure 7.3 demonstrates that since the image of the ensemble in $\mathbb{R}^{n_c}$ is not a hypersphere under $[\cdot]_{\mathcal{E}_c^*}$, it is more difficult to determine the precise distribution of $\underline{\Delta}\Psi_{\Phi,c}^2(M_{\text{prox}})$ than it was to determine that of $\underline{\Delta}\chi_c^2(M_{\text{prox}})$ under $[\cdot]_{\mathcal{E}_c}$. Since this hyperellipse does not have spherical symmetry, the distribution of $\underline{\Delta}\Psi_{\Phi,c}^2(M_{\text{prox}})$ effectively depends on the relative orientations of the model's and ensemble's images (i.e., the direction of the vector that would connect $D_{c,\text{true}}$ and $M_{c,\text{prox}}$). Had the model's image under $[\cdot]_{\mathcal{E}_c^*}$ been such that its depiction in the right-hand panel was a horizontal line that passes above the ensemble, for instance, the distribution of $\underline{\Delta}\Psi_{\Phi,c}^2(M_{\text{prox}})$ would have been narrower (i.e., the expected value of $\underline{\Delta}\Psi_{\Phi,c}^2(M_{\text{prox}})$ would have been smaller). Fortunately, the estimate the Monte Carlo algorithm yields for this distribution, is sensitive to this directionality. Therefore, while it is likely possible to determine this distribution analytically, doing so is beyond the scope of this work.

Chapter 8

The Effectiveness of $\Psi_{\Phi,c}^2$ for Synthetic LC Pairs

As a first test of how effectively $\Psi_{\Phi,c}^2$ eliminates single-band dominance in parameter estimation, we derive parameter point estimates for a set of synthetic observed LC pairs using both χ_c^2 and $\Psi_{\Phi,c}^2$, and compare the results using the tools developed in Chapter 6. Conveniently, the resulting set of point estimates can also be leveraged to demonstrate the convergence of χ_c^2 and $\Psi_{\Phi,c}^2$ (with regard to the point estimates they yield). The set of synthetic observed LC pairs was generated using the model outputs (for the TPC and Cone single-band models) for a pulsar with a period of 0.1668 s and a period derivative of $1.68 \times 10^{-14} \text{ s.s}^{-1}$; the period and period derivative of PSR J0742–2822. While, in principle, this test could also be performed using the outputs for the OG γ -ray model instead of those for the TPC γ -ray model, we opted to use the latter model’s outputs since it yields non-flat LCs for a larger number of parameter combinations.

Note that, as per the discussion in Section 3.2.2, this test is not designed to assess the accuracy of the dual-band parameter estimates that either statistic yields; rather, it is designed to assess the degree to which these parameter estimates are radio-dominated.

8.1 The Set of Synthetic Observed LC Pairs

Since, as described in Section 3.2.2, radio dominance only occurs when both single-band best-fit non-colocation and uncertainty disparity are present, the set of synthetic observed LC pairs was generated such that this is the case for a large fraction of the LC pairs that comprise it. Additionally, since these LC pairs are intended to serve as synthetic observed LC pairs, as opposed to synthetic true LC pairs, the raw model LC pairs were not used directly. Instead, each raw model LC pair was, in effect, subjected to a synthetic observation procedure.

8.1.1 Synthetic Observation

The intensity in each LC bin (across both LCs in each LC pair) was taken from a normal distribution with average equal to the corresponding normalised intensity in the raw model LC, and standard deviation equal to 17% of the intensity of that LC’s peak (i.e., since the raw model LCs are each normalised, an absolute standard deviation of 0.17). This value for the uniform standard deviation was chosen to ensure that none of the fits to the synthetic observed LC pairs would suffer from overfitting (see Section 8.1.3). As a last step, each of the synthetic observed LCs was normalised such that its peak intensity was 1.

8.1.2 Synthetic Non-colocation

To simulate the effect of non-colocation, the restriction that α and ζ have the same values for both single-band models (which is enforced for the sets of LC pairs that the TPC+Cone and OG+Cone models produce) was relaxed for the generation of the synthetic observed LC pairs; i.e., the values for both α and ζ were allowed to be different for the two




Figure 8.1: An example synthetic LC pair (plotted in grey in panels (c) and (d)). Here, the renormalisation of the individual constant- ζ cuts through the phaseplots (in panels (a) and (b)) has already been implemented (i.e., the phaseplots are normalised such that at each point on each of them, the intensity plotted is relative to the maximum intensity of the relevant constant- ζ cut). The plotted synthetic LCs are generated by adding Gaussian noise to a pair of model LCs, with standard deviation equal to 0.17 (chosen to be larger than the largest choice for ε_γ). Unlike for the LC pair in Figure 2.1, though, α and ζ are not the same for the two LCs. This facilitates the simulation of single-band best-fit non-colocation.

wavebands. Each synthetic observed LC pair is therefore associated with four parameter values: α_r and ζ_r for the radio component, and α_γ and ζ_γ for the γ -ray component. In accordance with the discussion in Section 7.2.2, these parameters were confined to lie in the interval $[0^\circ, 90^\circ]$, and for each parameter this interval was probed at a 2° resolution.

Only those parameter combinations where neither component raw model LC was flat were considered in the generation of the final set of synthetic observed LC pairs. Since the Cone model produces non-flat LCs for only 782 combinations of α_r and ζ_r (for the chosen values for P and $\dot{P}$), and the TPC model for only 1923 combinations of α_γ and ζ_γ , the set of synthetic observed LC pairs contains $782 \times 1923 = 1\,503\,786$ LC pairs (see Figure 8.2).

Figure 8.1 shows one such synthetic observed LC pair (prior to the renormalisation of its component LCs' peak intensities; plotted in grey in panels (c) and (d)), associated with the parameters $(\alpha_r, \zeta_r) = (6^\circ, 8^\circ)$ and $(\alpha_\gamma, \zeta_\gamma) = (48^\circ, 36^\circ)$. As this LC pair demonstrates, we assume a background level of 0 for both the radio and γ -ray bands. The phaseplots in panels (a) and (b) (with the relevant constant- ζ cuts indicated in yellow) differ from those in Figure 2.1 in that they are normalised such that the peak intensity for each cut is 1.

8.1.3 Synthetic Uncertainty Disparity

To simulate the effect of a large uncertainty disparity, and to investigate what happens as it dissipates, each synthetic observed LC pair was fit multiple times: once for each of five different combinations of the radio and γ -ray uncertainties ε_r and ε_γ . Since, for uniform uncertainties,

$$\lambda_{r\gamma} \propto \frac{\varepsilon_\gamma}{\varepsilon_r}, \quad (8.1)$$

we chose a fixed ε_r of 0.04 for all the fits, and assigned to ε_γ each of the values 0.01, 0.02, 0.04, 0.08, and 0.16, in turn; hence including values for the ratio $\varepsilon_\gamma/\varepsilon_r$ that cover the range $[1/4, 4]$. Notice that the uniform standard deviation of 0.17 used in the generation of the synthetic LC pairs is larger than any of the values used for ε_γ .

Figure 8.2 shows maps of the $\log(\Phi_r)$ and $\log(\Phi_\gamma)$ values (*left* and *right* panels, respectively; at a 1° resolution in α and ζ) that result when the radio and γ -ray uncertainties are equal (i.e., when $\varepsilon_\gamma = \varepsilon_r = 0.04$); with all the possible




Figure 8.2: Radio (*left*) and γ -ray (*right*) scaled fluxes Φ_r and Φ_γ (plotted logarithmically) as a function of viewing geometry (α and ζ), for the entire set of possible synthetic LCs (generated at a 1° resolution; with $\varepsilon_\gamma = \varepsilon_r = 0.04$). The set of synthetic LCs used was generated by sampling (α, ζ) -space at a resolution of 2° (excluding parameter combinations for which $\alpha = 0^\circ$ or $\zeta = 0^\circ$), as indicated by the grey dots. The locations of the two synthetic LCs shown in Figure 8.1 are also indicated (a square for the synthetic radio LC and diamond for the synthetic γ -ray LC, both black). For this pair of LCs (and choice for ε_γ), $\log(\Phi_r) \simeq 2.19$ and $\log(\Phi_\gamma) \simeq 2.04$; therefore, $\log(\lambda_{r\gamma}) \simeq 0.15$ (as per Eq. (6.9)).

2° resolution choices for (α_r, ζ_r) and $(\alpha_\gamma, \zeta_\gamma)$ indicated. A useful advantage of plotting these maps is that the value of $\log(\lambda_{r\gamma})$ associated with each synthetic observed LC pair (which characterises uncertainty disparity associated with that LC pair, as per the discussion in Section 6.2.1) can be easily recovered by subtracting the colour on the γ -ray map from the colour on the radio map (in accordance with Eq. (6.9)). For the synthetic observed LC pair plotted in Figure 8.1, for example (associated with the parameter combinations for the TPC and Cone models indicated in Figure 8.2), $\log(\Phi_r) \simeq 2.19$ and $\log(\Phi_\gamma) \simeq 2.04$ (with $\varepsilon_r = \varepsilon_\gamma = 0.04$), which means that $\log(\lambda_{r\gamma}) \simeq 2.19 - 2.04 = 0.15$ for the corresponding fits.

8.1.4 Coverage of $(\eta_s, \lambda_{r\gamma})$ -space

As described in Section 3.2.2, χ_c^2 and $\Psi_{\Phi,c}^2$ tend to disagree more (in terms of the parameter point estimates they yield) for dual-band fits where the single-band best fits are more non-colocated, and the LC uncertainty disparity is larger; i.e., η_c tends to grow as η_s and $|\log(\lambda_{r\gamma})|$ become larger. Therefore, since we intend to also demonstrate the convergence of χ_c^2 and $\Psi_{\Phi,c}^2$ (in terms of the parameter point estimates they yield), it is important that:

1. the synthetic observed LC pairs generated here cover a sufficiently large area in $(\eta_s, \lambda_{r\gamma})$ -space, and
2. the region they cover includes the region around the point $(\eta_s = 0, \lambda_{r\gamma} = 1)$ in this space, otherwise χ_c^2 and $\Psi_{\Phi,c}^2$ will either never disagree or never have the opportunity to converge.

Fortunately, the latter of these two requirements is satisfied almost by construction, since $\eta_s = 0$ for combinations where $(\alpha_r, \zeta_r) = (\alpha_\gamma, \zeta_\gamma)$; and $\lambda_{r\gamma} \approx 1$ for a large proportion of the fits thanks to the range of values chosen for ε_γ . The former of these two requirements turns out to also be satisfied a posteriori, as the results of this test will reflect.

8.2 The Set of Model LC Pairs

As described above, the raw model LC pairs serve, in effect, as true counterparts to the synthetic observed LC pairs. Therefore, for these fits (as opposed to real-world fits), both D_{true} and the value of ϕ_μ are effectively known. This means




Figure 8.3: Joint frequency distribution of f_r and $\log(\lambda_{r\gamma})$ for the fits conducted using χ_c^2 and $\Psi_{\Phi,c}^2$ (top and bottom rows), for each of three of the chosen values for ε_γ (0.01, 0.04, and 0.16; left, middle, and right columns). The six black dots locate the six joint fits associated with the example synthetic observed LC pair plotted in Figure 8.1.

that the set of model LC pairs used to fit each of the synthetic LC pairs, need not be imbued with as much freedom as those constructed in Section 7.3. In fact, for this test it is sufficient to use the raw model LC pairs themselves to fit each of the synthetic observed LC pairs. The only restriction (which stems from the conglomerated nature of the TPC+Cone dual-band model) that remains in place is that, for each model LC pair it should be the case that $(\alpha_r, \zeta_r) = (\alpha_\gamma, \zeta_\gamma)$.

Therefore, the set of model LC pairs used for these fits consists of all the LC pairs the TPC+Cone model produces where neither component LC is flat; 782 LC pairs in total, since the TPC model produces non-flat LCs everywhere where the Cone model does. Note that the resolution at which α and ζ are probed to construct the set of model LC pairs is 1° , as opposed to the 2° resolution at which α_r , ζ_r , α_γ and ζ_γ were probed to construct the set of synthetic observed LC pairs.

8.3 Results

Given the five values used for ε_γ , each fit to a synthetic LC pair yielded a total of 16 best-fit parameter point estimates: one radio-only estimate, five γ -ray-only estimates, and ten joint estimates (five each for χ_c^2 and $\Psi_{\Phi,c}^2$). Based on the GOF values associated with these 16 estimates, each joint fit was then characterised using the tools introduced in Chapter 6.

8.3.1 The Effectiveness of $\Psi_{\Phi,c}^2$

Figure 8.3 shows the joint frequency distribution of f_r and $\log(\lambda_{r\gamma})$ for the fits conducted using χ_c^2 and $\Psi_{\Phi,c}^2$ (top and bottom rows), for each of three of the chosen values for ε_γ (0.01, 0.04, and 0.16; left, middle, and right columns). These distributions are plotted as binned heat maps, with 400 bins along the $\log(\lambda_{r\gamma})$ axis (which cover the range of represented $\log(\lambda_{r\gamma})$ values in each plot) and 200 bins along the f_r axis. Each of the three values for ε_γ is associated with a different colour, and the colour intensity in each bin indicates how many fits it contains. Each synthetic observed LC pair is associated with exactly one fit on each of the six plots in this figure: the fit to that synthetic LC wherein the stated statistic and value for ε_γ were used. Since for each of these six fits $\lambda_{r\gamma} \propto \varepsilon_\gamma$ (as per Eq. (8.1)); since ε_r is the same for all the fits,




Figure 8.4: Plots of F as function of $\varepsilon_\gamma/\varepsilon_r$ for the five sets of fits conducted using χ_c^2 (a black “X” for each set of fits) and $\Psi_{\Phi,c}^2$ (a black dot for each set of fits). For each of the sets of fits conducted using $\Psi_{\Phi,c}^2$, and the set of fits conducted using χ_c^2 for which $\varepsilon_\gamma/\varepsilon_r = 1$ (i.e., for which the LC errors are balanced), $|F| < 0.5$, which indicates that χ_c^2 is undominated when the LC uncertainties are balanced, and $\Psi_{\Phi,c}^2$ is undominated regardless of relative sizes of the LC uncertainties. For the four remaining sets of fits $F > 0.5$, which indicates that χ_c^2 suffers from precision-determined dominance for these sets of fits, tending to favour the more precise band in each case (since $F > 0$).

they are shifted horizontally with respect to each other across the three columns; specifically, by a (logarithmic) distance of $\Delta \log(\lambda_{r\gamma}) = \log(4) \simeq 0.6$. Additionally, since the $\log(\lambda_{r\gamma})$ value associated with a fit is independent of the statistic used to conduct it, the six fits associated with each synthetic LC are, at most, shifted along the f_r axis with respect to each other across the two rows. Therefore, within each column, the marginal distribution of $\log(\lambda_{r\gamma})$ is identical for both plots, and only the marginal distribution of f_r is sensitive to the choice of statistic.

As an example, the six black dots in Figure 8.3 locate the six joint fits associated with the example synthetic observed LC pair plotted in Figure 8.1. Notice that $\log(\lambda_{r\gamma}) \simeq 0.15$ for both dots in the middle column, in accordance with the discussion in Section 8.1.3. For the three fits conducted using χ_c^2 (i.e., those represented in the distributions in the top row), $f_{r,\text{pcs}} = -1, -0.19$, and 1 , respectively. For all three of the fits conducted using $\Psi_{\Phi,c}^2$ (represented in the distributions in the bottom row), $f_{r,\text{sfs}} = -0.19$.

Note that these plots, which are a realisation of the plots shown in Figure 6.2, do not carry information on the η_s value associated with each synthetic observed LC pair (which is independent of the choice for ε_γ). In effect, the two-dimensional distributions shown in these plots are the result of marginalisation of the full three-dimensional distributions, over η_s . They are, however, sufficient for this test, since it is primarily focused on the effect of uncertainty disparity.

The Fits Conducted using χ_c^2

These fits display clear single-band dominance when the uncertainty in either band is much larger than that in the other, with fits clustering near $f_r = -1$ for $\varepsilon_\gamma < \varepsilon_r$ (higher γ -ray precision; panel (a)) and near $f_r = 1$ for $\varepsilon_\gamma > \varepsilon_r$ (higher radio precision; panel (c)). When the uncertainties in the two bands are equal, as in the middle column, the fits appear undominated; i.e., the f_r values are more evenly distributed in panel (b). This behaviour is indicative of precision-determined priority, as predicted for χ_c^2 in Section 6.3.

Figure 8.4 plots the values for F_{pcs} and F_{sfs} (see Section 6.3) as a function of the ratio $(\varepsilon_\gamma/\varepsilon_r)$. The plot for F_{pcs} corroborates the above assessment, since for the two plots in Figure 8.3 where single-band dominance is evident (panels (a) and (c)), F_{pcs} is comparatively close to 1 (0.930 and 0.860, respectively), and for the plot in Figure 8.3 where the fits appear undominated (panel (b)), $F_{\text{pcs}} = 0.135$, which is comparatively close to 0. Taking the χ_c^2 fits across all five choices for ε_γ as a single, monolithic set of fits yields $F_{\text{pcs}} = 0.763$ (for the sets of fits where $\varepsilon_\gamma = 0.02$ and 0.08 , $F_{\text{pcs}} = 0.712$ and 0.555 , respectively), which again confirms that χ_c^2 is prone to precision-determined single-band dominance.




Figure 8.5: Plots of η_c as a function of $\log(\eta_s)$ and $\log(\lambda_{r\gamma})$ for ε_γ choices of 0.01, 0.04, and 0.08 (with $\varepsilon_r = 0.04$). Red dots are joint fits where χ_c^2 gives the radio GOF higher priority than $\Psi_{\Phi,c}^2$ does, with the reverse being the case for blue dots. For each colour, the colour intensity increases as η_c increases. The three black circles locate the fits to the synthetic LC pair in Figure 8.1.

The Fits Conducted using $\Psi_{\Phi,c}^2$

Unlike for the fits conducted using χ_c^2 , the marginal distribution of f_r for these fits is identical across the three columns in Figure 8.3; i.e., it is independent of the choice for ε_γ . Moreover, since this distribution (similarly to that in panel (b)) is roughly even and symmetrical about $f_r = 0$ (at least compared to those in panels (a) and (c) in Figure 8.3), it appears that all three sets of fits conducted using $\Psi_{\Phi,c}^2$ are undominated. This assessment is corroborated by the values for F_{sfs} plotted in Figure 8.4, which are all comparatively close to 0, with $F_{\text{sfs}} = -0.112$ for the $\Psi_{\Phi,c}^2$ fits as a whole.

Interestingly, $F_{\text{sfs}} < 0$ for four of the choices of ε_γ , as well as for the $\Psi_{\Phi,c}^2$ fits as a whole. In particular, it is at its most extreme ($F_{\text{sfs}} = -0.311$) for $\varepsilon_\gamma = \varepsilon_r$. This indicates that, while the radio priority associated with a $\Psi_{\Phi,c}^2$ fit is independent of the choice for ε_γ , it still depends on the value for $\lambda_{r\gamma}$. Looking more closely at the distribution plotted in panel (e), we see that f_r tends to be closer to 1 for fits where $\log(\lambda_{r\gamma}) > 0$, and closer to -1 for fits where $\log(\lambda_{r\gamma}) < 0$. While an investigation of this second-order dependence is somewhat outside the scope of this test, it seems to be related to the duty cycles of the component observed LCs. In essence, it seems that the waveband with the shorter duty cycle is less capable of allowing compromise, since the range of possible LCs for which $\Psi_{\Phi}^2 > 0$ when fit to it, is more limited; hence, the waveband with the shorter duty cycle tends to be favoured.

Considering panel (e) within the context of Figure 8.2 reveals that the synthetic observed LC pairs for which $\log(\lambda_{r\gamma}) > 0$ in panel (e), do indeed result from the combination of radio LCs with longer duty cycles (that lie toward the lower left on the radio plot in Figure 8.2) with γ -ray LCs with shorter duty cycles (that lie toward the upper right on the γ -ray plot in Figure 8.2; as per Eq. (6.9)), and vice versa for the synthetic observed LC pairs for which $\log(\lambda_{r\gamma}) < 0$.

8.3.2 The Convergence of χ_c^2 and $\Psi_{\Phi,c}^2$

Figure 8.5 shows η_c as a function of η_s (plotted logarithmically) and $\log(\lambda_{r\gamma})$ for three of the five chosen γ -ray uncertainties: 0.01, 0.04, and 0.16. On each of the three plots, each dot corresponds to a different synthetic observed LC pair. Each synthetic observed LC pair corresponds to a dot on each of the three plots since it was fit once for each of the three γ -ray uncertainties. Since η_s does not depend on the choice of γ -ray uncertainty, these three dots appear in the same vertical position on each plot, with only the horizontal position changing since $\lambda_{r\gamma} \propto \varepsilon_\gamma$. Dots that appear to the left of the yellow dotted vertical line correspond to synthetic observed LCs for which $\Phi_r < \Phi_\gamma$, and vice versa for dots to the right.

For each dot, the colour intensity encodes the value of η_c , where fully white dots correspond to colocated dual-band best fits (i.e., $\eta_c = 0$), and fully red or blue dots correspond to fits for which η_c is a maximum. Red dots are fits for which the radio priority is higher for χ_c^2 than for $\Psi_{\Phi,c}^2$ (i.e., for which $f_{r,\text{pcs}} > f_{r,\text{sfs}}$), while blue dots are fits for which the radio priority is lower (i.e., for which $f_{r,\text{pcs}} < f_{r,\text{sfs}}$). In all three panels, dots that lie to the right of the yellow dotted vertical line are invariably either fully white or red, and vice versa for dots that lie to the left of that line. Note that in Figure 8.5, in




Figure 8.6: A more detailed, component-wise breakdown of the plots in Figure 8.5. Each coloured, dashed line (across all four panels) corresponds to a different choice for ε_γ (including the two choices omitted in Figure 8.5), and the thick black line corresponds to the entire set of synthetic fits (i.e., includes the fits for all five choices for ε_γ). In panel (a), the fraction N_{agree}/N of fits where χ_c^2 and $\Psi_{\Phi,c}^2$ agree (i.e., for which M_c^{pcs} and M_c^{sfs} are colocated and the corresponding dot in Figure 8.5 is fully white) is plotted as a function of $\log(\lambda_{\text{r}\gamma})$. This fraction is close to 0 at the extremes of $\log(\lambda_{\text{r}\gamma})$, and close to 1 at $\log(\lambda_{\text{r}\gamma}) = 0$ (as expected for convergence). In panel (b), the average dual-band best-fit non-colocation $\eta_{c,\text{ave}}$ (in the form of $\log(\eta_{c,\text{ave}})$; for the fits where χ_c^2 and $\Psi_{\Phi,c}^2$ disagree and the corresponding dot in Figure 8.5 is either red or blue) is plotted as a function of $\log(\lambda_{\text{r}\gamma})$. As expected for convergence, $\eta_{c,\text{ave}}$ approaches 0 as $|\log(\lambda_{\text{r}\gamma})| \rightarrow 0$. In panels (b) and (d) the same quantities are plotted as in panels (a) and (c), but as a function of η_s (in the form of $\log(\eta_s)$). As expected for convergence, $N_{\text{agree}}/N \rightarrow 1$ and $\eta_c \rightarrow 0$ as $\eta_s \rightarrow 0$. As indicated by the arrows in panels (b) and (d), $N_{\text{agree}}/N = 1$ and $\eta_c = 0$ for the 12990 fits for which $\eta_s = 0$ (i.e., at $-\infty$ on these plots).

the interest of clarity, more intensely coloured dots are plotted on top of less intensely coloured (and fully white) dots.

As an example, the three black circles in Figure 8.5 enclose the dots associated with the three pairs of dual-band fits conducted for the synthetic observed LC pair plotted in Figure 8.1. For the pair of fits conducted with $\varepsilon_\gamma = 0.01$ (i.e., associated with the circle on the left-hand plot in Figure 8.5, and represented in the distributions in the left-hand column of Figure 8.3), $f_{\text{r,pcs}} = -1 < -0.19 = f_{\text{r,sfs}}$ (and $\eta_c = 0.002$), and hence the plotted dot is blue. For the pair of fits conducted with $\varepsilon_\gamma = 0.16$ (the circle on the right-hand plot in Figure 8.5, and those in the right-hand column of Figure 8.3), the plotted dot is red since $f_{\text{r,pcs}} = 1 > -0.19 = f_{\text{r,sfs}}$ (with $\eta_c = 0.120$). For the pair of fits conducted with $\varepsilon_\gamma = 0.04$ (the circle on the middle plot, and those in the middle column of Figure 8.3), the plotted dot is white since $\eta_c = 0$ (and therefore lies below the more intensely red dots plotted above it).

For all three plots, the colour intensity of the dots tends toward zero as $\log(\lambda_{\text{r}\gamma}) \rightarrow 0$, with all the dots that lie near the yellow dotted vertical line being uniformly fully white (i.e., χ_c^2 and $\Psi_{\Phi,c}^2$ identify the same LC pair as the joint best fit for the dots close to the yellow dotted line). Figure 8.6 (a) and (c) demonstrate this trend in greater detail by breaking it up into two components: the fraction of fits where χ_c^2 and $\Psi_{\Phi,c}^2$ agree (N_{agree}/N) and the average η_c for the fits where they do not agree (plotted logarithmically). In each pane, the coloured lines correspond to the indicated values for ε_γ , and the thick black line corresponds to the average across all five γ -ray uncertainties. As is expected for convergence, $N_{\text{agree}}/N \rightarrow 1$ and $\eta_{c,\text{ave}} \rightarrow 0$ as $\log(\lambda_{\text{r}\gamma}) \rightarrow 0$.

Similarly, for all three plots in Figure 8.5, the colour intensity of the dots tends toward zero as $\eta_s \rightarrow 0$. Figure 8.6 (b) and (d) demonstrate this trend in greater detail by plotting the same quantities as before as a function of $\log(\eta_s)$. The

greyed-out portions on each plot indicate where $\eta_s > 1$, where either the radio or γ -ray component fit is worse than a null fit. The arrows in the top left and bottom left corners of panels (b) and (d), respectively, indicate that the respective trends become limits at $\eta_s = 0$. In total, $\eta_s = 0$ for 12990 of the fits (i.e., for 2598 synthetic observed LC pairs, each fit five times), and for each of these fits χ_c^2 and $\Psi_{\Phi,c}^2$ agree. As is expected for convergence, then, $N_{\text{agree}}/N \rightarrow 1$ and $\eta_{c,\text{ave}} \rightarrow 0$ as $\eta_s \rightarrow 0$, just as for $\log(\lambda_{r\gamma})$.

Interestingly, Figure 8.6 (b) and (d) reveal a trend with regard to the effect of the chosen γ -ray uncertainty. Specifically, the vertical positions of the lines corresponding to the various γ -ray uncertainties seem to roughly correlate with how much that uncertainty deviates from the fixed radio uncertainty, i.e., with $|\log(\varepsilon_r/\varepsilon_\gamma)|$.¹ On both plots, the purple and red lines, which both correspond to $|\log(\varepsilon_r/\varepsilon_\gamma)| \simeq 0.6$, tend to perform more poorly in terms of agreement and dual-band best-fit colocation than the green lines do, which correspond to $|\log(\varepsilon_r/\varepsilon_\gamma)| = 0$. In effect, this trend for $|\log(\varepsilon_r/\varepsilon_\gamma)|$ is a weaker version of the convergence posited for $\log(\lambda_{r\gamma})$, and, as such, essentially amounts to a demonstration that χ_c^2 and $\Psi_{\Phi,c}^2$ roughly converge as the relative uncertainties become equal.

Crucially, the plots in Figure 8.3 also reveal that it is, in effect, rather the case that the results of χ_c^2 converge to those of $\Psi_{\Phi,c}^2$ as $\log(\lambda_{r\gamma}) \rightarrow 0$, than the other way around. Stated differently, these plots reveal that $\Psi_{\Phi,c}^2$ is able to recover the result that χ_c^2 would yield in the absence of an uncertainty disparity.

¹Note that $|\log(\varepsilon_r/\varepsilon_\gamma)| = |\log(\varepsilon_\gamma/\varepsilon_r)|$.

Part IV

APPLICATION

Chapter 9

Case Study: Dual-band Fits for PSR J2039–5617

As stated in the [Introduction](#), the $\Psi_{\Phi,c}^2$ statistic has already been used to derive published parameter estimates for a real-world pulsar: PSR J2039–5617. While the fits to this pulsar do not form part of the main body of this research (i.e., they have no bearing on the effectiveness of the $\Psi_{\Phi,c}^2$ statistic in terms of the elimination of radio dominance), they do provide an ideal opportunity to demonstrate what roles the single- and dual-band statistics introduced in this work (i.e., the Ψ_{Φ}^2 statistic for single-band fits in the form of $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$, and $\Psi_{\Phi,c}^2$ for dual-band fits) should play when conducting joint, dual-band fits aimed at deriving meaningful parameter estimates. Specifically, they demonstrate why each of these statistics should be used in concert with its counterpart χ^2 statistic (i.e., with χ_r^2 , χ_{γ}^2 and χ_c^2 , respectively; via their reduced equivalents $\chi_{\nu,r}^2$, $\chi_{\nu,\gamma}^2$ and $[\chi_c^2]_{\nu}$, in this case): In each case the Ψ_{Φ}^2 statistic encodes how well the observed LC is reproduced (or how well the observed LC pair as a whole is reproduced), while the χ^2 statistic encodes how significantly the model’s reproduction deviates from what is required. To this end, this chapter contains a condensed version of the account contained in [Corongiu et al. \(2020\)](#) of the dual-band fits conducted for PSR J2039–5617, with a brief discussion of what role each of the various statistics served.

9.1 Background

Subsequent to its identification as a source of high-energy γ -rays (0.1 – 100 GeV) in the second *Fermi* LAT source catalogue (2FGL; [Nolan et al., 2012](#)), the stellar system that PSR J2039–5617 belongs to was discovered to be a close optical/X-ray binary with an orbital period of ~ 5.5 hr by [Romani \(2015\)](#) and [Salvetti et al. \(2015\)](#).¹ The 2FGL source designation associated with this system is 2FGL J2039.6–5620; in the most recent (fourth) *Fermi* LAT source catalogue (4FGL; [Abdollahi et al., 2020](#)), the designation associated with this system is 4FGL J2039.5–5617 (hereafter we will refer to this system using its 4FGL designation). Based on the properties of the optical and X-ray LCs for 4FGL J2039.5–5617, [Romani](#) and [Salvetti et al.](#) concluded that this system is most likely an interacting MSP binary (containing a singular pulsar): the former (optical) LC showed evidence of heating in the companion star, likely due to pulsar emission and/or a pulsar wind, and the latter (X-ray) LC was consistent with an intrabinary shock, where the particle wind originating from the neutron star and the stellar wind originating from its companion, collide. [Salvetti et al.](#) further speculated that this system might be a new ‘redback’ system, i.e., an MSP binary where the pulsar’s non-degenerate companion (with a typical mass of approx. 0.1 – 0.4 $M_{\odot}$; [Roberts, 2013](#)) is being destroyed through the action of the pulsar’s emission or particle wind, or both. This hypothesis, however, could only be accepted if it could be shown that the neutron star was indeed an MSP; i.e., through the detection of pulsed emission from this system.

The detection of γ -ray pulsations (accomplished via a targeted pulsation search conducted using the *Einstein@Home* distributed volunteer computing system), with a period 2.65 ms by [Clark et al. \(2020\)](#) confirmed that 4FGL J2039.5–5617 is, indeed, an MSP binary. Via optical LC modelling, these authors confirmed the classification of this system as a redback

¹Note that pulsations were only detected later for this source; i.e., initially, there was no pulsar associated with this system.

Table 9.1: Single-band-only GOF values and parameter estimates for PSR J2039–5617. Higher values for χ^2_ν indicate decreased GOF.

Fit	α ($^\circ$)	ζ ($^\circ$)	ϕ_μ	β ($^\circ$)	χ^2_ν
Cone	58^{+3}_{-2}	20^{+3}_{-1}	−0.13	-38^{+1}_{-1}	1.68
OG	71^{+2}_{-2}	29^{+3}_{-2}	−0.09	-42^{+5}_{-4}	6.10
TPC	53^{+10}_{-5}	54^{+4}_{-17}	−0.09	1^{+9}_{-25}	5.74

binary, and constrained the neutron star and companion masses to be $1.1 - 1.6 M_\odot$ and $0.15 - 0.22 M_\odot$, respectively. In addition to these constraints, they found a value of $i \gtrsim 60^\circ$ for the binary inclination angle. Since the period of the pulsar was now known, the detection of radio pulsations occurred soon after, as reported by [Corongiu et al. \(2020\)](#), who detected pulsations at 1.4 GHz and 3.1 GHz in the data they gathered (via observations using the Parkes radio telescope), and at 0.7 GHz in archival data (from 2015). With both radio and γ -ray pulsar LCs available (upon this detection), it was now possible to constrain the viewing geometry of PSR J2039–5617 via joint, dual-band LC fitting. Since, under the assumption that the MSP’s short period is due to it being spun-up through the accretion of matter from its companion star, the spin axis of the pulsar is likely roughly aligned with the orbital axis of the binary system. This means that the constraints such LC fitting would yield for the observer angle (ζ) associated with the pulsar, could serve as independent, indirect constraints on the orbital inclination angle of the binary system (i). To this end, we were asked to leverage the pipeline we had developed for the fits described in Chapter 10, and derive constraints on the viewing geometry of PSR J2039–5617.

9.2 The Single- and Dual-band Fits

The observed radio and γ -ray LCs used for these fits had, respectively, $n_r = 64$ and $n_\gamma = 30$ equally-spaced phase bins, and the scaled fluxes were $\Phi_r \simeq 34$ and $\Phi_\gamma \simeq 95$; this yields $\log(\Phi_r) \simeq 1.98$, $\log(\Phi_\gamma) \simeq 1.53$ and $\log(\lambda_{r\gamma}) \simeq 0.44$ for PSR J2039–5617, which indicates that the radio and γ -ray LC uncertainties are only moderately disparate. Therefore, each model LC pair consisted of a radio LC with 64 phase bins and a γ -ray LC with 30 phase bins. As described in Section 7.3, the parameter space of the OG+Cone and TPC+Cone dual-band models was explored at 1° resolution in α and ζ (with $\alpha, \zeta \in (0^\circ, 90^\circ)$), and at a resolution equivalent to width of 2 radio LC bins in $\phi_\mu \in [0, 1]$. In total, upon excluding LC pair where either single-band model predicts no pulsar-associated emission, the OG+Cone and TPC+Cone fits comprised a total of 259 200 model LC pairs each.

Alongside the shared parameters α , ζ and ϕ_μ , each model LC pair is also associated with an optimal radio and optimal γ -ray amplitude (a_r^* and a_γ^* ; just as in Section 7.3). Therefore, each model LC pair is associated with a total of five model parameters. This means that the joint fits have $\nu_c = (64 + 30) - 5 = 89$ degrees of freedom, and the single-band-only fits have $\nu_r = 64 - 4 = 60$ and $\nu_\gamma = 30 - 4 = 26$ degrees of freedom, respectively.

The constraints on α and ζ were derived by applying the Monte Carlo described in Section 7.4.2. Instead of using $\delta_{2\sigma, \text{sfs}}$, however, these constraints were derived using the equivalent estimate for the 3σ confidence level (i.e., $\delta_{3\sigma, \text{sfs}}$) for the distribution of $\Delta\Psi_{\Phi, c}^2(M_c^{\text{sfs}})$.

9.3 Results

9.3.1 Single-band-only Fits

Table 9.1 lists the single-band-only parameter estimates, obtained by minimising χ^2 (or, equivalently χ^2_ν), which are indeed non-colocated in the models’ shared parameter space (when comparing the radio-only estimate to each of the γ -ray-only estimates). The radio- and γ -ray-only best-fit LCs, along with their respective γ -ray and radio counterparts, are plotted in Figure 9.1. The radio-only best-fit LC is accompanied by a background-only OG model LC, for which $\chi^2_{\nu, \gamma} = \Phi_\gamma^2/\nu_\gamma = 45.00$, and a TPC model LC that is a worse fit than a background-only LC ($\chi^2_{\nu, \gamma} = 80.40 > \Phi_\gamma^2/\nu_\gamma$). The γ -ray-only best-fit LC for the OG model is accompanied by a comparatively good radio LC ($\chi^2_{\nu, r} = 54.92 \simeq 0.36 \times \Phi_r^2/\nu_r$),




Figure 9.1: The radio- and γ -ray-only best-fit LCs for the Cone (red; bottom panel), OG (light blue; top panel), and TPC (dark blue; top panel) models, along with their respective γ -ray and radio counterparts. The observed radio (bottom panel) and γ -ray (top panel) LCs (both grey with black error bars) are those observed at 1.4 GHz and at > 0.1 GeV, respectively. The uniform radio LC error is taken to be the standard deviation of the intensities in the off-peak region (the grey bands).

Table 9.2: Dual-band GOF values and parameter estimates for PSR J2039-5617. Higher values for $\Psi_{\Phi,c}^2$, $\Psi_{\Phi,r}^2$, and $\Psi_{\Phi,\gamma}^2$ indicate increased GOF, while higher values for $[\chi_c^2]_\nu$, $\chi_{\nu,r}^2$, and $\chi_{\nu,\gamma}^2$ indicate decreased GOF.

Joint Fit (+Cone)	α ($^\circ$)	ζ ($^\circ$)	ϕ_μ	β ($^\circ$)	$\Psi_{\Phi,c}^2$	$[\chi_c^2]_\nu$	$\Psi_{\Phi,r}^2$	$\chi_{\nu,r}^2$	$\Psi_{\Phi,\gamma}^2$	$\chi_{\nu,\gamma}^2$
OG	70_{-7}^{+3}	31_{-5}^{+12}	-0.13	-39_{-8}^{+18}	0.883	6.02	0.967	5.01	0.799	9.02
OG ($\zeta > 50^\circ$)	36_{-16}^{+5}	67_{-2}^{+4}	-0.13	31_{-6}^{+21}	0.834	11.12	0.924	11.50	0.743	11.52
TPC (Main)	61_{-1}^{+2}	42_{-4}^{+2}	-0.09	-19_{-5}^{+2}	0.846	16.50	0.859	21.21	0.833	7.53
TPC (Alt.)	29_{-11}^{+12}	67_{-6}^{+3}	-0.13	38_{-16}^{+13}	0.841	10.46	0.929	10.69	0.753	11.12

which resembles the shape of the observed radio LC despite the radio peak occurring too late in phase. The best-fit LC for the TPC model is accompanied by a somewhat worse radio fit ($\chi_{\nu,r}^2 = 113.28 \simeq 0.75 \times \Phi_r^2/\nu_r$), with two radio peaks instead of one. The comparatively better performance of the OG best fit's counterpart is consistent with the greater agreement between the estimates it and the Cone best fit yield for $\beta = \zeta - \alpha$, the parameter that typically governs peak multiplicity in the Cone model.

9.3.2 Dual-band Fits

The best-fit LC pair for the OG+Cone fit, plotted in green in Figure 9.2, has a joint GOF value of $\Psi_{\Phi,c}^2 = 0.883$ and a reduced χ^2 value of $[\chi_c^2]_\nu = 6.02$, as listed in Table 9.2. The emission maps associated with this LC pair are shown in Figure 9.3 (top row). The radio-only GOF value for this fit is $\Psi_{\Phi,r}^2 = 0.967$ ($\chi_{\nu,r}^2 = 5.01$), and the γ -ray-only GOF value is $\Psi_{\Phi,\gamma}^2 = 0.799$ ($\chi_{\nu,\gamma}^2 = 9.02$). These LCs, coupled with the 3σ confidence regions shown in white in Figure 9.4a, correspond to a pulsar parameter estimate of $(\alpha, \zeta) = ((70_{-7}^{+3})^\circ, (31_{-5}^{+12})^\circ)$.

Comparing this LC pair to those that correspond to the relevant single-band-only fits puts this joint fit, and the compromise it represents, into its proper context: at the cost of some GOF in the radio band (as compared to the radio-only




Figure 9.2: The best-fit LC pairs associated with the OG+Cone and TPC+Cone dual-band models (solid green and solid brown LCs, respectively). For the TPC+Cone model, the best fit included in the alternative contour identified in Figure 9.4b is also plotted (dashed brown LC). The observed LCs are the same as those in Figure 9.1.

fit; $\chi^2_{\nu,r}$ increases from 1.68 to 5.01), the degree to which the γ -ray LC is reproduced is increased substantially; from a background-only LC to an LC that is very similar to the γ -ray-only best-fit LC ($\chi^2_{\nu,\gamma}$ decreases from 45.00 to 9.02). Or, equivalently (as compared to the γ -ray-only fit), the radio peak's phase is recovered in the joint fit ($\chi^2_{\nu,r}$ decreases from 54.92 to 5.01) at the cost of some GOF in the γ -ray band ($\chi^2_{\nu,\gamma}$ increases from 6.10 to 9.02).

The TPC+Cone fit yielded a pair of confidence regions in (α, ζ) -space: one for which $\beta < 0^\circ$ (the main fit) and one for which $\beta > 0^\circ$. The solid brown LCs in Figure 9.2 are the best fit included in the former of these regions (enclosed in a solid blue bounding box in Figure 9.4b), and correspond to a parameter estimate of $(\alpha, \zeta) = ((61^{+2}_{-1})^\circ, (42^{+2}_{-4})^\circ)$. Similarly, the dashed brown LCs in Figure 9.2 are the best fit included in the latter of these regions (enclosed in a dashed blue bounding box in Figure 9.4b), and correspond to a parameter estimate of $(\alpha, \zeta) = ((29^{+12}_{-11})^\circ, (67^{+3}_{-6})^\circ)$. The emission maps associated with these two fits are shown in the bottom two rows of Figure 9.3, and the associated GOF values are listed in Table 9.2. The overall best fit for this pair of models (TPC+Cone) lies inside the $\beta < 0^\circ$ contour.

Comparing the OG+Cone and TPC+Cone best fits again demonstrates the effect single-band best-fit non-colocation has on the GOF for the joint fits. For the single-band only fits, the TPC model outperforms the OG model ($\chi^2_{\nu,\gamma} = 5.74$ vs. 6.10), but for the joint fits the OG+Cone model outperforms the TPC+Cone model ($\Psi^2_{\Phi,c} = 0.883$ vs. 0.846). Understood in terms of compromise, the greater degree of non-colocation between the TPC and Cone estimates necessitates a more costly compromise than that made in the OG+Cone fit.

Since an orbital inclination angle of $i < 50^\circ$ (and hence $\zeta < 50^\circ$) implies an unrealistically high pulsar mass of $M_{\text{psr}} > 2.4 M_\odot$, a second OG+Cone fit was conducted, wherein only model LC pairs for which $\zeta > 50^\circ$ were considered. The best-fit LC pair thus yielded, which corresponds to a parameter estimate of $(\alpha, \zeta) = ((36^{+5}_{-16})^\circ, (67^{+4}_{-2})^\circ)$, is plotted in orange in Figure 9.2, and the associated confidence region (again derived using a Monte Carlo algorithm) is shown in yellow in Figure 9.4a (enclosed in an orange bounding box). The emission maps associated with this fit are shown in the second row of Figure 9.3, and the associated GOF values are listed in Table 9.2. While this fit is worse than the overall OG+Cone best fit, as indicated by the lower associated $\Psi^2_{\Phi,c}$ value, it is still a good fit to the overall structure of the observed LCs since its GOF value is (substantially) larger than 0.




Figure 9.3: Model radio and γ -ray emission maps (relative intensity as function of rotational phase ϕ and observer angle ζ) for the four joint fits reported in Table 9.2. The grey regions indicate where the models predict no pulsar-associated emission, i.e., where only background emission will be observed. Each pair (row) of emission maps is associated with a single pulsar inclination angle α , and the LC pairs plotted in Figure 9.2 are each associated with a constant- ζ cut through the appropriate pair of emission maps (indicated by the horizontal yellow lines). For each cut, the point in model phase that corresponds to 0 in observational phase is indicated by a small vertical tick mark (also yellow).




Figure 9.4: Maps of joint goodness of fit as function of α and ζ (marginalised over ϕ_μ) for the OG+Cone and TPC+Cone dual-band models (panels (a) and (b), respectively) as characterised using $\Psi_{\Phi,c}^2$. Only model LC pairs for which $\Psi_{\Phi,c}^2 > 0$ and where both wavebands predict pulsar-associated flux, are plotted. In both panels, the white contour demarcates the Monte Carlo-derived 3σ confidence region(s), and each region’s best fit is indicated with a blue dot. The smaller region for the OG+Cone model is considered an extension of the larger region, while the TPC+Cone regions are treated as separate regions. The parameter constraints listed in Table 9.2 are derived from the blue bounding boxes that enclose the various regions.

9.3.3 Implications for the Orbital Inclination

Since $\zeta > 50^\circ$ from the optical fits conducted in Clark et al. (2020), we estimate the observer angle to be $\zeta = (67^{+4}_{-2})^\circ$ (from the second OG+Cone fit) or $\zeta = (67^{+3}_{-6})^\circ$ (from the alternative TPC+Cone contour) from these fits. Upon again assuming that $\zeta \equiv i$, these estimates correspond to mass ranges of $1.17 M_\odot < M_{\text{psr}} < 1.46 M_\odot$ and $1.19 M_\odot < M_{\text{psr}} < 1.63 M_\odot$ for the OG+Cone and TPC+Cone models, respectively. These estimates for the orbital inclination angle and pulsar mass are qualitatively consistent with those found Clark et al., since a fairly low pulsar mass is preferred in both cases.

9.4 Discussion

The account we gave for these fits was intentionally conservative in the sense that the statistics introduced in this work were only invoked where absolutely necessary. This is particularly evident in Section 9.3.1, where the single-band-only fits are described using only $\chi^2_{\nu,r}$ and $\chi^2_{\nu,\gamma}$ (i.e., without invoking $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$). Since it is still helpful to consider the single-band GOF in light of the perturbations the radio and γ -ray LCs represent, however, the scaled fluxes were used in this section to establish a meaningful context for the reduced χ^2 values. For example, by adding that $\chi^2_{\nu,r} \simeq 0.36 \times \Phi_r^2 / \nu_r$ for the radio LC that accompanies the OG γ -ray-only best-fit LC, we establish that the OG best-fit LC is accompanied by a radio LC that yields a marginally favourable radio-only fit (i.e., that still outperforms a background-equivalent radio LC).

The discussion of the joint fits (in Section 9.3.2, in concert with Table 9.2), on the other hand, features the statistics introduced in this work much more prominently, since these statistics express GOF in more appropriate units; units wherein the radio- and γ -ray-band GOF can be sensibly compared. The reduced χ^2 values are not disregarded, though, since they express DFO more effectively. For example, in the first paragraph of Section 9.3.2 (and in Table 9.2), each instance of a GOF value (i.e., of $\Psi_{\Phi,r}^2$, $\Psi_{\Phi,\gamma}^2$ or $\Psi_{\Phi,c}^2$) is accompanied by the accompanying DFO value (i.e., by a value for the appropriate reduced χ^2 statistic). Note that the second and fourth paragraphs in Section 9.3.2 were included specifically to elucidate the effect of the single-band best-fit non-colocation on the joint fits (i.e., that it necessitates compromise; since the first of our prospective papers is still forthcoming), and to demonstrate that the compromises that accompany the use of the

$\Psi_{\Phi,c}^2$ statistic are arguably fair.

In summary, then, this account demonstrates that the complementarity of the concepts of DFO and GOF (as outlined in Section 4.1) extends to the statistics that express them most naturally; i.e., to the χ^2 and Ψ_{Φ}^2 families of statistics. It also demonstrates that both types of statistic is necessary for a comprehensive account of a joint, dual-band fit aimed at deriving parameter estimates, since the former encodes how useful the model is while the latter encodes how far in error the model is.

Chapter 10

Dual-band Fits for 23 *Fermi* LAT 2PC Pulsars

As a first real-world test of the $\Psi_{\Phi,c}^2$ statistic, we derive parameter estimates via joint, dual-band LC fitting (using the OG+Cone and TPC+Cone conglomerated dual-band models) for 23 radio-loud *Fermi* LAT pulsars. To evaluate the effectiveness with which $\Psi_{\Phi,c}^2$ eliminates radio dominance in these fits (using the tools introduced in Chapter 6), we also identify the relevant single-band best fits and the joint best fits that χ_c^2 would identify. The main result that this evaluation yields (alongside the parameter estimates themselves), is a set of distributions similar to those plotted in Figures 6.2 and 8.3.

As a secondary step, the estimates obtained using $\Psi_{\Phi,c}^2$ are compared to those obtained by Johnson et al. (2014) and Pierbattista et al. (2015), who used models similar to the geometric models used here, but attempted to eliminate radio dominance by other means (see Section 3.2.3). This comparison is limited to the estimates themselves (i.e., the parameter estimates for α and ζ), however, since we do not have access to the specific model outputs that either of these groups used to obtain their estimates. While the implementations they used for the OG and TPC models yield very similar outputs to the implementations we use, these are not identical. Also, these authors opted to introduce a core radio beam for some pulsars (e.g., PSR J0437–4715), which has some effect on the degree to which our fits agree with theirs where the dual-band estimates for $\beta = \zeta - \alpha$ are concerned (our fits will tend to yield larger values for $|\beta|$ for the pulsars where the other authors' modelling favoured a strong core component). It should also be noted here that these authors used Poisson likelihood as the γ -ray GOF statistic, and combined it with the radio χ^2 value by first converting the latter to a likelihood as well (instead of using χ_c^2 to characterise joint GOF; leveraging the radio uncertainty scaling they imposed to mitigate radio dominance). Since we will not be conducting a comparison of the GOF each of the models achieved (instead focusing exclusively on the parameter estimates they yield), we do not expect the difference in our approaches to substantially affect the comparisons we conduct in this chapter.

As a baseline reference, we also include the parameter estimates derived using χ_c^2 in this comparison. Additionally, as a last sanity check we derived estimates for these pulsars by eye and included these estimates in the comparison.

10.1 Pulsar Selection

The set of pulsars used in this test was not specifically constructed for it, but is rather a subset of a larger set that was constructed in multiple stages to serve the needs of multiple, related studies being conducted in our research group. The primary criteria that the ultimate set of 23 pulsars aims to satisfy (within the context of this test) are:

- Every included pulsar must have previously been fit by either Johnson et al. (2014) or Pierbattista et al. (2015).
- The values for $\log(\lambda_{r\gamma})$ for the included pulsars must be roughly evenly distributed over a maximally broad range.
- The observed radio and γ -ray LCs for the included pulsars must be roughly compatible (in terms of shape) with the geometric models we use (i.e., the observed γ -ray LC with the OG and TPC models and the observed radio LC with the Cone model). (This ensures that both single-band models are at least somewhat useful, and hence that a dual-band fit is indeed warranted within the context of the discussion in Section 3.2.2.)

The detection of the six pulsars marked *a* in Table 10.1 by the *Fermi* LAT was reported in Weltevrede et al. (2010), who derived estimates for α and $\beta = \zeta - \alpha$ from their radio data. Subsequently, these pulsars formed the basis for the

Table 10.1: Simulation parameters, LC characteristics and previously derived parameter estimates for the 23 selected pulsars. The estimates for PSR J0030+0451, PSR J0437–4715, PSR J1231–1411, PSR J1614–2230 and PSR J2241–5236 were reported by [Johnson et al. \(2014\)](#). The rest were reported by [Pierbattista et al. \(2015\)](#).

PSR	P (s)	$\dot{P}$ (ss ⁻¹)	Prev. Constraints:					OG+Cone		TPC+Cone	
			$\log(\lambda_{r\gamma})$	$\log(\Phi_r)$	$\log(\Phi_\gamma)$	ε_r (%)	$\varepsilon_{r,sfs}$ (%)	α (°)	ζ (°)	α (°)	ζ (°)
J0437–4715 ^{b,c}	0.0058	1.37E-20	2.41	3.75	1.34	0.04	9.97	76 ⁺¹ ₋₁	46 ⁺¹ ₋₁	35 ⁺¹ ₋₁	64 ⁺¹ ₋₁
J0908–4913	0.1068	1.51E-14	2.36	3.67	1.31	0.04	8.03	75 ⁺² ₋₂	90 ⁺² ₋₂	70 ⁺² ₋₂	90 ⁺² ₋₂
J0742–2822 ^a	0.1668	1.68E-14	2.16	3.27	1.11	0.13	18.57	76 ⁺² ₋₃	86 ⁺³ ₋₂	62 ⁺² ₋₂	77 ⁺² ₋₂
J0248+6021	0.2171	5.51E-14	1.72	3.13	1.41	0.24	12.68	56 ⁺² ₋₂	65 ⁺² ₋₂	60 ⁺² ₋₂	53 ⁺² ₋₂
J2241–5236 ^c	0.0022	5.09E-21	1.67	3.10	1.43	0.19	8.77	19 ⁺¹ ₋₁	75 ⁺¹ ₋₁	20 ⁺¹ ₋₁	76 ⁺¹ ₋₁
J0835–4510	0.0893	1.25E-13	1.53	4.33	2.80	0.01	0.30	71 ⁺² ₋₂	83 ⁺² ₋₂	45 ⁺² ₋₂	69 ⁺² ₋₂
J0659+1414 ^a	0.3849	5.49E-14	1.53	2.94	1.41	0.32	10.59	78 ⁺² ₋₂	73 ⁺² ₋₂	30 ⁺² ₋₂	32 ⁺² ₋₂
J1048–5832	0.1237	9.61E-14	1.50	3.28	1.79	0.13	4.21	87 ⁺² ₋₂	76 ⁺² ₋₂	62 ⁺² ₋₂	74 ⁺² ₋₂
J0631+1036 ^a	0.2878	1.05E-13	1.31	2.65	1.34	0.39	7.94	87 ⁺² ₋₂	72 ⁺² ₋₂	48 ⁺² ₋₂	67 ⁺² ₋₂
J0030+0451 ^{b,c}	0.0049	1.00E-20	1.27	2.89	1.62	0.55	10.17	88 ⁺¹ ₋₂	68 ⁺¹ ₋₁	74 ⁺² ₋₂	55 ⁺³ ₋₁
J1614–2230 ^{b,c}	0.0032	3.89E-21	1.26	2.55	1.28	0.70	12.83	64 ⁺⁸ ₋₂₀	88 ⁺² ₋₅	80 ⁺⁸ ₋₁₆	80 ⁺⁶ ₋₄
J1513–5908 ^b	0.1513	1.53E-12	0.92	2.04	1.12	4.01	33.08	60 ⁺² ₋₂	59 ⁺² ₋₂	50 ⁺² ₋₂	54 ⁺² ₋₂
J1718–3825 ^a	0.0747	1.32E-14	0.90	2.37	1.48	1.04	8.22	80 ⁺² ₋₂	55 ⁺² ₋₂	45 ⁺² ₋₂	65 ⁺² ₋₂
J1709–4429	0.1025	9.30E-14	0.84	3.15	2.31	0.24	1.65	73 ⁺² ₋₂	72 ⁺² ₋₂	42 ⁺² ₋₂	63 ⁺² ₋₂
J1420–6048 ^a	0.0682	8.32E-14	0.63	2.24	1.61	2.06	8.68	55 ⁺² ₋₂	57 ⁺² ₋₂	52 ⁺² ₋₂	53 ⁺² ₋₂
J1231–1411 ^{b,c}	0.0037	8.19E-21	0.55	2.26	1.71	2.14	7.67	88 ⁺¹ ₋₁	67 ⁺¹ ₋₁	26 ⁺³ ₋₄	69 ⁺¹ ₋₁
J1833–1034 ^b	0.0619	2.02E-13	0.50	1.92	1.42	1.77	5.65	65 ⁺² ₋₂	87 ⁺² ₋₂	55 ⁺² ₋₂	75 ⁺² ₋₂
J1509–5850 ^a	0.0889	9.17E-15	0.11	1.68	1.57	7.55	9.80	85 ⁺² ₋₂	76 ⁺² ₋₂	46 ⁺² ₋₂	66 ⁺² ₋₂
J1747–2958	0.0988	6.13E-14	0.02	1.67	1.64	4.77	5.05	87 ⁺² ₋₂	79 ⁺² ₋₂	56 ⁺² ₋₂	77 ⁺² ₋₂
J0205+6449 ^b	0.0657	1.94E-13	0.02	1.60	1.58	8.11	8.49	73 ⁺² ₋₂	90 ⁺² ₋₂	75 ⁺² ₋₂	86 ⁺² ₋₂
J2229+6114 ^b	0.0516	7.83E-14	–0.21	1.74	1.95	7.47	4.60	75 ⁺² ₋₂	55 ⁺² ₋₂	42 ⁺² ₋₂	55 ⁺² ₋₂
J1124–5916 ^b	0.1355	7.53E-13	–0.30	1.23	1.54	13.36	6.64	83 ⁺² ₋₂	88 ⁺² ₋₂	84 ⁺² ₋₂	89 ⁺² ₋₂
J2021+3651	0.1037	9.57E-14	–0.34	1.72	2.06	6.84	3.13	68 ⁺² ₋₂	90 ⁺² ₋₂	73 ⁺² ₋₂	83 ⁺² ₋₂

^aFirst set pulsar.

^bSecond set pulsar.

^cMillisecond pulsar.

study that preceded this one, where parameter estimates were derived for them by eye, since they had (at the time) broad, roughly single-peaked γ -ray LCs (which rendered them compatible with the OG and TPC models) and relatively simple single- or double-peaked radio LCs (which rendered them compatible with the Cone model). The fits in that study were conducted by eye out of necessity, since the radio fit dominates for the majority of these pulsars.

The nine pulsars marked *b* in Table 10.1 were selected based on their spin-down luminosities, as part of an on-going study aimed at investigating the possible relationship between conductivity and spin-down luminosity uncovered by Brambilla et al. (2015) for the FIDO family of pulsar models (see Section 1.5 and Section 11.1). They were selected such that the entire set of pulsars (i.e., the set constituted by their addition to the already established set of six pulsars) covers a wide range of spin-down luminosities.

The last set of (eight) pulsars (the unmarked pulsars in Table 10.1) were selected specifically for the test conducted in this chapter based on the value of $\log(\lambda_{\gamma})$ for each of them. PSR J2021+3651 ($\log(\lambda_{\gamma}) = -0.34$) and PSR J0437-4715 ($\log(\lambda_{\gamma}) = 2.41$) were selected to maximise the range of $\log(\lambda_{\gamma})$ values included. The former of these is of particular interest since it is an example of a pulsar for which the γ -ray GOF dominates. In all, four of the pulsars fit by either Johnson et al. or Pierbattista et al. had values of $\log(\lambda_{\gamma}) < 1$ (thus signifying a potential for γ -ray dominance). Two of these pulsars, PSR J2229+6114 ($\log(\lambda_{\gamma}) = -0.21$) and PSR J1124-5923 ($\log(\lambda_{\gamma}) = -0.30$), had already been included as part of the second set of pulsars, while the third (PSR J1741-2054; $\log(\lambda_{\gamma}) = -0.23$), had relatively poor radio statistics, and was consequently not included. Similarly, PSR J1747-2958 ($\log(\lambda_{\gamma}) = 0.02$) was included to probe what happens when $\log(\lambda_{\gamma}) \approx 0$. Since the agreement between the estimates obtained using $\Psi_{\Phi,c}^2$ and χ_c^2 is expected to increase as $\log(\lambda_{\gamma}) \rightarrow 0$, the fits for this pulsar, along with those for PSR J0205+6449 ($\log(\lambda_{\gamma}) = 0.02$; included as part of the second set), have the greatest potential to demonstrate this convergence in a real-world context. To probe the effectiveness of $\Psi_{\Phi,c}^2$ at high $\log(\lambda_{\gamma})$ values, PSR J0908-4913 ($\log(\lambda_{\gamma}) = 2.36$) was included. The Vela pulsar (PSR J0835-4510; $\log(\lambda_{\gamma}) = 1.53$) was included based on its comparatively good γ -ray statistics. Similarly, PSR J1048-5832 ($\log(\lambda_{\gamma}) = 1.50$) and PSR J1709-4429 ($\log(\lambda_{\gamma}) = 0.84$) had the best γ -ray statistics in their respective $\log(\lambda_{\gamma})$ neighbourhoods, thus warranting their inclusion.

10.2 Pulsar Data

The LCs used in this test, both radio and γ -ray, are those published in the 2PC for the 23 pulsars listed in Table 10.1.¹ For the γ -ray LCs, the uncertainties used here are those reported in the 2PC, while for the radio LCs an uncertainty ε_r (that is uniform in phase; listed in Table 10.1) is derived from the off-peak variability in the observed radio LC. For each pulsar, the off-peak region (identified by eye) used to derive ε_r is indicated on the associated plot in Appendix A. The radio uncertainties used here necessarily differ from those used by Johnson et al. and Pierbattista et al. since both of their approaches to mitigating the radio dominance rely on an intentionally different prescription for these uncertainties.

While more up-to-date γ -ray LCs could be used, this would weaken the power of the direct comparison to the estimates that were derived previously using either the same dataset as the 2PC (Johnson et al.) or the LCs published therein (Pierbattista et al.). Both other groups used the radio LCs that were published in the 2PC to derive their estimates on α and ζ .

It should be noted here that the reported γ -ray uncertainties are Poissonian, and not Gaussian as is assumed when χ_{γ}^2 is calculated. Since the main focus of this test is the extent to which single-band dominance is eliminated for the parameter point estimates (and not the extent to which it is eliminated for the parameter constraints), we have opted to treat these uncertainties as Gaussian for this test. While this simplification likely reduces the nominal accuracy of the parameter estimates we derive here slightly, it does not affect the effectiveness of this test. Of course, as the γ -ray statistics improves this simplification will become increasingly accurate since the γ -ray uncertainties will become Gaussian in nature. For the brighter pulsars included in this test, this is arguably already the case (that the Gaussian assumption holds water for the γ -ray uncertainties).

¹Specifically, those contained in version 4 of the FITS files that accompanied its publication; available at https://fermi.gsfc.nasa.gov/ssc/data/access/lat/2nd_PSR_catalog/2PC_auxiliary_files_v04.tgz

LC Characteristics Table 10.1 summarises some of the characteristics of these LCs that are most relevant to the fits conducted here. To reflect our focus on radio dominance, the entries in Table 10.1 (and all the other tables in this chapter) are sorted according to the value of $\log(\lambda_{r\gamma})$ in descending order. Consequently, the pulsars that are more prone to radio dominance are listed higher up in these tables. Interestingly, this ordering results in a roughly decreasing trend in the value of $\log(\Phi_r)$, with the pulsars associated with the most precise radio observations lying closer to the top. As a reference, the effective scaled radio relative uncertainty $\varepsilon_{r,\text{sfs}}$ (calculated using Eq. (5.50)) that would yield the same best fits as $\Psi_{\Phi,c}^2$ is also included in Table 10.1. While this uncertainty is never used (neither when identifying the best fits nor when deriving the parameter constraints), it demonstrates exactly how severe the uncertainty disparity is for each pulsar in more familiar terms than $\log(\lambda_{r\gamma})$ does. Notice that $\log(\lambda_{r\gamma}) > 0 \Rightarrow \varepsilon_{r,\text{sfs}} < \varepsilon_r$, as might be expected from Eq. (5.50) since $\log(\lambda_{r\gamma}) > 0 \Rightarrow \lambda_{r\gamma} > 1$.

10.3 By-Eye Fits

As a sanity check, and to aid in the comparison of our estimates with those derived by either Johnson et al. (2014) or Pierbattista et al. (2015), we also derive parameter estimates by eye (some of these fits were previously reported in Bezuidenhout et al., 2017). While the by-eye fits are by no means robust, and our aim is not to reproduce them or the parameter estimates they yield, they provide a useful point of comparison with regard to radio dominance. To ensure that these fits are not biased by those obtained using other methods, a somewhat blind strategy was devised wherein the person conducting the fit does not know which model parameters are associated with each fit they consider. Since probing the parameter space of the models at the chosen resolution of 1° in both α and ζ leads to a relatively large number of LC pairs (as many as 8000 per dual-band model for some pulsars), the process was split into three phases.

During the first phase, the parameter space is probed at a 5° resolution, which yields a considerably smaller set of LC pairs (compared to the full set of model LC pairs) that act as representatives of their respective neighbourhoods in the model's parameter space. This resolution was chosen because the LCs produced by the various single-band models tend to only significantly change with changes in either α or ζ that are larger than $\sim 10^\circ$. These LC pairs are presented to the person conducting the fit in a random order, and they are either accepted or rejected as promising candidates based on how well they fit the observed LC pair (jointly; in the person conducting the fit's opinion).

The candidate LC pairs identified in the first phase are used as starting points for the exploration of the parameter space at 1° resolution during the second phase. During each of an indeterminate number of iterations, the set of unevaluated LC pairs associated with parameter combinations that lie either above, below, to the left or to the right of an accepted LC pair (in (α, ζ) -space), is shown to the person conducting the fit in a random order.² Each of these LC pairs is either accepted or rejected based on how well they fit the observed LC pair. Once no accepted parameter combinations have unevaluated neighbours, the second phase ends. The set of parameter combinations that correspond to the accepted LC pairs constitute the so-called inclusion contours associated with the by-eye fit in question.

During the third phase of the by-eye fitting procedure, the various connected regions in the parameter space that make up the inclusion contours are freely explored, and the region that contains the model's overall best by-eye fit is identified. In the interest of simplicity, and to reflect the subjective nature of the by-eye evaluations, the parameters associated with the best fit thus identified (if any single LC is identified) are disregarded, and the best-fit parameters associated with the fit are taken to be the midpoint of the smallest rectangular region that contains the inclusion contour to which the model's best fit belongs. The associated parameter constraints are also derived from the extent of this rectangular region (i.e., similar to the constraints for the other fits), and the derived uncertainties are thus symmetric.

Note that, while ϕ_μ was considered a free parameter for the by-eye fits, the best-fit ϕ_μ values were not recorded for the by-eye fits.

²Note that even if an LC pair was evaluated during the first phase, it is considered unevaluated during the second phase (to avoid unwanted holes in the inclusion contours).

10.4 Results

Fitting the observed LC pair associated with each pulsar yielded a total of 9 parameter estimates: one radio-only estimate, two γ -ray-only estimates (one for the OG model and one for the TPC model) and three joint estimates for each of the dual-band models (one derived by eye, one derived using χ_c^2 and one derived using $\Psi_{\Phi,c}^2$). For each pulsar, the best-fit LC pairs yielded by the joint fits conducted using χ_c^2 and $\Psi_{\Phi,c}^2$ are plotted in a figure in Appendix A (a total of four LC pairs per pulsar). While a summary by-eye evaluation of the effectiveness of the $\Psi_{\Phi,c}^2$ statistic based on these LC pairs is conducted below (and is definitely instructive), we have opted to rather leverage the tools developed in Chapter 6 to more rigorously evaluate the $\Psi_{\Phi,c}^2$ statistic's effectiveness (and therefore to not include these LC pairs here, although the best-fit α and ζ values these LCs correspond to are listed in Tables 10.3 and 10.4), since the latter are arguably more objective. See Appendix A for a more detailed description of these LC plots.

Alongside the best-fit LC pairs plotted in Appendix A for each pulsar, the (α, ζ) -space maps of $\Psi_{\Phi,c}^2$ associated with each of the joint fits is plotted in Appendix B. The parameter constraints listed in Tables 10.3 and 10.4 for the joint fits conducted using $\Psi_{\Phi,c}^2$ were derived from these maps (as described in Section 7.4). Also included on these maps are the contours for the fits conducted using χ_c^2 (from which the associated parameter constraints are derived). See Appendix B for a more detailed description of these maps.

Note that since [Pierbattista et al.](#) only reported estimates of α and ζ (and not ϕ_μ), these two parameters are considered to be the parameters of interest. While the best-fit values for ϕ_μ are listed for the various single- and dual-band fits conducted here, they are only listed for reference.

10.4.1 Single-band-only Fits

Table 10.2 lists the DFO (via $\chi_{\nu,r}^2$ and $\chi_{\nu,\gamma}^2$), GOF (via $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$) and parameter estimates associated with each of the three single-band-only fits conducted for each pulsar. Just as in Table 10.1, the pulsars are listed in order of descending $\log(\lambda_{r\gamma})$ (i.e., pulsars listed earlier are more prone to radio dominance). For each single-band model, the statistic values and parameter estimates it yielded are listed in a dedicated block of the table (labelled using the model's name). In each block, the first and second columns list the DFO and GOF values associated with each of the relevant fits, and the third through fifth columns list the estimated values for α , ζ and ϕ_μ (the values for a_r^* and a_γ^* are omitted because they were found analytically in each case). Note that we have assumed a minimum uncertainty for α and ζ of 1° .

As might be expected from the trend for $\log(\Phi_r)$ (which is observed to decrease as $\log(\lambda_{r\gamma})$ decreases; i.e., as one moves down the list of pulsars), the values for $\chi_{\nu,r}^2$ are much greater than 1, and the radio-only parameter estimates more tightly constrained for the pulsars listed earlier in Table 10.2. In fact, for all the pulsars where $\log(\Phi_r) > 2$, the Monte Carlo algorithm yields no radio-only contours, and the parameter uncertainties default to their assumed minimum of $\pm 1^\circ$. The $\Psi_{\Phi,r}^2$ values for these fits reveal, however, that this trend does not correspond to an equally dramatic decrease in GOF as the precision of the radio observations increases (although $\Psi_{\Phi,r}^2$ does tend to decrease somewhat as $\log(\lambda_{r\gamma})$ decreases); rather, the radio-band GOF is fairly high for all 23 pulsars.

Overall, the Cone model fares best of the three single-band models, with a median GOF of 0.90. Of the two γ -ray models, the OG model fares best, with a median GOF of 0.87, compared to 0.67 for the TPC model. Notably, the latter γ -ray model performs worse than the background LC for 4 of the 23 pulsars (since $\Psi_{\Phi,\gamma}^2 < 0$; the $\Psi_{\Phi,\gamma}^2$ values for these pulsars are highlighted in grey). Interestingly, for the Cone and OG models the average GOF is close to the median GOF, but for the TPC model, the average GOF is much lower the median because of these four very bad fits.

10.4.2 Dual-band Fits

Tables 10.3 and 10.4 list the results for the by-eye, χ_c^2 and $\Psi_{\Phi,c}^2$ fits to the OG+Cone and TPC+Cone dual-band models, respectively. For the by-eye fits only the estimates for α and ζ are listed since no GOF statistic was used to conduct them. For the χ_c^2 fits, the joint DFO (via $[\chi_c^2]_\nu$; recall that it is equivalent to χ_c^2), implied joint GOF (via $[\Psi_c^2]_\Phi$; see Section 4.3), radio-band GOF (via $\Psi_{\Phi,r}^2$) and γ -ray-band GOF (via $\Psi_{\Phi,\gamma}^2$) are listed for the best fit obtained for each pulsar, along with the corresponding parameter estimates. For the $\Psi_{\Phi,c}^2$ fits, the joint DFO is omitted and only the joint GOF is listed (via $\Psi_{\Phi,c}^2$; introduced in Section 4.5), alongside the single-band GOF values and the corresponding parameter estimates. For

Table 10.2: Results for the Cone, OG and TPC single-band-only fits.

PSR	$\log(\Phi_r)$	$\log(\Phi_\gamma)$	Cone					OG					TPC				
			$\chi^2_{\nu,r}$	$\Psi^2_{\Phi,r}$	α ($^\circ$)	ζ ($^\circ$)	ϕ_μ	$\chi^2_{\nu,\gamma}$	$\Psi^2_{\Phi,\gamma}$	α ($^\circ$)	ζ ($^\circ$)	ϕ_μ	$\chi^2_{\nu,\gamma}$	$\Psi^2_{\Phi,\gamma}$	α ($^\circ$)	ζ ($^\circ$)	ϕ_μ
J0437-4715	3.75	1.34	4576.07	0.96	28^{+1}_{-1}	69^{+1}_{-1}	0.031	1.01	0.90	61^{+1}_{-3}	32^{+6}_{-1}	0.148	6.75	0.34	54^{+1}_{-1}	45^{+1}_{-1}	0.117
J0908-4913	3.67	1.31	13552.55	0.84	53^{+1}_{-1}	30^{+1}_{-1}	0.000	2.58	0.72	68^{+5}_{-5}	89^{+1}_{-2}	0.039	1.89	0.79	77^{+8}_{-7}	87^{+2}_{-1}	-0.461
J0742-2822	3.27	1.11	302.80	0.98	63^{+1}_{-1}	50^{+1}_{-1}	0.008	1.14	0.68	75^{+14}_{-36}	89^{+1}_{-15}	0.078	1.91	0.47	35^{+12}_{-12}	61^{+4}_{-8}	0.250
J0248+6021	3.13	1.41	700.37	0.90	8^{+1}_{-1}	25^{+1}_{-1}	0.023	1.14	0.92	72^{+1}_{-5}	20^{+2}_{-2}	0.125	1.48	0.89	25^{+11}_{-11}	66^{+3}_{-3}	0.039
J2241-5236	3.10	1.43	273.92	0.96	87^{+1}_{-1}	27^{+1}_{-1}	0.055	4.28	0.73	28^{+9}_{-13}	84^{+4}_{-7}	-0.289	2.76	0.83	73^{+4}_{-1}	46^{+4}_{-3}	-0.336
J0835-4510	4.33	2.80	73529.55	0.96	44^{+1}_{-1}	71^{+1}_{-1}	0.008	311.31	0.92	82^{+1}_{-1}	69^{+1}_{-1}	0.000	10815.03	-1.65	68^{+1}_{-1}	80^{+1}_{-1}	0.023
J0659+1414	2.94	1.41	25.93	0.99	13^{+1}_{-1}	26^{+1}_{-1}	0.000	0.85	0.94	75^{+2}_{-4}	24^{+8}_{-2}	-0.047	6.84	0.53	33^{+5}_{-3}	63^{+1}_{-2}	-0.117
J1048-5832	3.28	1.79	639.71	0.96	48^{+1}_{-1}	36^{+1}_{-1}	-0.008	4.86	0.87	82^{+1}_{-1}	68^{+1}_{-1}	-0.016	49.79	-0.29	68^{+1}_{-1}	80^{+1}_{-1}	-0.008
J0631+1036	2.65	1.34	229.08	0.71	29^{+1}_{-1}	13^{+1}_{-1}	0.000	1.27	0.88	81^{+6}_{-2}	19^{+3}_{-1}	0.156	1.27	0.88	28^{+14}_{-13}	66^{+6}_{-6}	0.086
J0030+0451	2.89	1.62	231.62	0.90	10^{+1}_{-1}	47^{+1}_{-1}	0.047	5.54	0.69	86^{+1}_{-1}	69^{+1}_{-1}	0.039	28.79	-0.61	72^{+1}_{-1}	60^{+1}_{-1}	0.078
J1614-2230	2.55	1.28	51.73	0.89	75^{+1}_{-1}	25^{+1}_{-1}	0.031	2.73	0.66	80^{+4}_{-5}	89^{+1}_{-2}	0.141	3.63	0.54	84^{+1}_{-1}	83^{+1}_{-1}	0.141
J1513-5908	2.04	1.12	2.21	0.95	8^{+1}_{-1}	23^{+1}_{-1}	-0.016	0.90	0.89	65^{+12}_{-38}	31^{+38}_{-16}	0.047	1.07	0.87	35^{+9}_{-15}	64^{+3}_{-3}	0.000
J1718-3825	2.37	1.48	35.33	0.84	14^{+1}_{-1}	38^{+1}_{-1}	-0.023	0.84	0.96	76^{+3}_{-2}	26^{+6}_{-6}	0.102	4.74	0.76	54^{+1}_{-1}	47^{+1}_{-2}	0.086
J1709-4429	3.15	2.31	45.23	0.99	33^{+1}_{-1}	13^{+1}_{-1}	0.000	18.58	0.96	31^{+1}_{-1}	68^{+1}_{-1}	0.016	146.51	0.66	56^{+1}_{-1}	45^{+1}_{-1}	0.117
J1420-6048	2.24	1.61	16.96	0.86	34^{+1}_{-1}	27^{+1}_{-1}	0.016	3.82	0.90	85^{+4}_{-3}	26^{+5}_{-2}	0.070	2.87	0.92	60^{+2}_{-7}	45^{+4}_{-7}	0.094
J1231-1411	2.26	1.71	4.15	0.97	54^{+1}_{-1}	18^{+1}_{-1}	0.016	15.26	0.44	82^{+1}_{-1}	67^{+1}_{-1}	0.078	120.19	-3.41	71^{+1}_{-1}	59^{+1}_{-1}	0.109
J1833-1034	1.92	1.42	4.70	0.96	41^{+2}_{-1}	12^{+2}_{-1}	0.000	3.41	0.77	81^{+8}_{-15}	83^{+6}_{-8}	0.031	14.75	0.00	67^{+1}_{-1}	80^{+1}_{-1}	0.031
J1509-5850	1.68	1.57	1.98	0.79	10^{+1}_{-1}	31^{+1}_{-1}	0.023	2.52	0.92	29^{+8}_{-9}	75^{+1}_{-1}	0.070	2.68	0.91	34^{+5}_{-5}	68^{+1}_{-1}	0.062
J1747-2958	1.67	1.64	2.93	0.92	24^{+2}_{-1}	11^{+3}_{-1}	0.000	5.39	0.87	82^{+5}_{-4}	43^{+7}_{-7}	0.031	4.36	0.90	67^{+2}_{-2}	53^{+3}_{-3}	0.062
J0205+6449	1.60	1.58	2.14	0.65	17^{+1}_{-2}	42^{+1}_{-1}	0.016	3.39	0.77	88^{+1}_{-1}	89^{+1}_{-1}	0.000	12.88	0.14	84^{+1}_{-1}	86^{+1}_{-1}	0.000
J2229+6114	1.74	1.95	1.60	0.87	34^{+2}_{-3}	11^{+9}_{-2}	-0.008	6.19	0.92	68^{+1}_{-1}	23^{+1}_{-1}	0.188	23.20	0.72	54^{+1}_{-1}	45^{+1}_{-1}	0.148
J1124-5916	1.23	1.54	1.81	0.63	29^{+58}_{-8}	12^{+64}_{-2}	0.000	3.65	0.71	89^{+1}_{-1}	89^{+1}_{-1}	0.060	1.96	0.84	88^{+1}_{-1}	88^{+1}_{-1}	0.060
J2021+3651	1.72	2.06	1.94	0.82	7^{+2}_{-1}	27^{+1}_{-1}	0.016	33.92	0.75	74^{+1}_{-1}	85^{+1}_{-1}	0.039	43.94	0.67	72^{+1}_{-1}	82^{+1}_{-1}	0.039
Ave. g.o.f.			0.88					0.82							0.29		
Median g.o.f.			0.90					0.87							0.67		

both the six sets of GOF values (in each table), values that correspond to a fit that is worse than the relevant background fit (i.e., the negative values) are again highlighted in grey.

The Fits Conducted using χ_c^2

For the joint fits conducted for the OG+Cone model, the values of $[\chi_c^2]_\nu$ exhibit a similar trend to the one that was observed for the Cone fit; i.e., the value of $[\chi_c^2]_\nu$ tends to decrease as the value of $\log(\lambda_{r\gamma})$ decreases. Just as was the case for the Cone fit, the associated implied joint GOF values do not exhibit an equally dramatic trend, and are fairly high for all 23 pulsars (although, again, a weaker trend does arguably exist). Similarly, the radio-band GOF values are fairly high for the majority of the pulsars; only for three of the pulsars is $\Psi_{\Phi,r}^2 < 0.70$. For the γ -ray-band GOF values, however, a definite trend exists; one which is indicative of radio dominance: the value of $\Psi_{\Phi,\gamma}^2$ tends to be lower for higher values of $\log(\lambda_{r\gamma})$, and is negative for the majority of the pulsars listed first (their values are highlighted in grey). This trend is, arguably, mirrored in the radio-band GOF values to some extent since $\Psi_{\Phi,\gamma}^2 > \Psi_{\Phi,r}^2$ for all three of the pulsars for which $\log(\lambda_{r\gamma}) < 0$, and since $\Psi_{\Phi,r}^2 < 0.5$ for two of those three pulsars.

The overall radio dominance is also evident in the imbalance in the average single-band GOF values (0.80 for $\Psi_{\Phi,r}^2$ vs. -0.36 for $\Psi_{\Phi,\gamma}^2$), and in the fact that the average implied joint GOF value (i.e., the average value of $[\Psi_c^2]_\Phi$) is roughly equal to the average radio-band GOF value (i.e., the average value of $\Psi_{\Phi,r}^2$). The values for $f_{r,pcs}$ for the OG+Cone fits, which are listed in the fourth column of Table 10.5, confirm that these fits are, indeed, single-band dominated, with $f_{r,pcs}$ close to or equal to 1 for the majority of the pulsars listed first (which are associated with higher values for $\log(\lambda_{r\gamma})$), and close to or equal to -1 for two of the three pulsars for which $\log(\lambda_{r\gamma}) < 0$.³ The plot of $f_{r,pcs}$ vs. $\log(\lambda_{r\gamma})$ in the top left panel in Figure 10.1 further confirms this, since it has the characteristic “S” shape associated with single-band dominance (at least, the portion of the “S” for which $\log(\lambda_{r\gamma}) > 0$; see Section 6.3). Note that the value of $\tilde{\eta}_s$ (which is 1 if $\eta_s \geq 1$ and equal to η_s otherwise; see Section 6.3) is encoded in the colour intensities of the dots (in all four plots in Figure 10.1): The darker the dot, the greater $\tilde{\eta}_s$ (with $\tilde{\eta}_s = 1$ for fully black dots). For the joint fits conducted for the OG+Cone dual-band model, $F_{pcs} = 0.922$.

A similar situation exists for the TPC+Cone fits, the results of which are listed in Table 10.4 (in an identical fashion to the OG+Cone fits): The implied joint GOF values (for the χ_c^2 fit) and radio-band GOF values are fairly high for all 23 pulsars, but $\Psi_{\Phi,\gamma}^2$ tends to be drastically lower for pulsars where $\log(\lambda_{r\gamma})$ is substantially greater than 0. These trends are again reflected in the average GOF values: The average value of $[\Psi_c^2]_\Phi$ is again approximately equal to the average value of $\Psi_{\Phi,r}^2$, and the average value of $\Psi_{\Phi,\gamma}^2$ is substantially smaller. Examining the values for $f_{r,pcs}$ for the TPC+Cone fit (listed in the tenth column of Table 10.5) confirms that the TPC+Cone fits conducted using χ_c^2 are also single-band dominated. In fact, such an examination reveals that this dominance is even more severe than that which is observed in the OG+Cone fits, since $|f_{r,pcs}| < 0.90$ for only 5 of the 23 fits, as opposed to 9 out of the 23 fits for the OG+Cone fits. This difference is likely caused by the overall worse performance of the TPC model (as compared to the OG model), both for the single-band fits and for the dual-band fits (as evidenced by the lower median γ -ray GOF for the TPC+Cone fits; -0.75 vs. 0.47).

The plot of $f_{r,pcs}$ vs. $\log(\lambda_{r\gamma})$ for the TPC+Cone fits, in the top right panel of Figure 10.1 confirms this assessment, since it too has the characteristic “S” shape. In fact, just as one might expect from the stronger trend in the value of $f_{r,pcs}$ for the TPC+Cone fits, this plot resembles the archetypical “S” shape even more strongly. For the joint fits conducted for the TPC+Cone dual-band model, $F_{pcs} = 0.963$; higher than for the OG+Cone fits.

Overall, the OG+Cone and TPC+Cone dual-band models fare roughly equally well for the fits conducted using χ_c^2 (according to the respective median $[\Psi_c^2]_\Phi$ values). This parity is likely due to the general radio dominance, since for the majority of the pulsars the radio-only best-fit LC is effectively competing with itself.

The Fits Conducted using $\Psi_{\Phi,c}^2$

Curiously, the trend that the radio-band GOF (i.e., $\Psi_{\Phi,r}^2$) tends to be higher than the γ -ray-band GOF (i.e., $\Psi_{\Phi,\gamma}^2$) when $\log(\lambda_{r\gamma}) > 0$ seems to still be present for the OG+Cone fits. The radio-band GOF does not dominate as severely, though, since $\Psi_{\Phi,\gamma}^2 > 0$ for all the pulsars (in contrast to the values for the fit conducted using χ_c^2 , where $\Psi_{\Phi,\gamma}^2 < 0$ for 10 of the 23 pulsars). This means that the radio dominance that plagues the χ_c^2 seems to have been mitigated in the fits conducted

³To aid quick comprehension, values of $f_r = 1$ or -1 are highlighted in grey in Table 10.5, for both the χ_c^2 and $\Psi_{\Phi,c}^2$ fits.

Table 10.3: Results for the OG+Cone dual-band fits.

PSR	$\log(\lambda_{\gamma})$	M_c^{eye}		M_c^{pcs}							M_c^{sis}					
		α ($^{\circ}$)	ζ ($^{\circ}$)	$[\chi^2]_{\nu}$	$[\Psi^2]_{\Phi}$	$\Psi^2_{\Phi,r}$	$\Psi^2_{\Phi,\gamma}$	α ($^{\circ}$)	ζ ($^{\circ}$)	ϕ_{μ}	$\Psi^2_{\Phi,c}$	$\Psi^2_{\Phi,r}$	$\Psi^2_{\Phi,\gamma}$	α ($^{\circ}$)	ζ ($^{\circ}$)	ϕ_{μ}
J0437−4715	2.41	31^{+10}_{-10}	62^{+4}_{-4}	3834.72	0.96	0.96	−1.30	28^{+1}_{-1}	69^{+1}_{-1}	0.031	0.70	0.92	0.48	23^{+1}_{-1}	64^{+1}_{-1}	0.031
J0908−4913	2.36	70^{+1}_{-1}	85^{+2}_{-2}	11758.99	0.84	0.84	−1.70	43^{+1}_{-1}	66^{+1}_{-1}	0.000	0.54	0.74	0.34	67^{+1}_{-1}	89^{+1}_{-1}	0.008
J0742−2822	2.16	74^{+4}_{-4}	85^{+4}_{-4}	254.75	0.98	0.98	−1.25	63^{+1}_{-1}	50^{+1}_{-1}	0.008	0.57	0.88	0.26	77^{+1}_{-9}	88^{+7}_{-7}	0.016
J0248+6021	1.72	44^{+4}_{-4}	57^{+3}_{-3}	865.79	0.85	0.85	0.47	45^{+1}_{-1}	52^{+1}_{-1}	0.023	0.78	0.84	0.71	48^{+1}_{-1}	55^{+1}_{-1}	0.023
J2241−5236	1.67	79^{+1}_{-1}	25^{+5}_{-5}	231.85	0.96	0.96	−0.04	87^{+1}_{-1}	27^{+1}_{-1}	0.055	0.70	0.87	0.54	20^{+1}_{-1}	78^{+1}_{-1}	0.062
J0835−4510	1.53	68^{+2}_{-2}	83^{+2}_{-2}	55860.19	0.96	0.96	−1.75	44^{+1}_{-1}	71^{+1}_{-1}	0.008	0.84	0.86	0.81	82^{+1}_{-1}	58^{+1}_{-1}	0.008
J0659+1414	1.53	57^{+3}_{-3}	49^{+2}_{-2}	142.76	0.94	0.94	−0.16	48^{+1}_{-1}	55^{+1}_{-1}	0.000	0.50	0.90	0.09	54^{+1}_{-1}	48^{+1}_{-1}	−0.008
J1048−5832	1.50	61^{+9}_{-9}	79^{+7}_{-7}	550.13	0.95	0.95	−4.12	53^{+1}_{-1}	42^{+1}_{-1}	−0.008	0.88	0.92	0.84	80^{+1}_{-1}	68^{+1}_{-1}	−0.008
J0631+1036	1.31	54^{+10}_{-10}	53^{+9}_{-9}	201.74	0.70	0.70	0.34	54^{+1}_{-1}	62^{+1}_{-1}	0.000	0.64	0.69	0.59	69^{+8}_{-3}	61^{+7}_{-3}	0.000
J0030+0451	1.27	60^{+8}_{-8}	78^{+4}_{-4}	195.09	0.89	0.89	−2.26	38^{+1}_{-1}	57^{+1}_{-1}	0.047	0.74	0.83	0.64	86^{+1}_{-1}	67^{+1}_{-1}	0.039
J1614−2230	1.26	54^{+6}_{-6}	84^{+4}_{-4}	50.29	0.88	0.89	−2.97	26^{+1}_{-1}	75^{+1}_{-1}	0.039	0.37	0.23	0.50	88^{+1}_{-14}	84^{+5}_{-5}	0.125
J1513−5908	0.92	58^{+4}_{-4}	42^{+4}_{-4}	7.51	0.83	0.83	0.58	44^{+1}_{-1}	53^{+1}_{-1}	−0.016	0.84	0.81	0.87	47^{+1}_{-1}	57^{+1}_{-1}	−0.016
J1718−3825	0.90	55^{+4}_{-4}	55^{+4}_{-4}	40.08	0.79	0.79	0.50	67^{+2}_{-2}	47^{+3}_{-1}	−0.039	0.73	0.77	0.69	67^{+3}_{-1}	56^{+3}_{-3}	−0.023
J1709−4429	0.84	62^{+4}_{-4}	45^{+5}_{-5}	182.04	0.97	0.97	0.84	48^{+1}_{-1}	60^{+1}_{-1}	0.008	0.92	0.97	0.87	49^{+1}_{-1}	61^{+1}_{-1}	0.008
J1420−6048	0.63	60^{+6}_{-6}	61^{+6}_{-6}	23.37	0.78	0.78	0.79	66^{+1}_{-1}	51^{+1}_{-1}	0.062	0.79	0.78	0.80	67^{+1}_{-1}	52^{+1}_{-1}	0.062
J1231−1411	0.55	85^{+2}_{-2}	59^{+3}_{-3}	21.78	0.79	0.94	−1.13	85^{+1}_{-1}	59^{+1}_{-1}	0.031	0.42	0.41	0.43	82^{+1}_{-1}	66^{+1}_{-1}	0.078
J1833−1034	0.50	85^{+5}_{-5}	64^{+7}_{-7}	12.15	0.83	0.86	0.47	82^{+1}_{-1}	65^{+1}_{-1}	0.000	0.69	0.63	0.74	67^{+5}_{-3}	80^{+4}_{-4}	0.031
J1509−5850	0.11	58^{+5}_{-5}	59^{+6}_{-6}	3.59	0.71	0.74	0.65	54^{+3}_{-3}	65^{+3}_{-3}	0.016	0.70	0.73	0.67	76^{+4}_{-6}	65^{+6}_{-6}	0.016
J1747−2958	0.02	72^{+4}_{-4}	54^{+8}_{-8}	8.45	0.78	0.70	0.86	63^{+2}_{-2}	72^{+2}_{-2}	0.031	0.78	0.70	0.86	63^{+1}_{-1}	72^{+1}_{-1}	0.031
J0205+6449	0.02	65^{+5}_{-5}	88^{+2}_{-2}	3.67	0.57	0.58	0.56	73^{+1}_{-1}	86^{+1}_{-1}	0.000	0.57	0.58	0.56	73^{+1}_{-1}	86^{+1}_{-1}	0.000
J2229+6114	−0.21	69^{+5}_{-5}	46^{+6}_{-6}	6.81	0.78	0.72	0.81	64^{+1}_{-1}	50^{+1}_{-1}	0.016	0.78	0.78	0.77	64^{+1}_{-1}	50^{+1}_{-1}	0.008
J1124−5916	−0.30	56^{+8}_{-8}	86^{+3}_{-3}	3.59	0.61	0.25	0.71	89^{+5}_{-5}	89^{+1}_{-1}	0.060	0.48	0.26	0.70	88^{+1}_{-28}	89^{+11}_{-11}	0.060
J2021+3651	−0.34	70^{+4}_{-4}	85^{+4}_{-4}	14.10	0.68	0.42	0.74	75^{+1}_{-1}	84^{+1}_{-1}	0.039	0.63	0.61	0.65	87^{+1}_{-1}	75^{+1}_{-1}	0.023
Ave. g.o.f.				0.83	0.80	−0.36					0.68	0.73	0.63			
Median g.o.f.				0.83	0.85	0.47					0.70	0.78	0.67			

Table 10.4: Results for the TPC+Cone dual-band fits.

PSR	$\log(\lambda_{\text{r}\gamma})$	$M_{\text{c}}^{\text{eye}}$		$M_{\text{c}}^{\text{pcs}}$						$M_{\text{c}}^{\text{sis}}$						
		α ($^{\circ}$)	ζ ($^{\circ}$)	$[\chi_{\text{c}}^2]_{\nu}$	$[\Psi_{\text{c}}^2]_{\Phi}$	$\Psi_{\Phi,\text{r}}^2$	$\Psi_{\Phi,\gamma}^2$	α ($^{\circ}$)	ζ ($^{\circ}$)	ϕ_{μ}	$\Psi_{\Phi,\text{c}}^2$	$\Psi_{\Phi,\text{r}}^2$	$\Psi_{\Phi,\gamma}^2$	α ($^{\circ}$)	ζ ($^{\circ}$)	ϕ_{μ}
J0437–4715	2.41	57^{+9}_{-9}	26^{+10}_{-10}	3835.95	0.96	0.96	-2.09	28^{+1}_{-1}	69^{+1}_{-1}	0.031	0.37	0.50	0.23	48^{+1}_{-1}	53^{+1}_{-1}	0.094
J0908–4913	2.36	64^{+11}_{-11}	82^{+7}_{-7}	11349.47	0.84	0.84	-1.26	53^{+1}_{-1}	30^{+1}_{-1}	0.000	0.56	0.47	0.64	79^{+2}_{-2}	84^{+1}_{-1}	0.031
J0742–2822	2.16	88^{+5}_{-5}	74^{+5}_{-5}	254.48	0.98	0.98	-0.75	63^{+1}_{-1}	50^{+1}_{-1}	0.008	0.53	0.87	0.19	66^{+8}_{-4}	79^{+6}_{-1}	0.016
J0248+6021	1.72	43^{+3}_{-3}	57^{+3}_{-3}	596.29	0.90	0.90	-3.62	8^{+1}_{-1}	25^{+1}_{-1}	0.023	0.82	0.85	0.79	47^{+1}_{-1}	54^{+1}_{-1}	0.023
J2241–5236	1.67	30^{+2}_{-2}	59^{+2}_{-2}	233.66	0.96	0.96	-0.79	87^{+1}_{-1}	27^{+1}_{-1}	0.055	0.70	0.87	0.53	20^{+1}_{-1}	78^{+1}_{-1}	0.062
J0835–4510	1.53	54^{+7}_{-7}	75^{+7}_{-7}	60797.61	0.95	0.96	-6.17	44^{+1}_{-1}	71^{+1}_{-1}	0.008	-0.54	0.67	-1.76	66^{+1}_{-1}	80^{+1}_{-1}	0.016
J0659+1414	1.53	34^{+6}_{-6}	43^{+6}_{-6}	37.97	0.98	0.99	-6.31	13^{+1}_{-1}	26^{+1}_{-1}	0.000	0.46	0.89	0.02	53^{+1}_{-1}	46^{+1}_{-1}	-0.008
J1048–5832	1.50	64^{+2}_{-2}	77^{+3}_{-3}	558.11	0.95	0.95	-6.45	53^{+1}_{-1}	41^{+1}_{-1}	-0.008	0.32	0.94	-0.29	68^{+1}_{-1}	80^{+1}_{-1}	0.000
J0631+1036	1.31	54^{+7}_{-7}	52^{+9}_{-9}	195.88	0.70	0.71	-1.58	29^{+1}_{-1}	13^{+1}_{-1}	0.000	0.60	0.69	0.52	62^{+5}_{-3}	54^{+4}_{-3}	0.000
J0030+0451	1.27	59^{+6}_{-6}	78^{+5}_{-5}	208.21	0.88	0.90	-5.13	39^{+1}_{-1}	58^{+1}_{-1}	0.047	-0.07	0.51	-0.64	74^{+1}_{-1}	60^{+1}_{-1}	0.070
J1614–2230	1.26	85^{+5}_{-5}	54^{+7}_{-7}	51.99	0.87	0.89	-4.36	26^{+1}_{-1}	75^{+1}_{-1}	0.039	0.33	0.19	0.46	84^{+2}_{-6}	83^{+2}_{-3}	0.133
J1513–5908	0.92	58^{+8}_{-8}	43^{+7}_{-7}	2.95	0.93	0.95	-0.48	8^{+3}_{-1}	23^{+1}_{-1}	-0.016	0.82	0.81	0.83	47^{+1}_{-1}	56^{+1}_{-1}	-0.008
J1718–3825	0.90	47^{+7}_{-7}	55^{+8}_{-8}	40.47	0.78	0.80	0.03	37^{+1}_{-1}	59^{+1}_{-1}	-0.031	0.54	0.74	0.34	47^{+1}_{-1}	54^{+1}_{-1}	0.008
J1709–4429	0.84	62^{+4}_{-4}	44^{+5}_{-5}	264.66	0.95	0.97	0.15	48^{+1}_{-1}	60^{+1}_{-1}	0.008	0.63	0.83	0.43	41^{+1}_{-1}	64^{+1}_{-1}	0.008
J1420–6048	0.63	59^{+2}_{-2}	53^{+2}_{-2}	22.48	0.79	0.84	-0.23	33^{+4}_{-6}	41^{+4}_{-5}	0.023	0.83	0.78	0.88	61^{+1}_{-1}	47^{+1}_{-1}	0.062
J1231–1411	0.55	75^{+5}_{-5}	46^{+7}_{-7}	79.90	0.23	0.55	-3.94	65^{+1}_{-1}	79^{+1}_{-1}	0.086	-1.58	0.46	-3.62	61^{+1}_{-1}	69^{+1}_{-1}	0.102
J1833–1034	0.50	53^{+8}_{-8}	78^{+2}_{-2}	17.12	0.75	0.86	-0.31	63^{+1}_{-1}	79^{+1}_{-1}	0.000	0.32	0.63	0.00	67^{+1}_{-1}	80^{+1}_{-1}	0.031
J1509–5850	0.11	54^{+2}_{-2}	52^{+2}_{-2}	3.74	0.70	0.75	0.61	48^{+2}_{-2}	59^{+2}_{-2}	0.016	0.68	0.74	0.62	48^{+2}_{-2}	60^{+2}_{-2}	0.016
J1747–2958	0.02	55^{+5}_{-5}	71^{+4}_{-4}	8.48	0.78	0.71	0.85	61^{+1}_{-1}	70^{+1}_{-1}	0.031	0.78	0.71	0.85	61^{+1}_{-1}	70^{+1}_{-1}	0.031
J0205+6449	0.02	67^{+4}_{-4}	87^{+3}_{-3}	5.82	0.32	0.59	0.02	72^{+1}_{-1}	86^{+1}_{-1}	0.000	0.31	0.59	0.02	72^{+1}_{-1}	86^{+1}_{-1}	0.000
J2229+6114	-0.21	62^{+10}_{-10}	48^{+8}_{-8}	14.83	0.52	0.01	0.72	54^{+1}_{-1}	45^{+1}_{-1}	0.148	0.38	0.82	-0.07	51^{+1}_{-1}	65^{+2}_{-2}	0.000
J1124–5916	-0.30	67^{+8}_{-8}	84^{+5}_{-5}	2.59	0.72	0.25	0.84	89^{+1}_{-2}	89^{+1}_{-2}	0.060	0.56	0.42	0.69	83^{+6}_{-13}	80^{+9}_{-4}	0.040
J2021+3651	-0.34	84^{+4}_{-4}	67^{+5}_{-5}	16.61	0.63	0.41	0.67	72^{+1}_{-1}	82^{+1}_{-1}	0.039	0.55	0.47	0.63	75^{+1}_{-1}	82^{+1}_{-1}	0.039
Ave. g.o.f.				0.79	0.77		-1.72				0.39	0.67	0.10			
Median g.o.f.				0.84	0.86		-0.75				0.54	0.71	0.43			

using $\Psi_{\Phi,c}^2$ to some extent. As discussed in Section 6.2.2, though, a direct comparison such as this, of the values of $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$, may result in a misleading characterisation of the degree to which the radio-band GOF enjoys priority. This trend persists for the more robust $f_{r,sfs}$ values (with $F_{sfs} = 0.372$), however, which confirms that (for the OG+Cone fits) the radio-band GOF does, indeed, tend to enjoy higher priority for pulsars associated with higher $\log(\lambda_{r\gamma})$ values.

The plot of $f_{r,sfs}$ vs. $\log(\lambda_{r\gamma})$ in the lower left panel of Figure 10.1, however, tells a somewhat different story. While it is true that $f_{r,sfs}$ is close to 1 for the dots at the far right end of the plot, the other dots are comparatively evenly distributed about $f_{r,sfs} = 0$ (i.e., it appears as if $f_{r,sfs} > 0$ for just as many pulsars as $f_{r,sfs} < 0$, if the dots to the far right were ignored). Re-examining the values for $f_{r,sfs}$ listed Table 10.5 (in the fifth column) reveals that $f_{r,sfs} > 0$ just as frequently as $f_{r,sfs} < 0$ for the 16 pulsars with the lowest $\log(\lambda_{r\gamma})$ values. While this does not mean that there is no residual radio dominance, it does indicate that the radio dominance is likely effectively eliminated for a comparatively broad range of $\log(\lambda_{r\gamma})$ values ($\log(\lambda_{r\gamma}) \lesssim 1.5$).

For the TPC+Cone fits, a direct comparison of $\Psi_{\Phi,r}^2$ and $\Psi_{\Phi,\gamma}^2$ values for each of the fits again hints at some residual radio dominance (since $\Psi_{\Phi,r}^2 > \Psi_{\Phi,\gamma}^2$ for 16 of the 23 pulsars). In this case, however, the trend arguably does not manifest for the values of $f_{r,sfs}$ (listed in the eleventh column of Table 10.5), since $f_{r,sfs} > 0$ roughly as frequently as $f_{r,sfs} < 0$, with like values not clustering (i.e., with it not being true that the $f_{r,sfs} > 0$ for the majority of the pulsars listed first, and vice versa). The plot of $f_{r,sfs}$ vs. $\log(\lambda_{r\gamma})$ for the TPC+Cone fits (in the lower right panel of Figure 10.1) confirms the absence of residual radio dominance, with $F_{sfs} = -0.107$.

Overall, the OG+Cone model (for which the median joint GOF is 0.70) fares better than the TPC+Cone model (for which it is 0.54). For both models, the average radio-band GOF is somewhat reduced compared to the average radio-band GOF for the χ_c^2 fits, by 0.07 for the OG+Cone model and 0.10 for the TPC+Cone model. However, these reductions in radio-band GOF are far outweighed by the accompanying boost in γ -ray-band GOF, by 0.99 for the OG+Cone model and 1.82 for the TPC+Cone model. In essence, the fits conducted using $\Psi_{\Phi,c}^2$ compromise on the radio-band GOF in order to (more greatly) improve the γ -ray-band GOF.

The Impact of Non-colocation

As discussed in Section 6.2.1, single-band best-fit non-colocation enables single-band dominance by inducing a degeneracy in the identity of the model's joint best fit. Table 10.5 lists (in its sixth and twelfth columns) the value of η_s (which characterises this non-colocation), as well as the average loss in single-band-only GOF it results in (via κ_{pcs} and κ_{sfs} in its seventh, eighth, thirteenth and fourteenth columns; see Section 6.2.2), for each dual-band fit. The median non-colocation for the OG+Cone fits (i.e., the median value for η_s) is 0.77, while for the TPC+Cone fits the median non-colocation is 4.06; recall that $\eta_s > 1$ implies that either the radio counterpart to the γ -ray-only best-fit LC or the γ -ray counterpart to the radio-only best-fit LC (or both) is a worse fit than the relevant background LC (i.e., that $\Psi_{\Phi}^2 < 0$ for at least one of these two counterparts). This indicates that the degree to which the TPC and Cone best fits are non-colocated is significantly higher than the degree to which the OG and Cone best fits are non-colocated. Looking at the individual values for η_s , this difference becomes even more striking since $\eta_s < 1$ for only 2 of the 23 TPC+Cone fits and for 19 of the 23 OG+Cone fits.

For the OG+Cone model, the median non-colocation penalty incurred by the χ_c^2 statistic is 0.25, whereas for the $\Psi_{\Phi,c}^2$ statistic this median is 0.15; i.e., the median average loss in single-band GOF (due to compromise) associated with the $\Psi_{\Phi,c}^2$ statistic is 0.10 lower than that associated with the χ_c^2 statistic. For the TPC+Cone model, both these medians and their difference is higher, at 0.61, 0.20 and 0.41. This is consistent with the difference in median η_s values. Interestingly, $\kappa_{sfs} < 0.10$ for the same number of pulsars for both the OG+Cone and TPC+Cone fits.

Qualitative By-eye Evaluation of the Best-fit LC Pairs

While the tools developed in Chapter 6 provide powerful insight into the effectiveness of $\Psi_{\Phi,c}^2$, a by-eye evaluation of the best-fit LC pairs can still be instructive. Evaluating the best-fit LC pairs plotted in the figures in Appendix A reveals that the $\Psi_{\Phi,c}^2$ statistic tends to more often identify best-fit LC pairs that roughly reproduce both of the observed LC pairs than does the χ_c^2 statistic. Particularly good examples of such cases are the fits for PSR J2241–5236, PSR J0835–4510, PSR J1048–5832 and PSR J0030+0451. For all four of these pulsars, χ_c^2 identifies best-fit LC pairs where, in effect, only

Table 10.5: Fit characteristics for the dual-band fits.

PSR	$\log(\lambda_{\gamma})$	OG + Cone						TPC + Cone					
		$\log(\eta_c)$	$f_{r,pcs}$	$f_{r,sfs}$	η_s	κ_{pcs}	κ_{sfs}	$\log(\eta_c)$	$f_{r,pcs}$	$f_{r,sfs}$	η_s	κ_{pcs}	κ_{sfs}
J0437–4715	2.41	−1.45	1.00	0.88	2.12	1.10	0.23	−0.37	1.00	−0.72	2.33	1.21	0.28
J0908–4913	2.36	−1.05	1.00	0.68	0.60	1.21	0.24	−0.55	1.00	−0.50	1.73	1.03	0.26
J0742–2822	2.16	−1.17	1.00	0.72	1.89	0.97	0.26	−1.35	1.00	0.55	1.20	0.61	0.20
J0248+6021	1.72	−2.94	0.87	0.67	0.83	0.25	0.13	−0.97	1.00	0.41	4.06	2.25	0.08
J2241–5236	1.67	−1.66	1.00	0.45	0.74	0.39	0.14	−1.26	1.00	0.63	1.55	0.81	0.19
J0835–4510	1.53	−0.94	1.00	0.09	1.28	1.34	0.10	−0.24	1.00	−0.51	1.94	2.26	0.20
J0659+1414	1.53	−2.35	0.95	0.87	0.93	0.58	0.47	−0.54	1.00	0.76	6.78	3.42	0.30
J1048–5832	1.50	−1.19	1.00	0.05	0.17	2.50	0.04	−1.48	1.00	−0.79	0.39	3.08	0.01
J0631+1036	1.31	−2.98	0.97	0.90	0.62	0.28	0.15	−1.72	1.00	0.92	1.74	1.23	0.19
J0030+0451	1.27	−1.14	1.00	−0.18	0.13	1.48	0.06	−0.12	1.00	−0.90	27.60	2.26	0.21
J1614–2230	1.26	−0.04	1.00	−0.70	7.26	1.82	0.41	0.15	1.00	−0.86	15.63	2.46	0.39
J1513–5908	0.92	−2.65	0.52	−0.82	0.85	0.22	0.08	−1.14	1.00	−0.64	1.29	0.68	0.09
J1718–3825	0.90	−2.82	0.87	0.67	0.80	0.25	0.17	−2.16	0.92	0.69	4.29	0.39	0.26
J1709–4429	0.84	−5.19	0.76	0.67	0.95	0.07	0.06	−2.04	0.94	0.24	11.11	0.27	0.20
J1420–6048	0.63	−5.50	0.18	0.08	0.77	0.09	0.09	−3.25	0.98	−0.33	0.87	0.58	0.06
J1231–1411	0.55	−0.71	0.98	−0.97	0.26	0.80	0.29	−2.13	0.15	−0.51	70.16	0.47	0.36
J1833–1034	0.50	−2.44	0.59	−0.88	0.69	0.20	0.18	−2.13	0.60	−1.00	5.58	0.20	0.16
J1509–5850	0.11	−5.13	0.80	0.69	0.72	0.15	0.15	−6.14	0.85	0.81	2.32	0.17	0.17
J1747–2958	0.02	...	−0.92	−0.92	0.80	0.11	0.11	...	−0.71	−0.71	2.56	0.13	0.13
J0205+6449	0.02	...	0.59	0.59	0.41	0.14	0.14	...	0.37	0.37	20.51	0.09	0.09
J2229+6114	−0.21	−4.13	−0.14	0.35	0.80	0.13	0.12	−2.36	−1.00	0.93	11.20	0.43	0.42
J1124–5916	−0.30	−5.50	−1.00	−0.99	0.27	0.19	0.19	−2.99	−0.99	−0.18	5.22	0.19	0.18
J2021+3651	−0.34	−2.68	−0.97	−0.41	0.34	0.20	0.16	−3.64	−1.00	−0.84	20.89	0.21	0.20

the radio LC shape is reproduced, while $\Psi_{\Phi,c}^2$ identifies LC pairs where the shape of both of the single-band observed LCs are roughly reproduced. This indicates that $\Psi_{\Phi,c}^2$ is, indeed, less prone to radio dominance than is χ_c^2 .

Notably, for 4 of the 23 pulsars (PSR J0437–4715, PSR J0908–4913, PSR J1614–2230 and PSR J1231–1411), the radio component of the TPC+Cone best-fit LC pair that was identified using $\Psi_{\Phi,c}^2$ has a broad double-peak structure while the observed LC has a roughly narrow single-peak structure. In all four cases the leading model peak coincides with the single observed radio peak. This indicates that these model LC pairs were favoured because ϕ_μ is larger for them, and they more accurately reproduce the radio-to- γ -ray phase lag (i.e., the amount by which the main γ -ray peak lags behind the main radio peak in phase). For the OG+Cone model, this happens for two pulsars (PSR J1614–2230 and PSR J1231–1411). For the fit to the observed LC pair for PSR J1420–6048 (excluding the TPC+Cone fit conducted using χ_c^2), something similar happens: The radio component of the best-fit LC pair has a somewhat broad, single peak, which coincides with the trailing peak in the double-peaked observed radio LC. Additionally, for these fits to PSR J1420–6048 the peak in the γ -ray component of each of the best-fit LC pairs coincides precisely with the peak in the observed γ -ray LC. This again indicates that these models cannot accurately reproduce both the profile shapes and the radio-to- γ -ray phase lag simultaneously for these cases, and likely need to be further refined.

The Convergence of χ_c^2 and $\Psi_{\Phi,c}^2$

As established in Section 8.1.4, the degree to which the region around the point where $\eta_s = 0$ and $\lambda_{\gamma} = 1$ (i.e., the point where $\eta_s = 0$ and $\log(\lambda_{\gamma}) = 0$) is covered for a given set of fits greatly impacts how effectively the convergence of χ_c^2 and $\Psi_{\Phi,c}^2$ (with regard to the best fits they identify) can be tested. Unfortunately, in this test we have somewhat less opportunity to ensure that this coverage is optimal (compared to the test using synthetic observed LC pairs), since it is difficult to ensure that $\eta_s \simeq 0$ for some number of the pulsars included in the set of 23 pulsars (without actually doing the fits, at least). However, as discussed in Section 10.1, we were able to ensure that the $\log(\lambda_{\gamma})$ values for the 23 selected pulsars cover as large a range as possible, and that they are distributed as evenly as possible within that range.




Figure 10.1: Plots of radio priority f_r vs. scaled flux ratio $\log(\lambda_{r\gamma})$ (plotted logarithmically) for 23 *Fermi* LAT pulsars. Each row corresponds to a different joint GOF statistic (χ_c^2 at the *top* and $\Psi_{\Phi,c}^2$ at the *bottom*) and each column corresponds to a different dual-band model (OG+Cone on the *left* and TPC+Cone on the *right*). The colour intensities of the filled circles (each associated with a pulsar) encodes the value of $\tilde{\eta}_s$ (the single-band best-fit non-colocation, with a higher colour intensity corresponding to a greater non-colocation) for that combination of pulsar and dual-band model. The farther right a circle appears, the more potential there is for radio dominance, and vice versa for γ -ray dominance. The strongly S-shaped distributions for the χ_c^2 fits indicate that the χ_c^2 statistic is single-band dominated, while the more even distributions for the $\Psi_{\Phi,c}^2$ fits indicates that the $\Psi_{\Phi,c}^2$ statistic is likely undominated.




Figure 10.2: Plots of the difference between the parameter point estimates derived by [Johnson et al. \(2014\)](#) and [Pierbattista et al. \(2015\)](#) and those derived here. The plots in the top row are for dual-band fits using the OG+Cone model, while the plots in the bottom are for fits using the TPC+Cone model. The plots in the left-most column (the yellow plots) are for the fits conducted using χ^2_c , the plots in the middle column (the brown plots) are for the fits conducted by eye, and the plots in the right-most column (the green fits) are for the fits conducted using $\Psi^2_{\Phi,c}$. For each point estimate, the uncertainties associated with it are indicated using like-coloured error bars. The point estimates inside grey diamonds are deemed outliers with respect to the other 22 point estimates (in the same plot). For each scatterplot, the 2σ covariance ellipse is plotted, with any outliers ignored. To enable easy comparison between panels, each panel is equipped with a set of light grey reference lines (identical in each plot).

Of the two dual-band models, the coverage of the region around η_s is poorest for the TPC+Cone model, since $\eta_s < 1$ for only 2 of the 23 pulsars. For the OG+Cone model, $\eta_s < 1$ for 19 of the 23 pulsars, which indicates that the degree to which the OG and Cone single-band best fits are non-colocated tends to be substantially lower than the degree to which the TPC and Cone single-band best fits are non-colocated. The value of η_s is only less than 0.5 for 6 of the 23 OG+Cone fits, though, with the lowest value being 0.13. Therefore, the region around $\eta_s = 0$ is likely not covered sufficiently well for either dual-band model, and this test can likely only test the degree to which χ^2_c and $\Psi^2_{\Phi,c}$ converge as $\log(\lambda_{r\gamma}) \rightarrow 0$.

The values for $\log(\eta_c)$ listed in the third and ninth columns of Table 10.5 (for the OG+Cone and TPC+Cone fits, respectively), demonstrate that for these fits χ^2_c and $\Psi^2_{\Phi,c}$ tend to agree more closely for pulsars with lower $\log(\lambda_{r\gamma})$ values. Additionally, for the two pulsars with $\log(\lambda_{r\gamma}) \simeq 0$ (PSR J1747–2958 and PSR J0205+6449), χ^2_c and $\Psi^2_{\Phi,c}$ identify the same model LC pair as the model’s best fit (which means that $\eta_c = 0$, and $\log(\eta_c)$ is not defined; as indicated by the ellipses in Table 10.5). This seems to confirm, at least, that χ^2_c and $\Psi^2_{\Phi,c}$ converge with regard to the best fits they identify as $\log(\lambda_{r\gamma}) \rightarrow 0$.

10.4.3 Comparing Parameter Estimates

The differences between the parameter estimates derived here (for both dual-band models, via the by-eye approach and using the statistics χ^2_c and $\Psi^2_{\Phi,c}$) and those reported by [Johnson et al. \(2014\)](#) and [Pierbattista et al. \(2015\)](#) are plotted in

Figure 10.2 in the form of $(\Delta\alpha, \Delta\zeta)$ -space scatterplots (with uncertainties derived via propagation). In effect, for each combination of pulsar, dual-band model and fit type (grouped into panels based on the latter two categories; i.e., for each fit), the position (in (α, ζ) -space) of the estimate derived here is plotted relative to that of the previously-derived estimate. For each scatterplot, the 2σ covariance ellipse is also plotted. Important to note is that these ellipses do not take the uncertainties on each of the differences into account, and therefore characterise only the covariance of the point estimate differences. The fits inside the grey diamonds were considered to be outliers during the construction of these ellipses since they lie outside the 3σ confidence ellipses (not shown) for the other 22 pulsars in each case.

In the two top-most panels of Figure 10.3, the (normalised) marginal distributions of the OG+Cone differences as function of $\Delta\alpha$ and $\Delta\zeta$ are plotted, for all three sets of differences. The same marginal distributions are plotted for the TPC+Cone model in the top-most panels of Figure 10.4. These distributions were constructed via kernel density estimation; specifically, by treating each component of the parameter estimate as a normalised, asymmetrical triangular kernel that peaks at the point estimate for that parameter, with 2σ confidence level critical values derived from the constraints for that parameter. In the two bottom-most panels of Figures 10.3 and 10.4, two related distributions are plotted: those of the differences as function of $\Delta|\beta|$ and $\Delta(\alpha + \zeta)$. These were also constructed via kernel density estimation. Note that, while the marginal distributions as function of $\Delta(\alpha + \zeta)$ relate to the scatterplots in Figure 10.2 (see below), the marginal distributions as function of $\Delta|\beta|$ do not. These distributions attempt to crudely capture the relative degree to which the radio-band GOF was prioritised by leveraging the roughly circular (and hence symmetrical) shape of the radio emission cone on the (ϕ, ζ) -space phaseplots (e.g., the radio phaseplot in Figure 2.1(a); at least for $\alpha \gtrsim 15^\circ$). Note that the outliers were not removed before calculating the marginal distributions in Figures 10.3 and 10.4.

The scatterplots in the left-most column in Figure 10.2 show that the fits conducted using χ_c^2 tend to yield point estimates that either lie fairly close to the previously-derived point estimates or lie below and/or to the left of those point estimates (i.e., at lower α and/or ζ values). For the OG+Cone fits, $\Delta\zeta$ tends to be larger than $\Delta\alpha$ for the point estimates that are farther from the previously-derived point estimates. The reason for this is that the TPC model typically predicts pulsed emission for a larger range of ζ values than does the OG model (especially for small values of α). A quantity that captures this trend well (since it increases as one moves from the lower left corner of one of these plots to the upper right corner), is the sum $(\Delta\alpha + \Delta\zeta) = \Delta(\alpha + \zeta)$. Looking at the marginal distributions associated with this quantity (in the lower right panels of Figures 10.3 and 10.4), we see that $\Delta(\alpha + \zeta)$ does, indeed, tend to be less than zero for the χ_c^2 fits (the yellow bars). In terms of this quantity, the fits conducted using χ_c^2 tend to yield point estimates that either have roughly the same value for $\Delta(\alpha + \zeta)$ (than the previously-derived point estimates do) or have a lower value. Since the average value of $(\alpha + \zeta)$ for the radio-only point estimates is much lower than for the γ -ray-only point estimates (as compared to the results for either γ -ray model; 63.0° for the Cone estimates vs. 130.3° and 124.3° for the OG and TPC models), this trend indicates that the relative weight of the γ -ray-band GOF was higher for the fits conducted by Johnson et al. and Pierbattista et al. than for the fits conducted using χ_c^2 ; i.e., that the fits these authors conducted were less radio-dominated.

The covariance ellipses for the by-eye fits (in the middle column in Figure 10.2) are slightly smaller than those for the χ_c^2 fits. There is more uncertainty in $\Delta\alpha$ and $\Delta\zeta$ for the by-eye fits, however, since the by-eye parameter constraints are typically larger than those for the χ_c^2 -derived estimates. This means that the agreement between the estimates derived by-eye and by the other authors is possibly greater than is implied by the covariance ellipses. The large uncertainties on the differences, coupled with the way that point estimates were derived for the by-eye fits (i.e., that the by-eye point estimate do not identify a singular best fit, but instead marked the centre of the rectangle that encloses the inclusion contour; see Section 10.3), make it difficult to evaluate the trends (for the point estimates) in these two scatterplots, however. Looking instead at the marginal distributions plotted in Figures 10.3 and 10.4 (the brown bars; which take these uncertainties into account), we see evidence that a similar trend toward lower $\Delta(\alpha + \zeta)$ values does, indeed, exist for the by-eye fits for the OG+Cone model. For the TPC+Cone model, these distributions appear to be more centralised about the previously-derived parameter estimates. Therefore, based upon these marginal distributions, one might argue that for the OG+Cone fits the by-eye procedure granted the radio-band GOF higher priority than did the procedure used by the other authors, and that for the TPC+Cone fits it did not. It is difficult to draw such conclusions confidently, though, because the other authors generated the sets of model LC pairs they used using different γ -ray models to those we used. Notably, the marginal distribution for $\Delta|\beta|$ also appears roughly centralised for both dual-band models (just as was the case for the χ_c^2 fits), which indicates that the by-eye fits and the fits conducted by the other authors tend to agree on the value of $|\beta|$.




Figure 10.3: Smoothed distributions for the differences $\Delta\alpha$, $\Delta\zeta$, $\Delta|\beta|$ and $\Delta(\alpha + \zeta)$, between the parameter estimates derived here (via χ^2_c , $\Psi^2_{\Phi,c}$ and the by-eye procedure) and by [Johnson et al. \(2014\)](#) and [Pierbattista et al. \(2015\)](#) for the OG+Cone dual-band model.




Figure 10.4: Smoothed distributions for the differences $\Delta\alpha$, $\Delta\zeta$, $\Delta|\beta|$ and $\Delta(\alpha + \zeta)$, between the parameter estimates derived here (via χ^2_c , $\Psi^2_{\Phi,c}$ and the by-eye procedure) and by [Johnson et al. \(2014\)](#) and [Pierbattista et al. \(2015\)](#) for the TPC+Cone dual-band model.

The covariance ellipses for the fits conducted using $\Psi_{\Phi,c}^2$ (in the right-most column in Figure 10.2) are substantially smaller than those for the fits conducted using χ_c^2 . This indicates that the point estimates derived using $\Psi_{\Phi,c}^2$ are in closer agreement with those derived by the other authors than are the χ_c^2 -derived point estimates. Interestingly, $\Delta\zeta < 0$ for the majority of the OG+Cone fits, and $\Delta\alpha > 0$ for the majority of the TPC+Cone fits. This is likely related to a difference between the model implementations used here and those used by the other authors.

Just as was the case for the χ_c^2 and by-eye fits, the estimates derived for the OG+Cone model using the $\Psi_{\Phi,c}^2$ statistic tend to have lower $(\alpha + \zeta)$ values than those derived by the other authors. This is perhaps further evidence that the implementations used for the OG model yielded markedly different LCs. For the TPC+Cone model, all four marginal distributions are roughly centralised.

The Undominated Estimates vs. the By-eye Estimates

As discussed in Section 10.3, the by-eye fits can be used as a sanity check of sorts. The correlation plots in Figures 10.5 and 10.6 compare the degree to which the estimates derived here using the $\Psi_{\Phi,c}^2$ statistic and the estimates derived by the other authors (the two undominated statistical fits), correlate with the relevant by-eye estimates, for the four quantities α , ζ , $|\beta|$ and $(\alpha + \zeta)$. Each of the filled green dots is associated with a dual-band fit conducted using $\Psi_{\Phi,c}^2$, where higher colour intensities correspond to higher values for $\log(\lambda_{r\gamma})$. Each of the filled purple diamonds is similarly associated with a fit conducted by either Johnson et al. (2014) or Pierbattista et al. (2015). In each case, the horizontal position of a dot or diamond is determined by the value of the relevant quantity for the by-eye fit to that pulsar, and its vertical position is determined by the value for the relevant quantity for the appropriate statistical fit. Thus, the dot and diamond associated with a particular pulsar (connected by a grey line; for a given plot) will have the same horizontal position. Notice that in each case the colour intensities of the dot and diamond are the same, since $\log(\lambda_{r\gamma})$ is model independent.

For the OG+Cone model, the point estimates that the $\Psi_{\Phi,c}^2$ fits yield for α correlate more strongly (than those that the other authors' fits yielded) with the point estimates that the by eye fits yield; the correlation coefficient ρ for the $\Psi_{\Phi,c}^2$ fits is higher by 0.337. For ζ , the correlation is slightly worse for the $\Psi_{\Phi,c}^2$ fits (0.580 vs. 0.596). While these comparisons provide some insight, those for $|\beta|$ and $(\alpha + \zeta)$ are arguably more telling. For the former quantity, the previously-reported estimates correlate more strongly with the by-eye estimates. This indicates that the other authors' fits likely tended to give the radio-band GOF greater relative weight than the by-eye fits effectively did (recall that the by-eye fits were not conducted using a ξ -like statistic, and therefore did not grant the radio band weight via the same mechanism; i.e., via a choice for the relative radio weight ω_r). This is corroborated by the stronger correlation between the $\Psi_{\Phi,c}^2$ estimates and the by-eye estimates for $(\alpha + \zeta)$ (0.892 vs. 0.673), since this indicates that the $\Psi_{\Phi,c}^2$ fits likely tended to give the radio-band GOF lesser relative weight than the by-eye fits effectively did.

A similar situation again exists for the TPC+Cone fits with regard to the correlations for $|\beta|$ and $(\alpha + \zeta)$, but not for the correlations for α and ζ . For both of the latter two quantities, the correlation between the $\Psi_{\Phi,c}^2$ estimates and the by-eye estimates is stronger than between the other authors' estimates and the by-eye estimates.




Figure 10.5: Correlation plots for the OG+Cone dual-band fits that compare the parameter estimates derived by the other authors and by using the $\Psi^2_{\Phi,c}$ statistic to those derived by eye, via the quantities α , ζ , $|\beta|$ and $(\alpha + \zeta)$. On each plot, each pair of connected points (a pair always comprises a filled green circle and a filled purple diamond) corresponds to a single dual-band pulsar. The horizontal position of a pair of points corresponds to the value of the relevant quantity for the by-eye fit, while the vertical positions of the two constituent points correspond to the value derived for the same quantity by the other authors (for the purple points) and using the $\Psi^2_{\Phi,c}$ statistic (for the green points); i.e., to the values derived via the statistical fits. The colour intensity of the points in each pair (and that of the line connecting them) is determined by the value of $|\log(\lambda_{rY})|$ for the relevant pulsar (a more intense colour indicates a higher $|\log(\lambda_{rY})|$, and hence an increase in the potential for single-band dominance).




Figure 10.6: Correlation plots for the TPC+Cone dual-band fits that compare the parameter estimates derived by the other authors and by using the $\Psi^2_{\Phi,c}$ statistic to those derived by eye, via the quantities α , ζ , $|\beta|$ and $(\alpha + \zeta)$. On each plot, each pair of connected points (a pair always comprises a filled green circle and a filled purple diamond) corresponds to a single dual-band pulsar. The horizontal position of a pair of points corresponds to the value of the relevant quantity for the by-eye fit, while the vertical positions of the two constituent points correspond to the value derived for the same quantity by the other authors (for the purple points) and using the $\Psi^2_{\Phi,c}$ statistic (for the green points); i.e., to the values derived via the statistical fits. The colour intensity of the points in each pair (and that of the line connecting them) is determined by the value of $|\log(\lambda_{\text{ry}})|$ for the relevant pulsar (a more intense colour indicates a higher $|\log(\lambda_{\text{ry}})|$, and hence an increase in the potential for single-band dominance).

Chapter 11

Probative Extensions

Alongside the two applications discussed in Chapters 9 and 10 (which have specific bearing on the main thesis of this work), we also conducted two probative extensions of the implementation of the $\Psi_{\Phi,c}^2$ statistic presented in this work. The first of these extends the fitting procedure outlined in Chapter 7 to allow for joint fits (aimed at parameter estimation) that use the FIDO γ -ray model instead of the geometric OG or TPC models. The second extends (or attempts to extend) the domain over which the $\Psi_{\Phi,c}^2$ statistic can be applied, beyond joint pulsar LC fits, to joint fits to datasets associated with pulsar wind nebulae (PWNe).

Both of these extensions came about as the result of collaborative research opportunities, similar to the fits conducted for PSR J2039–5617 in Chapter 9. These opportunities, however, presented themselves somewhat earlier than that associated with PSR J2039–5617, before the development of the $\Psi_{\Phi,c}^2$ statistic had reached its conclusion. This means that an earlier incarnation of the $\Psi_{\Phi,c}^2$ statistic (specifically, the $\chi_{c,\text{sfn}}^2$ statistic defined in Appendix C) was used in the work presented in this chapter. Appendix C gives a brief account of the evolution of the $\Psi_{\Phi,c}^2$ statistic throughout this research, and lists each of its earlier incarnations.

11.1 Dual-band LC Fits using a FIDO γ -ray Model: Magnetospheric Conductivity and Spin-down Luminosity

Brambilla et al. (2015), conducting γ -ray-only fits to the LCs and phase-resolved spectra of eight bright *Fermi* LAT pulsars using the FIDO γ -ray emission model (Kalapotharakos et al., 2014, see Section 1.5), reported an apparent correlation between the best-fit value for the macroscopic conductivity σ (measured in units of the pulsar’s rotation frequency $\Omega = 1/P$) outside the light cylinder and the pulsar’s spin-down luminosity $\dot{E}$ (roughly speaking, the rate at which the pulsar’s rotational energy is dissipated via electromagnetic emission and particle acceleration; typically measured in ergs/s). Since we had access to a set of FIDO model LCs, this presented an ideal opportunity to extend the fitting procedure developed for this work to fits that involve the FIDO model. We therefore set out to investigate this possible correlation via γ -ray-only and joint fits to 2PC pulsar LCs (Abdo et al., 2013) using the FIDO γ -ray model and the conglomerated dual-band model constituted by the Cone radio model and the FIDO γ -ray model. We note that the Cone model is based on a different magnetic field geometry (the retarded vacuum dipole) than the FIDO model. This means that the conglomerated dual-band model is not entirely consistent, but since these two magnetic field geometries do not differ substantially at low altitudes, where the radio emission originates, this inconsistency can be tolerated for the purposes of this probative study.

The set of FIDO model LCs The code that was used to generate the set of model LCs we used effectively takes five parameters: α , ζ , P , B_{surf} (the magnetude of the magnetic field at the NS surface) and σ . The set of model LCs we used corresponds to the output of this code given the paramater combinations defined by the following values for each parameter:

- Eighteen values for α , that cover the interval $(0^\circ, 90^\circ]$ at a 5° resolution,
- One hundred values for ζ , that cover the interval $[0^\circ, 180^\circ)$ and correspond to evenly spaced $\cos(\zeta)$ values.

- Five values for P : $P \in \{0.01 \text{ s}, 0.03 \text{ s}, 0.1 \text{ s}, 0.3 \text{ s}, 1 \text{ s}\}$,
- Six values for B_{surf} : $B_{\text{surf}}/10^{12} \text{ G} \in \{0.5, 1, 2, 4, 8, 30\}$, and
- Five values for σ : $\sigma \in \{1 \Omega, 10 \Omega, 100 \Omega, 1000 \Omega, 10000 \Omega\}$.

Conducting the fits

The first step was to determine whether or not the conglomeration of these two single-band models is feasible, specifically in light of the fact that (unlike the OG and TPC models) the Cone and FIDO models assume different magnetic fields. Since the magnetic field close to the NS surface (where radio emission occurs) approximates a dipolar field in both the retarded vacuum dipole and force-free magnetic fields (the two magnetic fields relevant to this conglomeration at low altitudes), the conglomeration was deemed to indeed be reasonable.

The second step was to identify the 2PC pulsars for which fits would be conducted. Since we did not intend to conduct any by-eye fits, and since we had already developed an automated LC fitting pipeline to conduct the joint fits in Chapter 10, we decided to initially include all 117 2PC pulsars. Once the first round of fits had been concluded, however, we discovered that $\chi^2_\gamma < \nu_\gamma$ (the number of degrees of freedom associated with the γ -ray fit) for a large proportion of the γ -ray-only fits, and $\chi^2_\gamma \simeq 0$ for a few. Upon further investigation, we discovered that this apparent overfitting occurred most frequently for the pulsars with the lowest precision $\rho_\gamma = \Phi_\gamma^2/n_\gamma$ in their γ -ray LCs (see Section 4.2). Therefore, for the final round of fits (both γ -ray-only and joint) we excluded those pulsars for which $\rho_\gamma < 10$. Figure 11.1 demonstrates this trend for the statistic

$$\chi_{\text{sfn},\gamma}^2(M_\gamma) = \frac{\Phi_\gamma^2 - \chi_\gamma^2(M_\gamma)}{\Phi_\gamma^2 - \nu_\gamma} = \frac{\Phi_\gamma^2}{\Phi_\gamma^2 - \nu_\gamma} \Psi_{\Phi,\gamma}^2(M_\gamma), \quad (11.1)$$

a previous incarnation of $\Psi_{\Phi,\gamma}^2$ (see Appendix C). A useful attribute of this statistic was that $\chi_{\text{sfn},\gamma}^2 = 1$ implied that $\chi_\gamma^2 = \nu_\gamma$. Hence, a value of $\chi_{\text{sfn},\gamma}^2 > 1$ indicated that overfitting was taking place.¹

Results

The cut-off value (of 10) for ρ_γ (indicated by the vertical line in Figure 11.1) was decided upon by inspection of Figure 11.1, and left 65 pulsars to fit, of which only 34 were radio-loud. Figure 11.2 shows the results of the subsequent round of fits (with regard to the σ - $\dot{E}$ correlation). The results for the γ -ray-only fits are plotted in the upper panel, while results for the joint fits are plotted in the lower panel. In both panels, the blue circles indicate the results of the fits conducted by Brambilla et al. (2015) (with a clear upward trend in σ with increasing $\dot{E}$) and the black dots indicate the results of our fits. Note that the fits conducted by Brambilla et al. were exclusively γ -ray-only, and were not only based on LC fits (they also fit phase-resolved spectra, as mentioned earlier).

Upon comparing the distribution of the best fits we obtained to the trend apparent for the fits conducted by Brambilla et al. (2015), it becomes clear that the σ - $\dot{E}$ correlation did not manifest for our fits. The reason for this is probably that we only probed the σ - $\dot{E}$ correlation using LC shapes (that result mostly from spatial emission geometry and beaming), while Brambilla et al. (2015) obtained joint fits to LCs and spectra. The latter (especially the cut-offs) are more directly linked to the particle energetics, which are in turn influenced by σ since it determines $E_\parallel$ and hence the details of particle acceleration. Furthermore, the low resolution in σ associated with the set of model LCs we used may have hampered our search for a correlation since the shape of the FIDO model LCs is more sensitive to different values of σ when the value of σ is relatively low.

11.2 Joint Fits to Spectral and Spatial Data for Three Pulsar Wind Nebulae

As part of an on-going research effort, we and our collaborators conducted joint fits to spectral and spatial data for three pulsar wind nebulae (PWNe; PWN 3C 58, PWN G21.5–0.9 and PWN G0.9+0.1) using a spatiotemporal leptonic transport

¹Note that within the context of the vector space framework and the conceptual design presented in Chapter 7 (both of which came into being only after the definition of $\chi_{\text{sfn},\gamma}^2$), overfitting cannot occur since the observed uncertainties are assumed to be a true representation of the uncertainty in the identity of D^{true} .

code, developed by [van Rensburg et al. \(2018\)](#). The results of these fits are presented (and discussed in greater detail) in [van Rensburg et al. \(2020\)](#). The account given here focuses solely on the role played by the $\Psi_{\Phi,c}^2$ statistic, and on its performance for these fits.

The three PWNe listed above were selected (by our collaborators) for these joint fits because for each of them the following three observables are available: the broad-band spectral energy distribution (SED) and the X-ray photon index (XPI) and surface brightness (XSB) profiles (as function of radius, with respect to the geometric centre of the PWN on the sky). These datasets are plotted (in the form of scatterplots with errorbars) for each PWN in Figures 11.3, 11.4 and 11.5. The model used in [van Rensburg et al. \(2020\)](#) (an incarnation of the more general model described in [van Rensburg et al., 2018](#)) has a total of five free parameters (upon fixing a number of other parameters):² the present-day magnetic field strength B_{age} , the bulk flow normalisation V_0 , the age of the PWN t_{age} , the diffusion coefficient κ_0 and the radial magnetic field profile α_B . Just as is the case for the pulsar LC fits, the inclusion of multiple datasets in these fits promised to increase their constraining power substantially. Additionally, similar to the pulsar case, the three datasets associated each PWN (an SED, XPI and XSB) were not all equally precise. Unlike the pulsar fits, however, it was not feasible to generate an exhaustive set of model predictions for the three observables listed above since the parameter space of this model is very large. Therefore, our collaborators opted to conduct their fits, and therefore derive parameter point estimates, via a Markov chain Monte Carlo (MCMC) approach, wherein (for each joint fit) the model's parameter space is explored dynamically (by leveraging some GOF statistic; χ_c^2 in this case, defined to be the sum of the three single-dataset χ^2 statistics) to find the parameter combination that produces the best joint fit to the observed datasets.

Crucially, just as was the case for pulsar parameter estimation, this exploration can only be effective if the GOF statistic it relies on does, indeed, effectively characterise GOF. Upon the conclusion of the first round of joint fits, our collaborators realised that their fits may be susceptible to something akin to the single-band dominance that manifests for dual-band pulsar fits. Specifically, it may be the case that one of the datasets is significantly more precise than the other two (or significantly less precise). If this were the case and the GOF statistic the MCMC algorithm was equipped with were prone to granting each dataset's GOF value priority in accordance with its precision (as χ_c^2 does for χ_r^2 and χ_v^2 ; see Section 4.4), then some sort of single-dataset– or two-dataset–dominance might manifest. Additionally, upon a by-eye exploration of the model's parameter space, our collaborators were able to identify parameter combinations for each PWN that yielded fits that are arguably better than the MCMC-derived best fits, but that markedly differed from the MCMC-derived parameter point estimates. Given the risk that the MCMC-derived point estimates (and the by-eye–derived point estimates) might be biased, and the tentative evidence that they were, our collaborators enlisted our help, with the hope that the MCMC algorithm may be equipped with a statistic similar to $\Psi_{\Phi,c}^2$ (or, similar to the then-current incarnation, $\chi_{c,\text{sfn}}^2$; see Appendix C).

Adapting the $\chi_{c,\text{sfn}}^2$ Statistic for PWN Fits

Since there are three datasets involved in each fit (the SED, XPI and XSB), we needed to define four statistics: a single-dataset–only statistic for each dataset and a joint statistic. The first step in doing this was the identification of suitable equivalents to the background LCs in the pulsar case. This proved to be comparatively difficult, since the datasets being fit are somewhat less easily framed as perturbations (than pulsar LC are).³ After some trial and error, we decided to adopt the following standard for the background levels of the SED, XPI and XSB: 1% of the SED level for the SED fits, a constant value of 3 for the XPI fits and a constant value of 0.01 for the SB fits (in the appropriate units). Fortunately, for similar choices of the background levels, the parameter estimates the MCMC algorithm yields do not change substantially from those it yields with these background levels.

With these background levels in hand, we defined the three background-equivalent datasets B_{SED} , B_{XPI} and B_{XSB}

²Note that this model is not a conglomeration of three single-dataset models.

³At the time, we had only just started framing LC fitting in terms of perturbations, and had not yet realised the important role such a framing ultimately plays.

Table 11.1: Single-dataset-only and joint GOF values for the fits to the SED, X-ray photon index (XPI) and surface-brightness profile (XSB) of PWN 3C 58. For each statistic, a higher value corresponds to an increased GOF. For the single-dataset-only fits, the GOF value for the dataset being fit is printed in bold.

Fit	$\chi_{\text{sfn,SED}}^2$	$\chi_{\text{sfn,XPI}}^2$	$\chi_{\text{sfn,XSB}}^2$	$\chi_{\text{c,sfn}}^2$
SED-only	0.99	0.91	0.61	0.837
XPI-only	0.53	0.94	0.45	0.640
XSB-only	0.10	0.88	0.93	0.637
Joint: $\chi_{\text{c,sfn}}^2$	0.99	0.92	0.92	0.943
Joint: By-eye	0.99	0.94	0.72	0.883

analogously to B_r and B_γ , and defined the three single-dataset-only statistics

$$\chi_{\text{sfn,SED}}^2(M_{\text{SED}}) = \frac{\Phi_{\text{SED}}^2 - \chi_{\text{SED}}^2(M_{\text{SED}})}{\Phi_{\text{SED}}^2 - \nu_{\text{SED}}} \quad (11.2)$$

$$\chi_{\text{sfn,XPI}}^2(M_{\text{XPI}}) = \frac{\Phi_{\text{XPI}}^2 - \chi_{\text{XPI}}^2(M_{\text{XPI}})}{\Phi_{\text{XPI}}^2 - \nu_{\text{XPI}}} \quad (11.3)$$

$$\chi_{\text{sfn,XSB}}^2(M_{\text{XSB}}) = \frac{\Phi_{\text{XSB}}^2 - \chi_{\text{XSB}}^2(M_{\text{XSB}})}{\Phi_{\text{XSB}}^2 - \nu_{\text{XSB}}}, \quad (11.4)$$

where $\Phi_{\text{SED}}^2 = \chi_{\text{SED}}^2(B_{\text{SED}})$, $\Phi_{\text{XPI}}^2 = \chi_{\text{XPI}}^2(B_{\text{XPI}})$ and $\Phi_{\text{XSB}}^2 = \chi_{\text{XSB}}^2(B_{\text{XSB}})$ and ν_{SED} , ν_{XPI} and ν_{XSB} are the scaled fluxes and degrees of freedom associated with the SED, XPI and XSB datasets, respectively. These three single-dataset-only statistics, which are each 0 for a background-equivalent fit and 1 for a fit where $\chi^2 = \nu$, allowed for the natural definition of the joint fit statistic

$$\chi_{\text{c,sfn}}^2(M_c) = \frac{1}{3} \chi_{\text{sfn,SED}}^2(M_{\text{SED}}) + \frac{1}{3} \chi_{\text{sfn,XPI}}^2(M_{\text{XPI}}) + \frac{1}{3} \chi_{\text{sfn,XSB}}^2(M_{\text{XSB}}), \quad (11.5)$$

where $M_c(\theta) = (M_{\text{SED}}, M_{\text{XPI}}, M_{\text{XSB}})$ is the profile triplet associated with the model parameter combination θ . Since each of the single-dataset-only statistics are 0 for a background-equivalent fit, $\chi_{\text{c,sfn}}^2 = 0$ for a background-equivalent fit, and since each single-dataset-only statistic is $\simeq 1$ for a good fit (i.e., a fit for which $\chi^2 \simeq \nu$), $\chi_{\text{c,sfn}}^2 \simeq 1$ for a good joint fit (i.e., $\chi_{\text{c,sfn}}^2 \simeq 1$ for a joint fit where all three single-dataset-only fits are good fits).

Results

The best-fit profile triplets for each of the three fitting procedures (MCMC with χ_c^2 , MCMC with $\chi_{\text{c,sfn}}^2$ and by eye) are plotted for each PWN in Figures 11.3, 11.4 and 11.5. As might be expected given the size of the model's parameter space, all three best-fit profile triplets differ in all three cases. However, for none of the PWNe was one best-fit profile triplet clearly better than the others; invariably, a compromise between the various datasets apparently had to be reached. This indicated that the MCMC algorithm did not fail to find a good joint fit because it was equipped with a biased joint GOF statistic, but because the model likely does not produce a good joint fit anywhere in its parameter space.

Model Degeneracy and Single-dataset Best-fit Non-colocation

The apparent necessity for compromise between the three datasets suggests that the single-dataset best-fit parameter combinations for each PWN are likely non-located in the model's parameter space (like the single-band best-fit parameter combinations tended to be for the pulsar fits). Unfortunately, since there are three datasets involved, the tools developed in Chapter 6 cannot be simply extended to characterise these fits. However, we may be able to gain some insight into the character of these fits by performing a single-dataset-only fit for each dataset and comparing the resulting single-dataset-only GOF values to the GOF values for the joint fits; to this end, we performed single-dataset-only fits for PWN 3C 58.

Table 11.1 lists the single-dataset-only and joint GOF values for the three single-dataset-only fits and by-eye and $\chi^2_{c,sfn}$ joint fits. In each column, the highest GOF value is indicated in bold. As expected, each of the GOF values that lie on the main diagonal of the table is bolded (because the SED-only fit maximises the value of $\chi^2_{sfn,SED}$ etc.). For each of the single-dataset-only fits, however, the GOF values that do not lie on the main diagonal are markedly lower than the corresponding bolded GOF values in their columns; e.g., the GOF value for the XSB dataset is substantially lower for the XPI-only fit (0.45) than for the XSB-only fit (0.93; the $\chi^2_{sfn,XSB}$ column's maximum). This indicates that the single-dataset-only parameter estimates likely do not lie near each other in the model's parameter space; i.e., that there is significant single-dataset best-fit non-colocation. This means that the three-dataset analogue of the perimeter of the set of model-predicted LC pairs in Figure 6.1 likely comprises a larger number of model-predicted profile triplets, and that the MCMC would therefore yield a wide variety of best-fit profile triplets when equipped with joint GOF statistics that cover a wide variety of relative weightings for the three datasets; i.e., that there is a degeneracy in the identity of the models best fit. Just as was the case for the dual-band pulsar models, this degeneracy indicates that further model refinement is likely needed, since no statistic would find a profile triplet comprised of fits that are as good as the respective single-dataset-only best fits.

That being said, the $\chi^2_{c,sfn}$ fit does a commendable job, finding a profile triplet where the single-dataset-only GOF values deviate no more than 0.2 from those for the single-dataset-only fits (the average loss in GOF across the three datasets, or the value of the three-dataset analogue of κ , is only 0.01). Notably, for both the by-eye and $\chi^2_{c,sfn}$ fits, the GOF for the SED dataset is maximal.




Figure 11.1: Plot of γ -ray-band GOF (as measured using $\chi^2_{\text{sfm},\gamma}$) as function of γ -ray-band LC sensitivity for the first round of γ -ray-only FIDO fits. The value of $\chi^2_\gamma > \nu_\gamma$ for all the fits that lie below the thin horizontal line, while for all the fits that lie above it $\chi^2_\gamma < \nu_\gamma$ (which indicates that overfitting occurred). Since overfitting apparently occurred more frequently for pulsar with less precise observed LCs (i.e., lower ρ_γ), we excluded all the 2PC pulsars for which $\rho_\gamma < 10$ from the second round of fits.




Figure 11.2: Plots of best-fit σ vs. $\dot{E}$ for the γ -ray-only (*top* panel) and joint (*bottom* panel) FIDO fits. The black dots correspond to the results for our fits, and the blue circles to the results reported by [Brambilla et al. \(2015\)](#). Notice that σ tends to increase as $\dot{E}$ increases for the blue circles; this is the correlation [Brambilla et al. \(2015\)](#) reported.




Figure 11.3: SED, XPI and XSB fits for PWN 3C 58. For these plots, “ χ^2 ” means χ^2_c and “SFN” means $\chi^2_{c,\text{sfn}}$. The SED for this PWN comprises radio, X-ray and γ -ray observations. For more detail, see [van Rensburg et al. \(2020\)](#).




Figure 11.4: SED, XPI and XSB fits for PWN G21.5–0.9. For these plots, “ χ^2 ” means χ_c^2 and “SFN” means $\chi_{c,\text{sfn}}^2$. The SED for this PWN comprises radio, X-ray and γ -ray observations. For more detail, see [van Rensburg et al. \(2020\)](#).




Figure 11.5: SED, XPI and XSB fits for PWN G09+0.1. For these plots, “ χ^2 ” means χ^2_c and “SFN” means $\chi^2_{c,sfn}$. The SED for this PWN comprises radio, X-ray and γ -ray observations. For this PWN an additional “constant age” fit was performed, which we will not discuss here. For more detail, see [van Rensburg et al. \(2020\)](#).

Summary and Conclusion

As stated in the [Introduction](#) to this work, the overarching aim of this thesis is to demonstrate that the $\Psi_{\Phi,c}^2$ statistic is a more appropriate extension to Pearson's χ^2 statistic to joint, dual-band LC fits aimed at pulsar parameter estimation than is the χ_c^2 statistic. In Part I, we first gave a brief overview of the field of pulsar astrophysics and introduced the geometric models we would use in Parts III and IV. Then, we introduced Pearson's χ^2 statistic within the context of pulsar LC fitting (both its single- and dual-band incarnations), and examined their performance for single- and dual-band fits aimed at deriving pulsar parameter estimates more closely. For dual-band fits, we established that the χ_c^2 -equipped parameter estimation procedure is prone to single-band dominance in cases where the single-band best-fit parameter combinations are non-colocated in the dual-band model's parameter space and there exists a large uncertainty disparity between the observed LCs associated with the two bands.

In Part II, we took this examination further, and demonstrated that the χ_c^2 -equipped parameter estimation procedure does, indeed, fail to effectively derive meaningful parameter constraints for such single-band-dominated fits (in Section 4.1; via the examination of a real-world dual-band fit). This demonstration provided little insight into the reason for this failure. To gain this insight, we reframed the pulsar fitting problem with the perturbative nature of the observed LCs in mind. Specifically, we identified the scaled flux as an important quantity, and converted it into the scaled-flux normalised Ψ_{Φ}^2 statistic for single-band fits and the scaled-flux standardised $\Psi_{\Phi,c}^2$ statistic for joint, dual-band fits. In Chapter 5 we refined these ideas yet further, and constructed a rigorous mathematical foundation for the framework established in Chapter 4. Leveraging this framework, we discovered that the statistics that we had defined in Chapter 4 arise quite naturally given simple constraints; most notably, the constraint that the scaled flux must be standardised across the two wavebands in the case of $\Psi_{\Phi,c}^2$.

Next, we turned our attention toward gauging the effectiveness of the χ_c^2 and $\Psi_{\Phi,c}^2$ statistics (with regard to parameter estimation). In Chapter 6 we developed a suite of mathematical tools that would allow us to determine whether or not a given dual-band statistic yield joint fits that are single-band dominated. In Chapter 8, we leverage these tools for a set of synthetic observed LC pairs, and show that not only does the $\Psi_{\Phi,c}^2$ statistic yield the same, undominated parameter estimates regardless of the magnitude of the uncertainty disparity, but also that the χ_c^2 -derived parameter estimates converge to those derived using $\Psi_{\Phi,c}^2$. In Chapter 10, we did the same for a set of 23 *Fermi* LAT pulsars, and observed a similar convergence. Beyond this convergence, we also confirmed that the best-fit LC pairs that the $\Psi_{\Phi,c}^2$ statistic yielded do, indeed, more effectively reproduce the broad structure of the observed LC pairs.

The results of the tests conducted in Chapters 8 and 10 (specifically, the characterisations yielded by the tools developed in Chapter 6; reported in Table 10.5 and plotted in Figure 10.1) demonstrate that the $\Psi_{\Phi,c}^2$ statistic is, indeed, a more appropriate extension of Pearson's χ^2 statistic to joint, dual-band LC fits aimed at pulsar parameter estimation than is the χ_c^2 statistic. It should be noted here that this conclusion is made solely within the context of the framework developed in Chapter 6 (for understanding what single-band dominance is), and is contingent on neither the results of the [Qualitative By-eye Evaluation of the Best-fit LC Pairs](#) that the $\Psi_{\Phi,c}^2$ statistic yields (conducted in Section 10.4.2) nor the degree to which the associated dual-band parameter estimates correlate with those derived by eye (assessed in Section 10.4.3). The former of these merely serves as a sanity check to confirm that the improved performance of the parameter estimation procedure when equipped with the $\Psi_{\Phi,c}^2$ statistic (as per the tools developed in Chapter 6) does, indeed, entail an improvement in the perceived quality of the joint fits. Such a confirmation is informative since the results obtained using the tools developed in Chapter 6 most directly address an abstracted version of the [Problem Statement](#) (implicitly defined in Section 4.1): that the parameter estimation procedure is equipped with a GOF statistic that does not satisfy Eq. (4.1). Similarly, the latter (along

with the comparison of the $\Psi_{\Phi,c}^2$ -derived parameter estimates with those derived by other authors) serves as a sanity check to confirm that the results obtained via the introduction of the $\Psi_{\Phi,c}^2$ statistic are similar to those obtained using independent, arguably equally effective, radio-dominance mitigation strategies.

A Broader Perspective

With the effectiveness of the introduction of the $\Psi_{\Phi,c}^2$ established (with regard to its elimination of radio dominance), it is perhaps prudent to take a step back and more thoroughly consider how the $\Psi_{\Phi,c}^2$ statistic fits into the broader context of pulsar LC fitting, and indeed pulsar astrophysics as a whole.

First and foremost it should be reiterated that the $\Psi_{\Phi,c}^2$ statistic is intended specifically for use in dual-band fits aimed at parameter estimation, and that it in no way supplants the χ_c^2 statistic where GOF testing is concerned. Moreover, the only such dual-band fits for which the $\Psi_{\Phi,c}^2$ statistic is more appropriate are those where the model's best-fit LC pair does not pass the χ^2 GOF test, since it is only for those fits that different conceptions of goodness yield different, non-zero estimates for the true deviation (see Section 7.4.1). Of course, since these two statistics converge in their effect as the uncertainty disparity becomes smaller, the question of which is more appropriate only becomes important when there is a large uncertainty disparity; i.e., an uncertainty disparity that is large enough that radio dominance can manifest.

Second it should be stressed that, while the introduction of the $\Psi_{\Phi,c}^2$ statistic effectively eliminates radio dominance (or, more generally, single-band dominance) in parameter estimation, it is not necessarily the case that the $\Psi_{\Phi,c}^2$ statistic is always the most appropriate ξ -like statistic (see Section 6.1) to use when deriving parameter estimates (using strictly invalid models). In fact, within the context of the discussion on parameter estimate accuracy conducted in Section 3.2.2, the $\Psi_{\Phi,c}^2$ statistic is almost certainly not the most appropriate ξ -like statistic to use for any dual-band fit; the most appropriate such statistic would be the one that grants the waveband associated with the more accurate single-band-only parameter estimate proportionally more weight. However, since it is not possible to know which single-band-only parameter estimate is more accurate based solely on LC fitting, it is not possible in general to determine which ξ -like statistic is, in fact, most appropriate.

There exists, of course, a multitude of possible heuristics that aim to answer the question of which single-band-only estimate is more accurate, but these are all outside the scope of this work.

Future Work

The general nature of the framework we developed in Chapter 5 means that the Ψ_{Φ}^2 and $\Psi_{\Phi,c}^2$ statistics are likely very widely applicable. In the near-term, however, the most prudent next steps to take are to expand the reach of the $\Psi_{\Phi,c}^2$ statistic to more than two wavebands and to the most current pulsar emission models. The first, tentative steps in these directions have already been taken, as documented in Chapter 11.



Another potential avenue is to investigate the cause for the residual $\log(\lambda_{r\gamma})$ -dependence in the radio priority observed for the synthetic fits (reported in Section 8.3). It might be the case that an even more appropriate statistic (or family of statistics; for parameter estimation) is one that treats the on-pulse and off-pulse regions differently when determining GOF.



APPENDICES



Appendix A



Best-fit LC pairs for 23 *Fermi* LAT 2PC pulsars



Each plot in this appendix comprises five LC pairs: the observed LC pair in grey (with the LC uncertainties plotted as black errorbars), the best-fit LC pairs identified for the OG+Cone model using χ_c^2 and $\Psi_{\Phi,c}^2$ (the thin, light brown line and dark brown line, respectively) and the best-fit LC pairs identified for the TPC+Cone model using χ_c^2 and $\Psi_{\Phi,c}^2$ (the thin, light green line and dark green line, respectively). The dark grey dashed horizontal lines correspond to the radio and γ -ray background levels, and the light grey dashed horizontal line on the γ -ray plots corresponds to the minimum γ -ray LC amplitude (see Section 7.3). As a reference, a dotted vertical line is plotted for each best-fit LC pair at the value of ϕ_μ for that fit (in the same colour). Since ϕ_μ can be the same for two LC pairs (and these vertical lines can therefore overlap), the best-fit point estimates are listed in the legend to each figure, in the form of an $(\alpha, \zeta, \phi_\mu)$ triplet.



PSR J0437–4715 $\log(\lambda_{\text{r}\gamma}) = 2.41$ **PSR J0908–4913** $\log(\lambda_{\text{r}\gamma}) = 2.36$ 






































Appendix B


Maps of $\Psi_{\Phi,c}^2$ for 23 *Fermi* LAT 2PC Pulsars




Each figure in this appendix contains two $\Psi_{\Phi,c}^2$ maps; one for the OG+Cone fit and one for the TPC+Cone fit. On each map, the point estimates for α and ζ that result from the $\Psi_{\Phi,c}^2$ fit are indicated with a vertical and a horizontal dark grey line, and the estimates that result from the χ_c^2 fit are indicated with similar grey dashed lines. The contours that demarcate the 2σ confidence regions for the various fits are coloured as follows (if they exist; i.e., no minimum has been enforced yet for the uncertainty on the parameter estimates): orange for the χ_c^2 confidence region, white for the $\Psi_{\Phi,c}^2$ confidence region if $\Psi_{\Phi,c}^2 > 0.5$ and black for the $\Psi_{\Phi,c}^2$ confidence region if $\Psi_{\Phi,c}^2 \leq 0.5$.


Note that the main colourmap only covers the range $0 \leq \Psi_{\Phi,c}^2 \leq 1$, which means that those parameter combinations associated with fits for which $\Psi_{\Phi,c}^2 < 0$ are not coloured (as well as parameter combinations where either component single-band model predicts no pulsed emission). If $\Psi_{\Phi,c}^2 < 0$ for the $\Psi_{\Phi,c}^2$ best-fit LC pair, then a greyscale colourmap is used instead. Each figure has two colourbars, one for the OG+Cone map (left) and one for the TPC+Cone map (right). On each, the $\Psi_{\Phi,c}^2$ values for the $\Psi_{\Phi,c}^2$ - and χ_c^2 -identified best-fit LC pairs are plotted as a white diamond and an orange circle (if those values are greater than 0). For the $\Psi_{\Phi,c}^2$ fit, the value of $\delta_{2\sigma,\text{sfs}}$ (the critical value that gives rise to the $\Psi_{\Phi,c}^2$ contours; see Section 7.4.2) is plotted as a horizontal white line on the colourbar (if there are contours and if $\delta_{2\sigma,\text{sfs}} > 0$). If the $\Psi_{\Phi,c}^2$ value for the $\Psi_{\Phi,c}^2$ -identified best-fit LC pair is less than 0.5, then both the diamond and horizontal line on the colourbar is black (to ensure that they stand out).

























Appendix C

The Evolution of the $\Psi_{\Phi,c}^2$ Statistic

As might be expected, the (hopefully) logically consistent account given in the main body of this thesis (in Chapters 4 and 5) for the development of the $\Psi_{\Phi,c}^2$ statistic does not accurately represent the true, somewhat messy, development timeline. To give the reader a somewhat more accurate picture of the real-world process that led to the eventual creation of the $\Psi_{\Phi,c}^2$ statistic, this appendix details the various forms our proposed statistic took prior to what is presented in this work.

C.1 Blazing the Trail

The concept that emerged earliest was that the main issue behind the χ_c^2 statistic's failure was that the two single-band χ^2 statistics do not express GOF in the same units. At this early stage, however, we had neither realised that GOF and DFO are separate from each other nor realised the perturbative nature of LCs (and had most especially not yet identified Φ^2 as an important quantity). What we did know was that χ_r^2 was typically much larger for the worst radio-only fit than χ_γ^2 was for the worst γ -ray-only fit. This led to the definition of the “dynamic range scaled χ_c^2 ” statistic

$$\chi_{c,\text{drs}}^2(M_c) = \left(\frac{\chi_{\gamma,\text{max}}^2 - \chi_{\gamma,\text{min}}^2}{\chi_{r,\text{max}}^2 - \chi_{r,\text{min}}^2} \right) \chi_r^2(M_r) + \chi_\gamma^2(M_\gamma), \quad (\text{C.1})$$

where the differences in the numerator and denominator of the fraction in brackets are the γ -ray and radio dynamic ranges, respectively. The effect of this scaling was that the dynamic ranges of the radio and γ -ray components of the joint GOF was equal; or, that the worst radio-band fit was just as much worse than the best radio-band fit as the worst γ -ray-band fit was worse than the best γ -ray-band fit.

Upon application, this statistic seemed to yield fits that were no longer radio dominated. However, it soon became apparent that this statistic had a serious drawback: Its reliance on the GOF of the worst model fit meant that the eventual result of a given joint fit was sensitive to how good the worst single-band fits were. We were, in essence, in need of a more robust way to determine what change in χ_r^2 is equivalent to a given change in χ_γ^2 . It was in trying to solve this issue that the utility of the background LC became apparent. We reasoned that the background LC B_r was just as bad a fit to D_r as B_γ was to D_γ since a radio model whose best-fit LC was B_r was just as useless as a γ -ray model whose best-fit LC was B_γ (since both models wholly fail to predict the presence of a pulsar). This birthed the comparatively short-lived “modified dynamic range scaled χ_c^2 ” statistic

$$\chi_{c,\text{mdrs}}^2(M_c) = \left(\frac{\Phi_\gamma^2 - \chi_{\gamma,\text{min}}^2}{\Phi_r^2 - \chi_{r,\text{min}}^2} \right) \chi_r^2(M_r) + \chi_\gamma^2(M_\gamma). \quad (\text{C.2})$$

While this statistic proved to be considerably more robust than $\chi_{c,\text{drs}}^2$, it was quickly superseded by the statistic

$$\chi_{c,\text{sfm}}^2(M_c) = \frac{1}{2} \chi_{\text{sfm},r}^2(M_r) + \frac{1}{2} \chi_{\text{sfm},\gamma}^2(M_\gamma), \quad (\text{C.3})$$

where

$$\chi_{\text{sfn},r}^2(M_r) = \frac{\Phi_r^2 - \chi_r^2(M_r)}{\Phi_r^2 - \nu_r} \quad \text{and} \quad \chi_{\text{sfn},\gamma}^2(M_\gamma) = \frac{\Phi_\gamma^2 - \chi_\gamma^2(M_\gamma)}{\Phi_\gamma^2 - \nu_\gamma} \quad (\text{C.4})$$

are the associated single-band statistics (instances of the general single-band statistic χ_{sfn}^2), which was the first incarnation of the $\Psi_{\Phi,c}^2$ statistic.¹ At this point in time, however, we had not yet realised that what we were in fact trying to do was to *standardise* GOF across the two wavebands, and instead called it the “scaled-flux *normalised* χ_c^2 ” statistic, since in each waveband (to our mind at least) the statistic value was normalised (i.e., $\chi_{\text{sfn},r}^2(M_r)$ and $\chi_{\text{sfn},\gamma}^2(M_\gamma)$ were normalised).

Two core ideas led to the definition of $\chi_{\text{c,sfn}}^2$. The first was that, in principle, the χ^2 value associated with the background LC was not the only χ^2 value that could act as a fiducial point between the two wavebands. Since the expected value of χ^2 was ν for a model’s best-fit LC, ν_r and ν_γ surely correspond to equal levels of goodness, just like Φ_r^2 and Φ_γ^2 do. This insight gave rise to the form of the denominators in Eq. (C.4), which are, in essence, a refined incarnation of the dynamic range initially envisioned. The second core idea was that χ_c^2 measures GOF indirectly; i.e., that we could make $\chi_{\text{c,sfn}}^2$ into a true, direct GOF statistic by subtracting each of the component χ^2 values (χ_r^2 and χ_γ^2) from their associated scaled fluxes (Φ_r^2 and Φ_γ^2) before dividing them by the dynamic ranges. This idea, therefore, gave rise to the form of the numerators in Eq. (C.4).

The $\chi_{\text{c,sfn}}^2$ statistic proved to be remarkably robust, and consistently yielded joint fits that were not radio dominated (or, at least did not seem to be radio dominated). It was at this point that we started work on the paper aimed at introducing our new statistic and documenting its development. It is worth noting here that we had been developing the Monte Carlo algorithm, along with the other concepts outlined in Chapter 7, since the definition of $\chi_{\text{c,drs}}^2$, and had performed fits to 16 of the 23 pulsars fit in Chapter 10. These fits were initially intended as the only test of our new statistic (we had not yet conceived of any of the tools developed in Chapter 6, and had not yet had the idea to perform a test using synthetic observed LC pairs).

C.2 Crystalising our Understanding

While the quality of the joint fits conducted using the $\chi_{\text{c,sfn}}^2$ statistic could arguably have sufficed as evidence that it was indeed the statistic we were looking for, we realised in the writing of our first paper that the reasoning by which we arrived at it was fraught with murkily defined concepts and supported by borderline hand-waving argumentation. This forced us to bite the bullet and convert the novel geometric interpretation we had been toying with (of pulsar LC fitting) into a fully fledged mathematical framework within which to derive our statistic from first principles. What resulted was, in essence, the bulk of Part II of this thesis. Of course, the $\chi_{\text{c,sfn}}^2$ statistic did not survive this test, and was eventually supplanted by the $\Psi_{\Phi,c}^2$ statistic.

The first component of the $\chi_{\text{c,sfn}}^2$ statistic to fail was the form of the denominators. Stated most plainly, it became evident quite deep into the development of the vector space framework that the geometric interpretation for each of the terms in Eq. (C.3) did not extend elegantly to the statistic as a whole. Roughly speaking, when written in terms of $(\Phi_c^2 - \nu_c)$, the $\chi_{\text{c,sfn}}^2$ statistic does not resemble the two single-band components, which indicated that something was wrong. This came through most strongly when we tried to determine, geometrically, what the expected value for $\chi_{\text{c,sfn}}^2$ might be. While the answer to this question never came (although we now suspect that $\Psi_{\Phi,c}^2$ is Gamma-distributed in some fashion for $M_{\text{c,true}}$), asking it was sufficient for us to realise that the denominators in our statistic should just be the scaled fluxes. In essence this is equivalent to realising that zero serves as a perfectly fine second fiducial χ^2 value (alongside the scaled flux).

With this last realisation, the $\Psi_{\Phi,c}^2$ statistic had finally taken the form it has in Eq. (4.20).

¹Note that the notation used here for this statistic (as well as for the others) is different from the notation used at the time since that notation would conflict with what is used in the main body of the thesis.

Bibliography

- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *ApJS*, 187, 460 [page 23].
- Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, *ApJs*, 208, 17 [pages 23, 135].
- Abdollahi, S., Acero, F., Ackermann, M., et al. 2020, *ApJS*, 247, 33 [page 109].
- Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, *Nature*, 300, 728 [page 28].
- Andrae, R., Schulze-Hartung, T., & Melchior, P. 2010, arXiv e-prints, arXiv:1012.3754 [page 43].
- Arons, J. 1983, *ApJ*, 266, 215 [page 34, 34, 34].
- Atwood, W. B., et al. 2009, *ApJ*, 697, 1071 [pages 15, 22, 23].
- Avni, Y. 1976, *ApJ*, 210, 642 [pages 45, 95].
- Baade, W. 1942, *ApJ*, 96, 188 [page 21].
- Baade, W., & Zwicky, F. 1934, *Proceedings of the National Academy of Science*, 20, 259 [page 21].
- Baym, G., Hatsuda, T., Kojo, T., et al. 2018, *Reports on Progress in Physics*, 81, 056902 [page 25].
- Bevington, P. R. 1969, *Data reduction and error analysis for the physical sciences* (New York: McGraw-Hill, 1969) [page 43].
- Bezuidenhout, T., Venter, C., Seyffert, A., & Harding, A. K. 2017, in *5th Annual Conference on High Energy Astrophysics in Southern Africa*, 18 [pages 18, 120].
- Bhattacharya, D., & van den Heuvel, E. P. J. 1991, *Physics Reports*, 203, 1 [page 28].
- Bogovalov, S. V. 1999, *A&A*, 349, 1017 [pages 36, 37].
- Bowyer, S., Byram, E. T., Chubb, T. A., & Friedman, H. 1964, *Nature*, 201, 1307 [page 21].
- Bradt, H. V., Rothschild, R. E., & Swank, J. H. 1993, *A&A*, 97, 355 [page 22].
- Brambilla, G., Kalapotharakos, C., Harding, A. K., & Kazanas, D. 2015, *ApJ*, 804, 84 [pages 119, 135, 136, 136, 136, 136, 141, 141].
- Carroll, B. W., & Ostlie, D. A. 1996, *An introduction to Modern Astrophysics* (Addison-Wesley Publishing Company, Inc.) [page 25].
- Chen, A. Y., & Beloborodov, A. M. 2014, *ApJL*, 795, L22 [page 37].
- Cheng, K. S., Ho, C., & Ruderman, M. 1986a, *ApJ*, 300, 500 [page 35, 35, 35].
- . 1986b, *ApJ*, 300, 500 [pages 32, 81].

- . 1986c, *ApJ*, 300, 522 [pages 35, 35, 36, 36, 36].
- Clark, C. J., Nieder, L., Voisin, G., et al. 2020, *MNRAS* [pages 18, 18, 109, 114, 114].
- Cocke, W. J., Disney, M. J., & Taylor, D. J. 1969, *Nature*, 221, 525 [page 22].
- Corongiu, A., Mignani, R. P., Seyffert, A. S., et al. 2020, *MNRAS* [pages 18, 109, 110].
- Coroniti, F. V. 1990, *ApJ*, 349, 538 [page 37].
- Cromartie, H. T., Fonseca, E., Ransom, S. M., et al. 2019, *Nature Astronomy*, 4, 72 [page 25].
- Daugherty, J. K., & Harding, A. K. 1982, *ApJ*, 252, 337 [pages 32, 33].
- Daugherty, J. K., & Harding, A. K. 1983, *ApJ*, 273, 761 [page 32].
- Deutsch, A. J. 1955, *Annales d’Astrophysique*, 18, 1 [page 39].
- Duyvendak, J. J. L. 1942, *Publications of the Astronomical Society of the Pacific*, 54, 91 [pages 21, 26].
- Dyks, J., Harding, A. K., & Rudak, B. 2004, *ApJ*, 606, 1125 [page 40, 40].
- Dyks, J., & Rudak, B. 2003, *ApJ*, 598, 1201 [page 39].
- Dyks, J., & Rudak, B. 2003, *ApJ*, 598, 1201 [page 42].
- Erber, T. 1966, *Rev. Mod. Phys.*, 38, 626 [pages 31, 31, 32].
- Fichtel, C. E., Hartman, R. C., Kniffen, D. A., et al. 1975, *ApJ*, 198, 163 [page 22].
- Forman, W., Jones, C., Cominsky, L., et al. 1978, *ApJS*, 38, 357 [page 22].
- Fritz, G., Henry, R. C., Meekins, J. F., Chubb, T. A., & Friedman, H. 1969, *Science*, 164, 709 [page 22].
- Fruchter, A. S., Berman, G., Bower, G., et al. 1990, *ApJ*, 351, 642 [page 28].
- Gehrels, N., Fichtel, C. E., Fishman, G. J., Kurfess, J. D., & Schonfelder, V. 1993, *Scientific American*, 269, 68 [page 15].
- Giacconi, R., Gursky, H., Paolini, F. R., & Rossi, B. B. 1962, *Physical Review Letters*, 9, 439 [page 21].
- Gold, T. 1968, *Nature*, 218, 731 [pages 21, 21, 22, 29].
- . 1969, *Nature*, 221, 25 [page 22].
- Goldreich, P., & Julian, W. H. 1969, *ApJ*, 157 [pages 29, 29, 29, 29, 29, 30].
- Gonthier, P. L., Van Guilder, R., & Harding, A. K. 2004, *ApJ*, 604, 775 [page 39].
- Harding, A. K. 2013, *Frontiers of Physics*, 8, 679 [page 28].
- Harding, A. K., & Muslimov, A. G. 2003, in *Pulsars, AXPs and SGRs Observed with BeppoSAX and Other Observatories*, ed. G. Cusumano, E. Massaro, & T. Mineo, 121–126 [page 32].
- Hewish, A., et al. 1968, *Nature*, 217, 709 [page 22, 22].
- Hillier, R. R., Jackson, W. R., Murray, A., Redfern, R. M., & Sale, R. G. 1970, *ApJ*, 162, L177 [page 22].
- Johnson, T. J., Venter, C., & Harding, A. K. 2011, in *American Astronomical Society Meeting Abstracts*, Vol. 217, American Astronomical Society Meeting Abstracts #217, 215.04 [page 56].

- Johnson, T. J., Venter, C., Harding, A. K., et al. 2014, *The Astrophysical Journal Supplement Series*, 213, 6 [pages 10, 15, 15, 43, 56, 56, 91, 117, 117, 118, 119, 119, 119, 120, 129, 129, 130, 131, 131, 132].
- Kalapotharakos, C., Harding, A. K., & Kazanas, D. 2014, *ApJ*, 793, 97 [pages 36, 135].
- Kanbach, G., Bennett, K., Bignami, G. F., et al. 1980, *A&A*, 90, 163 [pages 22, 29].
- Kanbach, G., Bertsch, D. L., Fichtel, C. E., et al. 1989, *Space Science Reviews*, 49, 69 [page 22].
- Kaspi, V. M. 2018, in *Pulsar Astrophysics the Next Fifty Years*, ed. P. Weltevrede, B. B. P. Perera, L. L. Preston, & S. Sanidas, Vol. 337, 3–8 [page 28].
- Kijak, J., & Gil, J. 1998, "MNRAS", 299, 855 [pages 39, 39, 81].
- . 2003, *A&A*, 397, 969 [pages 39, 81].
- Lampton, M., Margon, B., & Bowyer, S. 1976, *ApJ*, 208, 177 [page 44].
- Li, J., Spitkovsky, A., & Tchekhovskoy, A. 2012, *ApJ*, 746, 60 [page 36].
- Luo, Q., Shibata, S., & Melrose, D. B. 2000, *MNRAS*, 318, 943 [page 32].
- Malina, R. F., Bowyer, S., Lampton, M., et al. 1982, *Optical Engineering*, 21, 764 [page 22].
- Mayall, N. U., & Oort, J. H. 1942, *Publications of the Astronomical Society of the Pacific*, 54, 95 [pages 21, 27].
- Minkowski, R. 1942, *ApJ*, 96, 199 [page 21].
- Misner, C. W., Thorne, K. S., & Wheeler, J. A. 1973, *Gravitation* (San Francisco: W.H. Freeman and Co., 1973) [page 21].
- Morini, M. 1983, *MNRAS*, 202, 495 [page 40].
- Nolan, P. L., Abdo, A. A., Ackermann, M., et al. 2012, *ApJS*, 199, 31 [page 109].
- Oppenheimer, J. R., & Volkoff, G. M. 1939, *Physical Review*, 55, 374 [page 21].
- Ostriker, J. P., & Gunn, J. E. 1969, *ApJ*, 157, 1395 [page 26].
- Pacini, F. 1967, *Nature*, 216, 567 [page 21].
- . 1968, *Nature*, 219, 145 [page 21].
- Padmanabhan, T. 2001, *Theoretical Astrophysics - Volume 2, Stars and Stellar Systems* (Cambridge University Press) [pages 21, 21, 25].
- Pearson, K. 1900, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50, 157 [page 43].
- Pétri, J., & Mitra, D. 2020, *MNRAS*, 491, 80 [page 15].
- Philippov, A. A., & Spitkovsky, A. 2014, *ApJL*, 785, L33 [page 37].
- Philippov, A. A., Spitkovsky, A., & Cerutti, B. 2015, *ApJL*, 801, L19 [page 37].
- Pierbattista, M., Harding, A. K., Grenier, I. A., et al. 2015, *A&A*, 575, A3 [pages 10, 15, 15, 43, 56, 56, 56, 56, 117, 117, 118, 119, 119, 119, 120, 121, 129, 129, 130, 131, 131, 132].
- Roberts, M. S. E. 2013, in *Neutron Stars and Pulsars: Challenges and Opportunities after 80 years*, ed. J. van Leeuwen, Vol. 291, 127–132 [page 109].
- Romani, R. W. 2015, *ApJL*, 812, L24 [page 109, 109].

- Rudak, B., & Dyks, J. 1999, MNRAS, 303, 477 [page 31].
- Rybicki, G. B., & Lightman, A. P. 1979, Radiative processes in astrophysics (Wiley-Interscience) [pages 30, 31].
- Salvetti, D., Mignani, R. P., De Luca, A., et al. 2015, ApJ, 814, 88 [page 109, 109, 109].
- Seyffert, A., Venter, C., Harding, A. K., Allison, J., & Schutte, W. 2015, in 3rd Annual Conference on High Energy Astrophysics in Southern Africa (HEASA2015), 28 [page 18].
- Seyffert, A. S., Venter, C., Harding, A. K., Allison, J., & Schutte, W. D. 2016, arXiv e-prints, arXiv:1611.01076 [page 18].
- Shapiro, S. L., & Teukolsky, S. A. 1983, Black holes, white dwarfs and neutron stars: The physics of compact objects (Wiley Interscience) [page 28].
- Staelin, D. H., & Reifenstein, I. E. C. 1968, Science, 162, 1481 [page 22].
- Story, S. A., Gonthier, P. L., & Harding, A. K. 2007, ApJ, 671, 713 [pages 39, 39, 39, 48, 81].
- Sturrock, P. A. 1971, ApJ, 164, 529 [page 31].
- Tavani, M., Barbiellini, G., Argan, A., et al. 2009, A&A, 502, 995 [page 15].
- Thompson, D. J. 2001, in AIP Conference Series, ed. F. A. Aharonian & H. J. Völk, Vol. 558, 103–114 [page 23].
- Tolman, R. C. 1939, Physical Review, 55, 364 [page 21].
- Truemper, J. 1982, Advances in Space Research, 2, 241 [page 22].
- Truemper, J., Pietsch, W., Reppin, C., et al. 1978, ApJ, 219, L105 [page 22].
- van Rensburg, C., Krüger, P. P., & Venter, C. 2018, MNRAS, 477, 3853 [page 137, 137].
- van Rensburg, C., Venter, C., Seyffert, A. S., & Harding, A. K. 2020, MNRAS, 492, 3091 [pages 18, 137, 137, 142, 143, 144].
- Venter, C., Harding, A. K., & Guillemot, L. 2009, ApJ, 707, 800 [pages 15, 15, 39, 40].
- Viganò, D., Torres, D. F., Hirotani, K., & Pessah, M. E. 2015, MNRAS, 447, 1164 [page 31].
- Weltevrede, P., et al. 2010, ApJ, 708, 1426 [pages 15, 24, 24, 25].
- Weltevrede, P., Abdo, A. A., Ackermann, M., et al. 2010, ApJ, 708, 1426 [page 117].
- Wheeler, J. A. 1966, Annual Review of Astronomy and Astrophysics, 4, 393 [page 21].