Bayesian methods for the analysis of ultra-high-energy cosmic rays

Alexander Khanin
Astrophysics Group

Department of Physics
Imperial College London

Thesis submitted for the Degree of Doctor of Philosophy to
Imperial College London

· 2016 ·

"Consider, for example, a globular world, a mere crust upon an inferno of molten rock and iron. An accidental world, made of the wreckage of old stars, the home of life which, nevertheless, in a most unhomely fashion, is regularly scythed from its surface by ice, gas, inundation or falling rocks travelling at 20,000 miles an hour."
-Terry Pratchett, *The Science of Discworld*

I dedicate this thesis to my Mom, my Dad, and my sister Asya.
Thanks for always supporting me, even when we're not on the same landmass.

Declaration and Copyright

My contribution to this work has been: i) the development of the Bayesian formalism, done in collaboration with my supervisor, Daniel Mortlock; ii) the development of the program that was used to apply the formalism to the data; iii) the data analysis; and iv) the interpretation of the results, in consultation with my supervisor. I also wrote the two papers that are the basis of Chapters 3 and 4 of this thesis.

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Alexander Khanin
May 2016

Publications

Khanin A., Mortlock D. J., 2016, *A Bayesian analysis of the 69 highest energy cosmic rays detected by the Pierre Auger Observatory* , Monthly Notices of the Royal Astronomical Society, in press
Khanin A., Mortlock D. J., 2014, *A Bayesian self-clustering analysis of the highest energy cosmic rays detected by the Pierre Auger Observatory*, Monthly Notices of the Royal Astronomical Society, Vol: 444, Pages: 1591-1599, ISSN: 0035-8711

Abstract

The origins of ultra-high-energy cosmic rays (UHECRs) are one of the open puzzles of astrophysics.

A number of plausible candidates, such as active galactic nuclei (AGNs) have been discussed, but no clear consensus has been reached. One way to assess the different hypotheses is by analyzing the UHECR arrival directions. Recently, a small number of studies have begun applying Bayesian methodologies to this problem, forming the first steps in the development of a comprehensive Bayesian framework for the study of UHECRs. In this work, we have developed two Bayesian methods to study this question, and have applied them to UHECRs from the Pierre Auger Observatory (PAO).

The first method was a Bayesian approach to studying the catalogue-independent clustering of UHECRs. Previously, this had been difficult as there is no well motivated clustered model that can be used in a Bayesian model comparison. We have resolved this difficulty by developing a multi-step approach that derives such a model from a sub-set of the data. This approach could have broad applications for anisotropy searches in other areas of astronomy. Our results were consistent with both isotropic and clustered models.

The second was a Bayesian method that was aimed to find associations between UHECR arrival directions and source catalogues. It was an extension of a previous Bayesian study, but analyzed a greater data set, used a more refined UHECR model, and was generalized to be applicable to a greater variety of source catalogues. Our results were broadly consistent with previous work, with the purely isotropic UHECR models being disfavoured for reasonable parameter ranges.

It will be of great interest to apply our methods to samples of greater size. The extended UHECR samples that will be available in the near future should be sufficient for our methods to determine the origins of the UHECRs.

Contents

List of Tables **9**

List of Figures **10**

Acknowledgements **11**

Thesis summary **12**

1 Introduction to ultra-high-energy cosmic rays **14**
1.1 Background . . . 15
1.2 Acceleration . . . 17
1.2.1 Fermi Acceleration . . . 17
1.2.2 Unipolar induction . . . 18
1.3 Propagation . . . 20
1.3.1 Interactions with cosmic backgrounds . . . 20
1.3.2 Interactions with magnetic fields . . . 22
1.4 Detection . . . 23
1.4.1 Cosmic ray observatories . . . 24
1.5 Principal observables and experimental results . . . 26
1.5.1 Composition . . . 26
1.5.2 Arrival energy . . . 27
1.5.3 The arrival directions of cosmic rays . . . 28

2 Bayesian methods **32**
2.1 Cox's theorem . . . 33
2.2 Bayes' theorem . . . 35
2.3 Priors . . . 36
2.4 Bayesian parameter inference . . . 37

2.5 Model comparison . . . 39
2.5.1 Model comparison with poorly specified prior parameters . . . 40
2.5.2 Calculation of the evidence . . . 41
2.5.3 Example case: On/Off measurements . . . 41
2.5.4 Example case: Gaussian likelihood . . . 44
2.6 Comparison with frequentist statistics . . . 47
2.6.1 Example 1: Confidence intervals . . . 49
2.6.2 Example 2: Significance tests . . . 53

3 A Bayesian self-clustering analysis 56
3.1 Introduction . . . 57
3.2 Statistical formalism . . . 57
3.2.1 Uniform model . . . 58
3.2.2 Non-uniform model . . . 60
3.2.3 Non-uniform exposure . . . 62
3.2.4 Illustration of the multi-step method . . . 63
3.3 Application to simulated UHECR samples . . . 64
3.3.1 Isotropic distribution of sources . . . 65
3.3.2 AGN sources . . . 65
3.3.3 Measurement . . . 68
3.3.4 Results of the simulations . . . 69
3.4 Analysis of the PAO data . . . 71
3.5 Conclusions . . . 71

4 A Bayesian analysis of UHECR source fractions 75
4.1 Introduction . . . 76
4.2 Data . . . 77
4.2.1 UHECR sample . . . 77
4.2.2 Source catalogues . . . 77
4.3 UHECR model . . . 79
4.3.1 Injection . . . 79
4.3.2 Energy loss during propagation . . . 79
4.3.3 Effective smearing . . . 80
4.3.4 Observed UHECR flux . . . 80
4.4 Statistical formalism . . . 81
4.4.1 Prior . . . 82
4.4.2 The likelihood . . . 83
4.4.3 The source fraction . . . 85

4.4.4 Model comparison 85
4.5 Simulations 87
4.6 Results 88
4.7 Conclusions 93

5 Thesis Conclusions **95**

Bibliography **97**

List of Tables

1.1 Comparison between surface detectors and fluorescent telescope detectors. 24
1.2 Current UHECR observatories. 25

2.1 Two example hypotheses. 54

4.1 Maximum a posteriori estimates and 68% credible intervals for F_{src}. . 91

List of Figures

1.1 Differential energy spectrum of the cosmic rays. 16
1.2 Sizes and magnetic field strengths of astrophysical accelerators. 19
1.3 Loss lengths from the UHECR energy loss processes. 22
1.4 Relative PAO exposure in Galactic coordinates. 26
1.5 UHECR energy spectra, as measured by different observatories. . . . 29

2.1 Robustness of Bayesian priors demonstrated with a coin example. . . 38
2.2 On/Off likelihood, prior regions, and posteriors. 45
2.3 Gaussian likelihood, prior regions, and posteriors. 46
2.4 Dependence of the Bayes factor on hyperparameter s for the cases of the On/Off likelihood and the Gaussian likelihood. 47
2.5 The truncated exponential. 51

3.1 Illustration of model creation. 59
3.2 Application of multi-step method to test cases. 62
3.3 Flowchart visualising the propagation model. 69
3.4 Results of the multi-step method applied to mock UHECR catalogues. 73
3.5 Application of the multi-step method to the PAO data. 74

4.1 Relative PAO exposure in Galactic coordinates. 77
4.2 Expected UHECR number flux for four source catalogues. 78
4.3 Loss lengths from the four energy loss processes. 80
4.4 Example likelihood and priors for the Poisson product likelihood. . . 82
4.5 Poisson product likelihood, prior regions, and posteriors. 86
4.6 Dependence of the Bayes factor on hyperparameter s for the cases of the On/Off, Gaussian, and Poisson product likelihoods. 87
4.7 Application of the method to mock catalogues. 89
4.8 Application of the method to 27 and 69 events from PAO. 92
4.9 Application of the method to PAO data. 94

Acknowledgements

First and foremost, I would like to thank Daniel Mortlock for having been a fantastic mentor and supervisor. Thank you for your invaluable guidance and support.

There are many people I would like to thank for their help with my research. Andrew Jaffe and Steven Warren provided valuable feedback on many stages of my work, and helped me greatly. I would like to thank Kumiko Kotera and Manlio De Domenico for our discussions on cosmic rays and aspects of my work. I am grateful to Andreas Berlind and Glennys Farrar for making their mock catalogues public, and would like to thank Todor Stanev for providing the results of his UHECR propagation models.

This work was greatly helped by the thoughtful comments of the two anonymous referees of my papers.

Finally, I would like to extend my gratitude to my family for their invaluable support, and to the Imperial Astrophysics postgrads for all the nice memories.

Thesis summary

This thesis is concerned with Bayesian methods for the analysis of the arrival directions of UHECRs.

In Chapter 1, we provide general background to UHECRs. We explain some of the difficulties that arise in the study of these particles, and explain how the study of the arrival directions of UHECRs provides insight into their sources. We also give background to the PAO and the 69 events that we analyze.

In Chapter 2 we discuss Bayesian inference, focusing on methods that we use in our analysis. This includes a discussion of cases of Bayesian model comparison with poorly motivated prior parameters, and how multi-step Bayesian approaches can be used to overcome this problem. We briefly review the advantages of Bayesian methods over frequentist approaches.

In Chapter 3, we discuss a Bayesian approach for the analysis of the anisotropy of UHECRs independent of any source-catalogue. Such catalogue-independent analysis has a number of advantages over catalogue-dependent methods, and until this point has not been conducted in a Bayesian way. We have developed a multi-step Bayesian approach to study this question, and have applied it to a sample of 69 events recorded by the Pierre Auger Observatory (PAO). The result of our analysis was that the PAO data are consistent with both uniform and anisotropic models. We have shown that data sets of 690 UHECRs would be sufficient to distinguish between these two models. Samples of this magnitude are expected to be produced by future experiments, such as the Japanese Experiment Module Extreme Universe Space Observatory (JEM-EUSO).

In Chapter 4, we discuss a Bayesian method that is aimed to determine the source fraction of a sample of UHECRs, which is the fraction of the sample that is expected to have originated at whichever source catalogue is under consideration. Our methodology can be regarded as an extension and a generalization of a previous study. We use our method to determine constraints on the source fraction of the 69 PAO events for several source catalogues: AGNs in the Veron-Cety & Verson (VCV) catalogue, as well as AGNs detected with the Swift Burst Alert Telescope (Swift-

BAT), galaxies from the 2MASS Redshift Survey (2MRS), and an additional volume-limited sample of 17 nearby AGNs. For fiducial values of the model parameters, we report 68% credible intervals for the fraction of source originating UHECRs of $0.09^{+0.05}_{-0.04}$, $0.25^{+0.09}_{-0.08}$, $0.24^{+0.12}_{-0.10}$, and $0.08^{+0.04}_{-0.03}$ for these respective catalogues. We find that for reasonable ranges of prior parameters, the Bayes factors disfavour a purely isotropic model.

Chapter 5 rounds off our discussion and suggests possible avenues for future work.

Chapter 1

Introduction to ultra-high-energy cosmic rays

1.1 Background

Cosmic rays (CRs) are high-energy particles that flow into the Solar System and reach the Earth. The study of UHECRs harkens back to the beginning of the 20th century. After radioactivity was discovered by Henri Becquerel and Marie Curie in 1896, it was discovered that between 10 and 20 ions are generated per cubic centimeter of air every second, which was believed to be due to the natural radioactivity of the Earth. This idea was disproved in 1911 by Domenico Pacini (Pacini 1912), and one year later independently by the better known Victor Hess (Hess 1912). Pacini developed an experimental technique for conducting radioactivity measurements underwater, and found that the rate of ionization underwater is lower than at sea level. Victor Hess carried out balloon experiments that demonstrated an increase in radiation at high altitudes when compared to the radiation at sea level. Hess conducted a similar balloon experiment during a near-total eclipse, during which he still could measure rising ionization rates at high altitudes, thus ruling out the Sun as the source of the radiation. These results independently led Pacini and Hess to the conclusion that radiation was entering the Earth's atmosphere from outer space. This radiation was dubbed "cosmic rays". In 1936, Hess received the Nobel Prize in Physics for this discovery. Jacob Clay found that the intensity of this radiation is lower near the equator than at northern latitudes. This led him to the conclusion that the cosmic rays are not photons, but charged particles that are affected by the geomagnetic field (Clay & Berlage 1932).

Today it is believed that CRs consist mainly of protons and atomic nuclei, and have energies in the range 10^9 eV to 10^{21} eV. The energies of the highest energy CRs are the highest of any particle observed in nature, and are substantially higher than energies that can be produced in accelerators on Earth. The "Oh-My-God" particle that was measured in 15 October 1991 had an estimated energy of 3×10^{20} eV, which is approximately 40 million times the energy of the highest energy protons in accelerators (Bird et al. 1993). The highest energy CRs make it possible to investigate particle physics at energies that otherwise are not accessible. Before the creation of particle accelerators, the study of CRs was the best way to do high energy particle physics. Many particles, such as the positron (Anderson 1933) and the muon (Neddermeyer & Anderson 1937), were discovered in observations of CRs.

Figure 1.1 shows the energy spectrum of CRs with observed energies $E_{\rm obs} > 10^{11}$ eV. The spectrum has three general features: the knee at $\sim 10^{15}$ eV; the ankle at $\sim 3 \times 10^{18}$ eV; and the cutoff above $\sim 10^{20}$ eV. The UHECR flux dN/dE follows a power law spectrum $\mathrm{dN/dE} \propto \mathrm{E}^{-2.7}$ below the knee. Between the knee and the

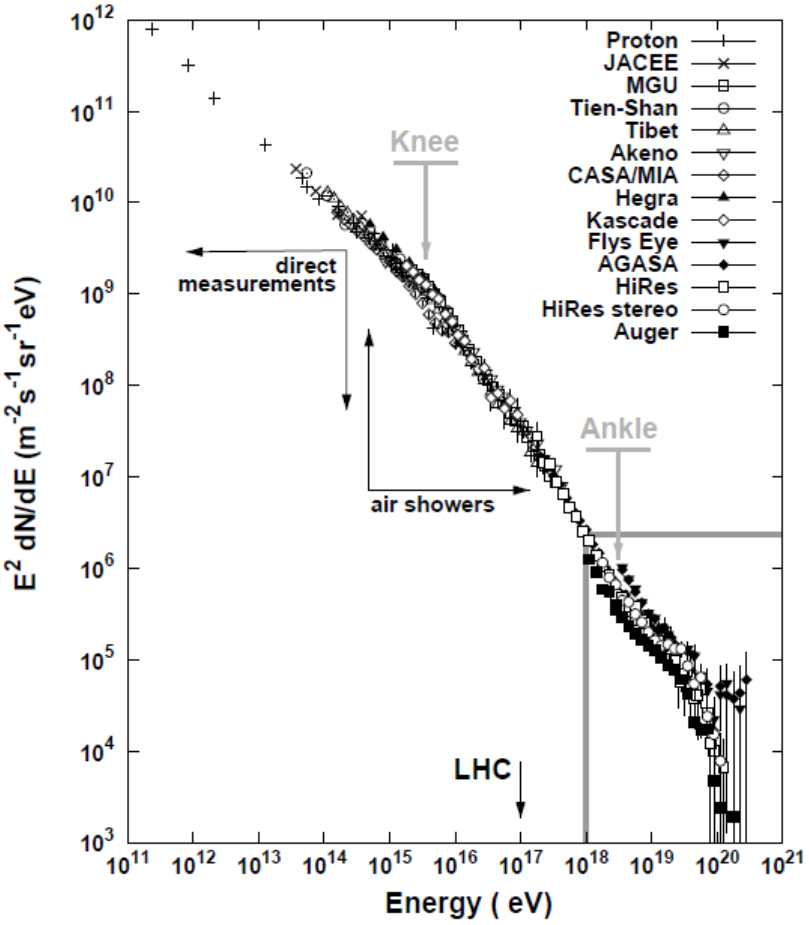


Figure 1.1: Differential energy spectrum of CRs of energy above 10^{11} eV. The energy spectrum has been multiplied by E^2 to emphasize the spectral shape. The plot distinguishes between data from different CR experiments. The locations of the knee and ankle are indicated with arrows. The graph shows that lower energy CRs can be directly observed, while higher energy CRs are detected by analyzing air showers. The LHC energies are shown for comparison. Figure taken from Letessier-Selvon & Stanev (2011).

ankle, the spectrum steepens, the power index increasing by ~ 0.3. Above the ankle, the spectrum flattens, and becomes similar to the spectrum before the knee. Above the cutoff, the spectrum falls off exponentially (Letessier-Selvon & Stanev 2011).

A number of open issues remain in the study of cosmic rays, especially with respect to ultra-high-energy cosmic rays (UHECRs) with arrival energies $E \gtrsim 10^{18}$ eV. In particular, the sources of these particles remain an open question. A number of plausible candidates have been proposed, such as active galactic nuclei (AGNs), gamma-ray bursts (GRBs) and pulsars, but no definite conclusion has been reached.

Several issues complicate the study of this question. One is the fact that UHECRs are charged particles that travel through a magnetized universe, so that

they experience magnetic deflection during their propagation to Earth. The arrival directions of lower energy CRs are thought to be largely independent of their point of origin, while the arrival directions of UHECRs are expected to be deflected by a few degrees (De Domenico & Insolia 2013). The second, and most important complication is the fact that the flux of UHECRs is extremely low, leading to sparse data sets. For example, the flux of UHECRs above $E \gtrsim 10^{20}\,\mathrm{eV}$ is ~ 1 event per square kilometer per century per steradian (Stanev 2004). The observed number flux of UHECRs falls off with energy, and falls off extremely rapidly above $\sim 5 \times 10^{19}\,\mathrm{eV}$, as UHECRs with energies above the threshold interact with the cosmic microwave background (CMB) to produce pions (see Section 1.3).

In this chapter, we provide a summary of the physics of UHECRs, their injection, propagation, and detection on Earth. We explain how the question of the UHECR sources can be investigated by analyzing the arrival directions of the rays, and summarize the work which has previously been done on this subject.

1.2 Acceleration

The accelerating mechanism of UHECRs is not settled. Accelerating microscopic particles to the required ultra-high-energies ($\sim 20\,\mathrm{J}$) is extremely challenging. Here, we briefly review two acceleration mechanisms that are commonly cited: Fermi acceleration and unipolar induction.

1.2.1 Fermi Acceleration

Fermi acceleration is based on the acceleration of charged particles through their interaction with magnetic inhomogeneities. Fermi originally discussed a mechanism in which a cosmic particle is repeatedly scattered off magnetized clouds (Fermi 1949). The fractional energy gain of this process is proportional to β^2, where β is the average velocity of the scattering centres in terms of c. This process is known today as second order Fermi acceleration. The average energy gain through this mechanism is fairly small. The first truly successful stochastic UHECR acceleration process was based on acceleration of charged particles inside shock waves, and is called first order Fermi acceleration (Axford et al. 1977). In this process, a charged particle in front of the shock wave passes through the shock, where it is scattered by magnetic irregularities, so that it is projected back through the shock, where it is scattered by magnetic irregularities again. In this way, the particle can bounce back and forth through the shock, gaining energy. The energy gain is thereby proportional

to β, where β here is the blast shock velocity in terms of c. This is known as first order Fermi acceleration.

In first order Fermi acceleration, it seems that the fractional energy gain in the first crossing is Γ_s^2, where Γ_s is the shock bulk Lorentz factor, while the fractional gain in subsequent crossings is $\beta \sim 1$ (Achterberg et al. 2001). Fermi acceleration leads to a power law UHECR spectrum with an index $\alpha = 2$ for non relativistic shocks, and $2.1 - 2.3$ for relativistic shocks (Achterberg et al. 2001), which is consistent with observations.

Sources that make use of the Fermi mechanism can be narrowed down significantly by making use of the Hillas criterion (Hillas 1984). The main idea behind this criterion is that when the Larmor radius of a charged particle inside a magnetic field r_{L} approaches the size of the electromagnetic accelerator, the particle can no longer be confined inside the accelerator. This puts a bound on the maximum energy that a charged particle can be accelerated to inside an accelerator. More formally, the criterion can be written as

$$E_{\max} \approx 2\beta Q B r_{\mathrm{L}} \tag{1.1}$$

where Q is the charge of the particle, B is the magnetic field, and r_{L} is the characteristic scale of the accelerator. A number of sources have been suggested that fit the Hillas criterion well and that reproduce the aforementioned UHECR properties: active galactic nuclei, accretion shocks in galaxy clusters, and gamma ray bursts. Lower energy, Galactic CRs are believed to be accelerated by non-relativistic shocks in supernova remnants (see, e.g., Hillas 2006).

Figure 1.2 shows a number of astrophysical sources of high energy particles. The figure shows diagonal lines for different isotopes of UHECRs: protons and iron nuclei (see Section 1.5.1 for a discussion of UHECR composition). The sources above each line are capable of accelerating CRs of the respective composition above $10^{20}\,\mathrm{eV}$. The sizes of the sources thereby range from kilometers to megaparsecs.

1.2.2 Unipolar induction

In unipolar induction, the cosmic ray is accelerated by a powerful electric field that is induced in rapidly rotating compact magnetized objects such as pulsars and black holes. Such objects generate relativistic outflows, where the electric field is induced due to a combination of the strong magnetic fields and the rotational energy: $\boldsymbol{E} = \boldsymbol{v} \times \boldsymbol{B}/c$, where $\boldsymbol{v}$ is the velocity of the outflowing plasma and $\boldsymbol{B}$ is the magnetic field.

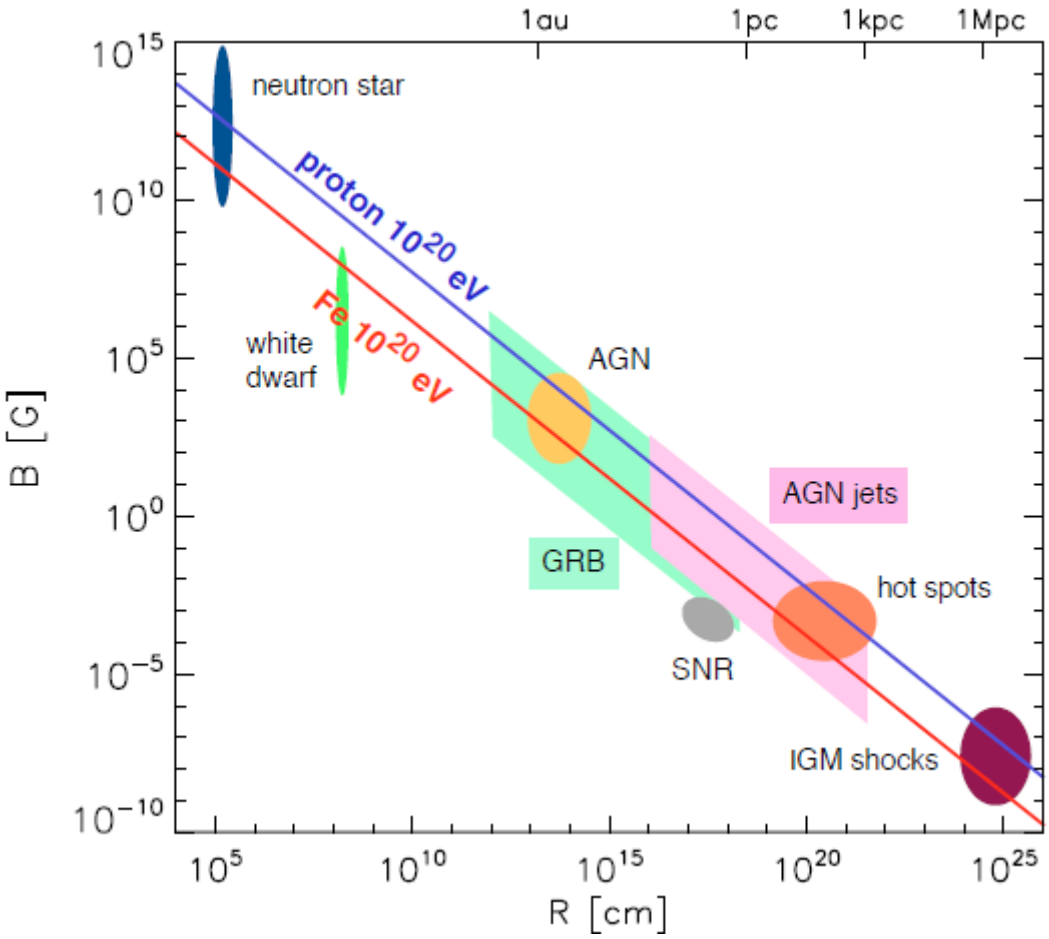


Figure 1.2: Sizes and magnetic field strengths of astrophysical accelerators. Two diagonal lines are shown for protonic and iron CRs. For each of the lines, the sources above the line are capable of accelerating CRs of the respective composition to 10^{20} eV. Figure taken from Kotera & Olinto (2011).

Unipolar induction was first discussed in the context of ordinary pulsars (Shapiro & Teukolsky 1983). However, ordinary pulsars can only accelerate particles to energies of 10^{15} eV. In Blasi et al. (2000), this model was discussed in the framework of magnetars, young neutron stars with millisecond rotation periods and extremely powerful magnetic fields. It was shown that such objects can accelerate iron nuclei to energies $\sim 10^{20}$ eV, making them a viable acceleration mechanism for UHECRs. This acceleration mechanism leads to a power law injection spectrum for the UHECRs, with an injection index $\alpha = 1$. Such a hard spectrum is challenging to reconcile with current observations, but that spectrum can be softened by an adequate distribution of initial voltages among magnetar winds (Kotera 2011).

1.3 Propagation

During their propagation, UHECRs experience two kinds of interactions:

(i) Interactions with cosmic backgrounds, which cause energy loss and alter the composition of the UHECRs.

(ii) Interactions with the Galactic and extra-Galactic magnetic fields, which alter the UHECR trajectories.

Here, we review both of these. We provide more detail for UHECRs of protonic composition, as those are the focus of this work.

1.3.1 Interactions with cosmic backgrounds

UHECRs interact with the CMB at the highest energies, and with the infrared-ultraviolet background at lower energies. For protonic UHECRs, these interactions generally take the form either of electron-positron pair production (also known as the Bethe-Heitler process) or of pion production. These two processes can be written as:

$$p\gamma \rightarrow pe^{+}e^{-} \tag{1.2}$$

$$p\gamma \rightarrow N + n\pi, \tag{1.3}$$

where p and γ are the protonic UHECR and the background photon, respectively, N is a nucleon and n is the number of pions produced. The energy loss processes that UHECRs experience can be characterized in terms of the loss length $L_{\rm loss} = -E(\mathrm{d}E/\mathrm{d}r)^{-1}$, which equals $\lambda_{\rm mfp}/K_{\rm loss}$, where $\lambda_{\rm mfp}$ is the mean free path of the respective scattering process, and $K_{\rm loss}$ is the mean energy loss in a collision with a photon. $L_{\rm loss}^{-1}$ is a characteristic scale of the energy loss.

For pure proton composition, $L_{\rm loss}$ obeys the expression

$$L_{\rm loss}^{-1} = \frac{1}{c}[\beta_{\rm GZK}(E,z) + \beta_{\rm BH}(E,z) + \beta_{\rm adi}(E,z) + \beta_{\rm IR}(E,z)], \tag{1.4}$$

where c is the speed of light and $\beta_{\rm GZK}(E,z)$, $\beta_{\rm BH}(E,z)$, $\beta_{\rm adi}(E,z)$ and $\beta_{\rm IR}(E,z)$ are terms corresponding to the main energy loss processes experienced by UHECRs of pure proton composition, as a function of energy and redshift (e.g. Stanev 2009; Kotera & Olinto 2011):

(i) the GZK scattering. This is a photo-pion production process off the CMB photons. It dominates the total energy loss at energies above $E > 5 \times 10^{19}$ eV;

(ii) Bethe-Heitler (BH) pair production off the CMB radiation, which dominates at lower energies;

(iii) the adiabatic energy loss due to the expansion of the Universe. This is not a scattering process and is different from the others, but it is part of the overall energy loss of the UHECRs;

(iv) a photo-pion production process due to scattering off photons in the infrared-ultraviolet (IR-UV) spectrum, which plays a role at lower energies.

A detailed discussion of the first three processes, including relevant mathematical expressions, can be found in De Domenico & Insolia (2013). For the pion production off the IR-UV photons, the relevant discussion can be found in Kotera & Olinto (2011). The loss lengths are shown as a function of energy in the left panel of Figure 1.3. The figure shows the energy loss lengths of the separate processes for protonic UHECRs, as well as the combined total loss length $L_{\rm tot}$. From the figure, it is evident that the GZK scattering becomes extremely powerful for energies $> 5 \times 10^{19}$ eV, where it rapidly comes to dominate the overall energy loss process. This is called the GZK cutoff, and is of crucial significance to the study of UHECRs. It is expected that due to this cutoff, UHECRs at energies $> 5 \times 10^{19}$ eV should not be capable of travelling farther than a particular "GZK horizon", which is roughly ~ 100 Mpc. This puts a strong constraint on possible sources of UHECRs above the GZK cutoff. Figure 1.3 shows $L_{\rm tot}$ plots for z values of 0.0 and 0.1, which correspond to distances of 0 and ~ 400 Mpc, thus covering the GZK horizon.

The right panel of Figure 1.3 shows $L_{\rm tot}$ as a function of energy for heavier UHECRs of various masses. The figure shows that the minimum values of $L_{\rm tot}$ are significantly lower than for the protonic case, which means that the maximum energy loss is significantly greater. For heavier nuclei, the interaction with the cosmic backgrounds includes an additional class of processes: In addition to the aforementioned energy losses, heavy nuclei undergo photo-disintegration (spallation) processes on CMB and IR-UV photons, losing nuclei as they propagate (see, e.g., Puget et al. 1976; Epele & Roulet 1998; Stecker & Salamon 1999; Bertone et al. 2002; Allard et al. 2005; Hooper et al. 2008; Aloisio 2008). This substantially changes the picture: For protonic UHECRs, it is possible to calculate the entire energy loss that a UHECR experiences during propagation from $L_{\rm loss}$ alone. For heavier nuclei, this is not possible, and a full model of the energy loss needs to include values of the loss length for all UHECR isotopes lighter than the nucleus that was injected at the source, as well as cross sections of losing different numbers of nuclei due to spallation.

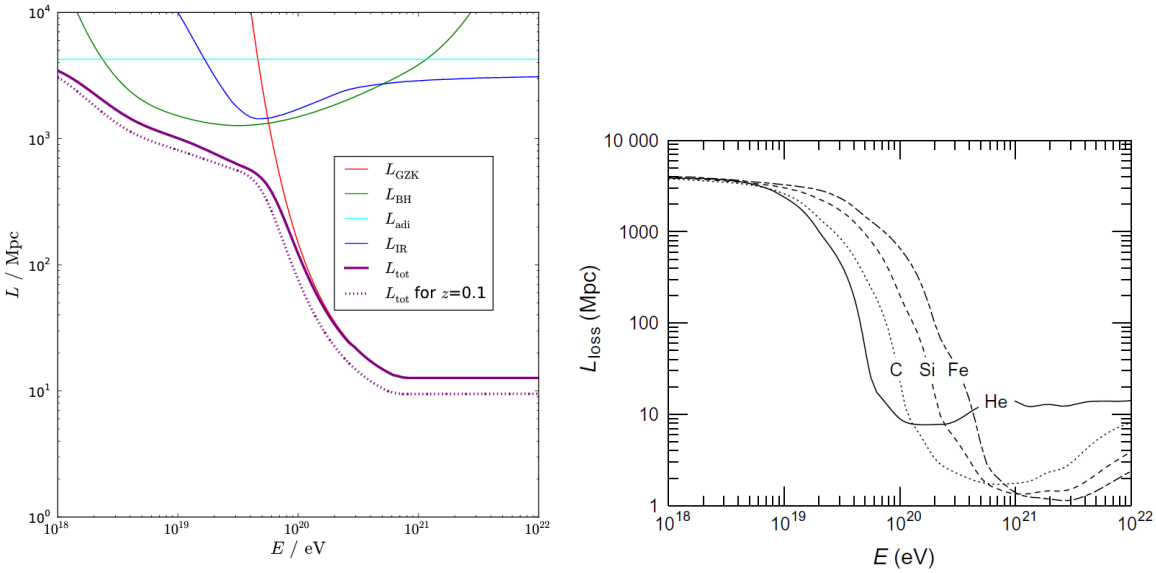


Figure 1.3: Left panel: Loss lengths from the four energy loss processes for UHECRs of protonic composition: GZK photo-production, BH pair-production, adiabatic energy losses, and photo-production off the infrared background. The four loss lengths are combined into a single total loss length $L_{\rm tot}$. The plots illustrate that after the GZK limit of $\sim 5 \times 10^{19}\,$eV, the GZK scattering becomes extremely powerful and dominates the overall energy loss. Right panel: $L_{\rm tot}$ plots for heavier UHECR nuclei, as calculated in Bertone et al. (2002). The right panel was taken from Stanev (2009).

1.3.2 Interactions with magnetic fields

As UHECRs are charged particles, their trajectories are altered by extra-Galactic and Galactic magnetic fields.

The fields of intergalactic space are weakly constrained, the most recent estimate of the upper bound being $\sim 2\,$nG (De Domenico & Insolia 2013). Estimates of the extent of magnetic deflection vary widely. For example, for a proton with $E > 10^{20}\,$eV, estimates of the magnetic deflection vary from less than a degree (Dolag et al. 2004) to $10 - 20^{\circ}$ (Sigl et al. 2004). The topology of the fields is also poorly known. It is possible that intergalactic fields have structure inside and around galaxy clusters and groups of galaxies (see, e.g., Ryu et al. 2008). Different topologies of the intergalactic B field could have a considerable impact on UHECR trajectories.

Significant progress has been made in recent years on observations of Galactic magnetic fields (see, e.g., Han et al. 1999; Han 2008; Jansson et al. 2009). The Galactic magnetic field is a combination of a regular and a random component. The regular field $B_{\rm reg}$ has a spiral structure which follows the structure of the Galaxy itself. The local strength of this field is $\sim 1.8\,\mu$G. Estimates of the random field are

in the range $B_{\rm rand} \sim 0.5 - 2\,B_{\rm reg}$. The correlation length of $B_{\rm rand}$ is $\sim 50 - 100\,$pc. It is likely that $B_{\rm rand}$ dominates the total magnetic field strength in the arms of the Galaxy, while $B_{\rm reg}$ dominates in the inter-arm space. Much work has been done on the effect of the Galactic magnetic fields on the propagation of UHECRs (Harari et al. 1999; Alvarez-Muñiz et al. 2002; Tinyakov & Tkachev 2002, 2005). These studies have shown that a particle of charge Z and energy E should not be deflected by more than $\sim 10^\circ Z(40{\rm EeV/E})$, so that a proton with $E = 10^{20}\,$eV would not be deflected by more than $\sim 4^\circ$.

In the small-deflection limit, as a UHECR traverses a distance L in a regular magnetic field B, it is deflected by an angle

$$\delta \simeq 6.4^\circ Z\Big(\frac{E}{50{\rm EeV}}\Big)^{-1}\Big|\int_L \frac{{\rm ds}}{3{\rm kpc}} \times \frac{B}{2\mu{\rm G}}\Big| \tag{1.5}$$

$$\delta \simeq 6.4^\circ Z\Big(\frac{E}{50{\rm EeV}}\Big)^{-1}\Big|\int_L \frac{{\rm ds}}{3{\rm Mpc}} \times \frac{B}{2{\rm nG}}\Big|, \tag{1.6}$$

where E and Z are the energy and atomic number of the UHECR, respectively (Harari et al. 2002). The first and second line reflect typical galactic and intergalactic scales, respectively. For a turbulent field, the deflection will be stochastic, with zero mean and a root-mean-square angular scale of

$$\delta_{\rm rms} \simeq 1.2^\circ \left(\frac{E}{50{\rm EeV}}\right)^{-1} \left(\frac{B_{\rm rms}}{4\mu{\rm G}}\right) \left(\frac{L}{3{\rm kpc}}\right)^{1/2} \left(\frac{l}{50{\rm pc}}\right)^{1/2} \tag{1.7}$$

$$\delta_{\rm rms} \simeq 2.3^\circ \left(\frac{E}{50{\rm EeV}}\right)^{-1} \left(\frac{B_{\rm rms}}{1{\rm nG}}\right) \left(\frac{L}{10{\rm Mpc}}\right)^{1/2} \left(\frac{l}{1{\rm Mpc}}\right)^{1/2}, \tag{1.8}$$

where the first and second lines once again represent typical galactic and intergalactic scales (Harari et al. 2002).

1.4 Detection

The number flux of UHECRs is extremely low, making it impossible to directly register the primary particles that interact with the upper atmosphere. UHECRs are measured indirectly through ground based installations that detect extended atmospheric showers (EAS) caused by the UHECRs. Still, the ground based detectors collect numbers of events that are negligible compared to, for example, the number of astrophysical photons measured in any band. The EAS mainly consist of electromagnetic, muonic and hadronic components, as well as electrons and positrons (see,

e.g., Keilhauer 2006).

The interaction energies involved lie far beyond energies that occur in the laboratory. The center-of-mass energies are of the order of several hundred TeV. Models that attempt to derive the particle properties from the shower development necessarily include extrapolation of our understanding of interaction physics to higher energies.

Two methods of UHECR detection are being used:

– surface detectors (SD). These detect the EAS particles on the ground. An array of detectors ($\sim 1\,\mathrm{km}$ spacing) determines the lateral distribution function of the shower's particle density. This method detects mostly the periferic part of the EAS.

– fluorescent telescope detectors (FD). These detect the ultraviolet emissions of the EAS by using a telescope. The UV radiation is caused by the excitation of atmospheric nitrogen by the cosmic ray particles. FD detects mostly the central part of the shower.

The advantages and disadvantages of the two types of detection are listed in Table 1.1.

Table 1.1: Comparison between surface detectors and fluorescent telescope detectors.

Type of detection	Advantages	Disadvantages
SD	Works independently of weather and time of day. Can detect different particle components of the shower.	Gives a lateral, 2-D picture of the shower.
FD	3-D picture, longitudinal as well as lateral components.	Only works on clear, moonless nights, which constitute about 10% of the total time. Only sensitive to the electron component of the EAS.

1.4.1 Cosmic ray observatories

The observatories that are currently operating are the Pierre Auger Observatory (PAO) (Abraham et al. 2004), Telescope Array (TA) (Nonaka et al. 2009), and Yakutsk (Ivanov 2009), where PAO and TA are hybrid detectors of SD and FD, while

Yakutsk is SD. The differences, strengths and weaknesses of these observatories are summarized in Table 1.2.

Table 1.2: Current UHECR observatories.

Observatory	Detection Method	Features
Yakutsk complex	SD detection. SD scintillators covering $\sim 10\,\mathrm{km}^2$.	Simultaneous detection of several EAS components. Very useful for the analysis of UHECR composition due to its large area muon detectors.
Telescope Array experiment (TA)	Hybrid observatory. Combination of SD plastic scintillators with an area of 680 km^2 and three FD telescopes.	In the hybrid regime, TA can analyse the same EAS through independent SD and FD reconstructions. FD detection can be done in stereo.
Pierre Auger Observatory (PAO)	Hybrid observatory. Combination of SD plastic scintillators of area $3000\,\mathrm{km}^2$ and 4 FD telescope sites, with a total of 27 FD telescopes.	Exposure superior to all other observatories. Stereo observation of FD is not possible. The detector stations are most sensitive to the muon component of the air shower, which is worst understood, leading to an increased uncertainty in the results.

Former observatories are AGASA (Chiba et al. 1992), the first cosmic ray observatory, and HiRes (Boyer et al. 2002), which introduced the FD method.

The focus of this research will be on data from the PAO. The PAO is a CR observatory located in Argentina, at a longitude of 69.5° W and a latitude 35.2° S. The observatory has SD plastic scintillators of a total area of 3000 km^2 and 4 FD telescope sites, with a total of 27 FD telescopes. PAO measures UHECR arrival directions with an uncertainty of $\sim 1\,\mathrm{deg}$ and arrival energies with a relative uncertainty of $\sim 12\%$ (Letessier-Selvon et al. 2014).

The primary advantage of PAO is its unprecedented exposure, which allows for very high statistics. In the past, the Pierre Auger collaboration has produced some of the most interesting results in the field of UHECR research. In this work, we analyze a sample of 69 UHECRs with observed energies $E_{\mathrm{obs}} \geq E_{\mathrm{thres}} = 5.5 \times 10^{19}\,\mathrm{eV}$ that were recorded by the PAO in the period January 2004 - November 2009, as documented in Abreu et al. (2010). The PAO's total exposure of this data-set is

$\epsilon_{\rm tot} = 20,370\,{\rm km}^2\,{\rm sr\,yr}$. This sample has previously been analyzed by a number of studies, such as Abreu et al. (2010) and Soiaporn et al. (2013).

Recently, the PAO reported a study of a greater sample of 231 events with $E_{\rm obs} \geq 5.2 \times 10^{19}\,{\rm eV}$, corresponding to an exposure of $66,000\,{\rm km}^2\,{\rm sr\,yr}$. This sample was recorded in the period January 2004 - March 2014 (Aab et al. 2015).

PAO's relative exposure per unit solid angle, $d\epsilon/d\Omega$, is illustrated in Figure 1.4 The relative exposure is directly proportional to $\Pr(\det|\boldsymbol{r})$, the probability that a UHECR will be detected if it arrives from a direction given by unit vector $\boldsymbol{r}$, but is normalized so that $\int (d\epsilon/d\Omega)\, d\Omega = \epsilon_{\rm tot}$. PAO can reliably detect UHECRs arriving from directions within 60° from its observatory zenith. It should be noted that the integrated exposure depends only on the declination, as one can see from the figure. This is due to the fact that the observatory zenith traces a circular path on the sky due to the Earth's rotation, so that the variation of the exposure with right ascension disappears as one integrates over time.

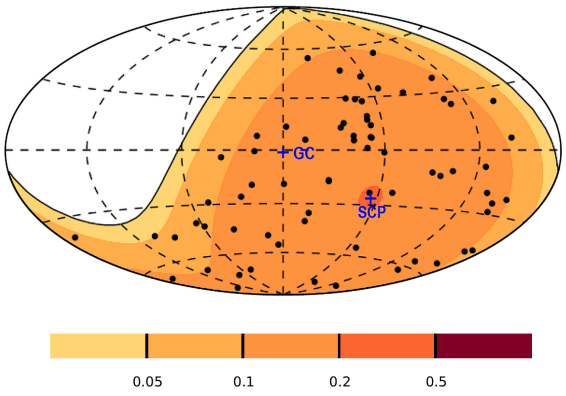


Figure 1.4: Relative PAO exposure in Galactic coordinates. The arrival directions of the 69 UHECRs are shown as black points. The Galactic centre (GC) and south celestial pole (SCP) are indicated.

1.5 Principal observables and experimental results

1.5.1 Composition

1.5.1.1 Detection

The composition of cosmic rays is inferred from the depth of the maximum EAS development $X_{\rm max}$, and the extent of its fluctuations. $X_{\rm max}$ is usually determined by

monitoring the shower development by the use of FD. In addition to using FD, Yakutsk determines $X_{\max}$ by measuring the Cherenkov light distribution on the ground (Berezhko et al. 2012), while PAO has the additional method of muon detection on the ground (García-Gámez 2013).

1.5.1.2 Results

There are no clear results on the cosmic ray composition at present. All experiments indicate an almost pure proton composition (see, e.g., Abbasi et al. 2010a; Jui et al. 2012; Berezhko et al. 2012), while Auger suggests an increase in UHECR mass for arrival energies $E_{\rm arr} \gtrsim 4\times10^{18}\,\mathrm{eV}$, culminating at iron for $E_{\rm arr} = 10^{19.5}\,\mathrm{eV}$ (Abraham et al. 2010). There is no clear resolution to this question, and various arguments have been put forward about the plausibility of different possible scenarios given the data and the theoretical background.

The origin of the disagreements has not been settled. It is possible that the primary composition depends strongly on the direction in the sky, if the UHECRs are produced in nearby sources. PAO is the only observatory that is located in the Southern Hemisphere, which could explain the difference in composition that PAO observes. This hypothesis is challenged by the fact that, in 2012, the PAO collaboration conducted an analysis of events in the equatorial part of the Northern Hemisphere (which is observable by PAO), and found no difference to its Southern Hemisphere results (Barcikowski et al. 2013). The northern observatories do not yet have sufficient statistics to conduct such an analysis. Another possibility for the observed difference in the result are methodological differences in the analysis of $X_{\max}$, which could lead to systematic errors. More data are needed to resolve this issue.

1.5.2 Arrival energy

1.5.2.1 Detection

The arrival energies of the UHECRs are reconstructed indirectly from properties of the EAS. In SD, the energy is inferred from the lateral distribution of the signal. FD determines the energy of the electrons and positrons in the shower core based on the fluorescent yield.

These methods of energy determination are model dependent, and there is a significant uncertainty that is caused by the fluctuations in the particle interactions in the upper atmosphere. In addition, there are complicating factors that relate to the respective method of observation. In SD, there is uncertainty in the modelling

of the lateral distribution as a function of energy. The FD uncertainty lies in the value of the fluorescent yield, as well as the estimation of the fraction of the energy that resides outside the shower core. For a concise summary of energy detection and related issues, see, e.g., Troitsky (2013).

1.5.2.2 Spectrum

Figure 1.5 shows the energy spectra for AGASA (Takeda et al. 2003), Yakutsk (Egorova et al. 2004), HiRes I (Abbasi et al. 2008b), PAO (Abreu et al. 2011) and TA (Abu-Zayyad et al. 2013).

It is clear from the figure that the energy spectra are not the same. The reasons suggested for this are systematic errors due to the choice of model. It has been attempted to compensate for the systematic error and to check the consistency of spectral shapes by adjusting the energy scales of the individual experiments, within energy scale uncertainties, to a common scale. The results were that in a broad interval, $10^{17.5} - 10^{19.5}$ eV the shifted spectra coincide, which is a strong confirmation of the approach. Spectra coincide worse at higher energies. The shifted spectra are also shown on Figure 1.5.

HiRes and PAO both have reported an exponential cutoff in their spectra (Abbasi et al. 2008b; Abraham et al. 2008b) at the highest energies, in agreement with the prediction of a GZK cutoff. While AGASA does not feature a suppression of the spectrum (Takeda et al. 1998), the AGASA experiment only has 2 ultra-high energy events (after rescaling), which turns its observations statistically insignificant. It is important to note that it cannot be said with certainty that the suppression of the spectrum is due to the GZK effect, or the maximum acceleration energy of the UHECRs.

1.5.3 The arrival directions of cosmic rays

1.5.3.1 Measurement

The arrival directions of the UHECRs can be inferred both from the SD and the FD data. SD reconstructs the direction from the times at which the individual SD detectors were triggered by the EAS front. FD can determine the arrival direction directly by observing the shower core in stereo. If stereo is not available, the development of the signal in time is used.

The arrival directions are the least model dependent observables that can be inferred from the EAS, as the problem is essentially geometrical.

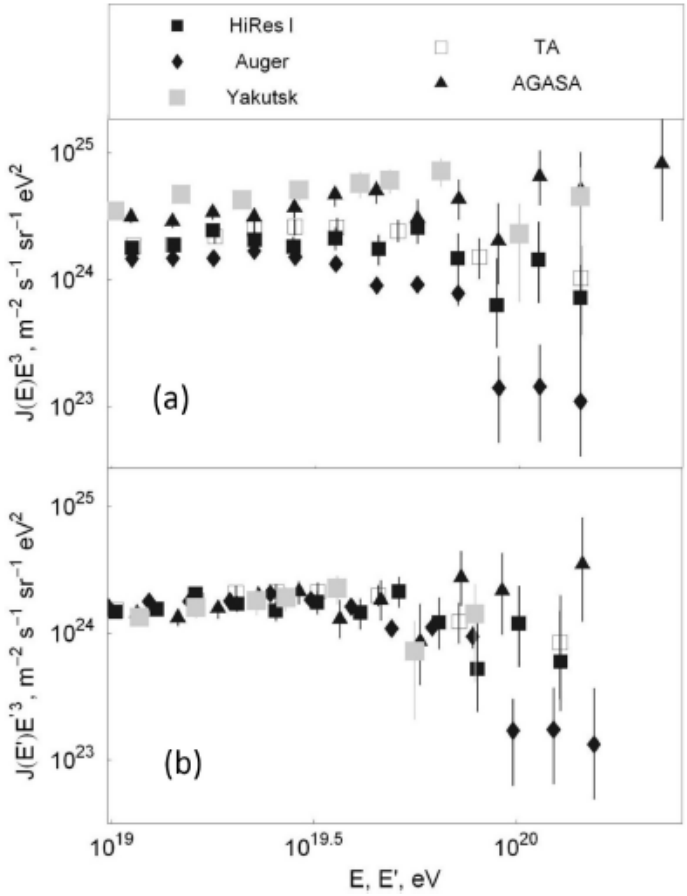


Figure 1.5: The UHECR energy spectra, measured by AGASA (Takeda et al. 2003), Yakutsk (Egorova et al. 2004), HiRes I (Abbasi et al. 2008b), PAO (Abreu et al. 2011) and TA (Abu-Zayyad et al. 2013). The top and bottom panels show the spectra before and after the energy scale-adjustments. Figure taken from Troitsky (2013).

1.5.3.2 Significance

The UHECR arrival directions could provide invaluable insights into the origin of UHECRs. Cosmic rays at ultra-high-energies are expected to be deflected only by a few degrees, so that the arrival directions should be closely tied to the source directions. As discussed in Section 1.3.1, UHECRs are expected to have come from a limited GZK horizon of $\sim 100\,\mathrm{Mpc}$. Inside this horizon, the number of potential sources is limited, and is distributed inhomogeneously. As a result, it should be possible to find a correlation between the arrival directions of UHECRs and catalogues of potential sources, and also to determine whether the flux of UHECRs on Earth matches what would be expected if the UHECR sources were following the local large scale structure of the universe.

1.5.3.3 Previous studies

A substantial amount of research has been done into the arrival directions of UHECRs. A number of attempts have been made to find correlations between UHECR arrival directions and catalogues of possible sources. Cross-correlation studies have been conducted with galaxy catalogues, such as the Two Micron All-Sky Survey (2MASS) Redshift Survey (2MRS) (Abraham et al. 2009; Abbasi et al. 2010b), as well as specific types of objects such as active galactic nuclei (AGNs) (Abraham et al. 2007; Abraham et al. 2008a; George et al. 2008; Pe'Er et al. 2009; Watson et al. 2011) and BL Lacertae objects (BL LAcs) (Tinyakov & Tkachev 2001).

Overall, no clear consensus has been reached. Different studies have reported different degrees of correlation, depending on the statistical approach, the UHECR sample, and the population of source candidates that was used.

Perhaps the best-known and most significant correlation that has been measured was reported by the PAO collaboration in 2007. PAO reported a correlation between the arrival directions of UHECRs with energies $E > 5.6 \times 10^{19}\,\mathrm{eV}$ and the positions of nearby active galaxies (Abraham et al. 2007). This was taken as evidence that UHECRs originate either in those galaxies or in objects that are similarly distributed. The result has been confirmed by Yakutsk (Ivanov 2009), but not by HiRes (Abbasi et al. 2008a) and by TA (Abu-Zayyad et al. 2012). A later study by the Pierre Auger collaboration has shown a much weaker correlation than before (Abreu et al. 2010). The PAO Collaboration's recent study Aab et al. (2015) found no statistically significant evidence of anisotropy.

Attempts to associate UHECRs with specific sources are hampered to some degree by large magnetic deflections, possibly transient sources and incomplete catalogues. An alternative approach is based on the idea that if the UHECR sources are distributed inhomogeneously inside the GZK horizon, it should be possible to detect self-clustering in the UHECR arrival directions, independent of any source catalogue. Numerous studies have been conducted investigating this question, but no clear consensus has been reached. AGASA reported clustering of events with $E > 40\,\mathrm{EeV}$, at an angular scale of 2.5° (Hayashida et al. 1996). This result was not confirmed by PAO or TA. PAO reported a significant excess of event pairs for energies $E > 57\,\mathrm{EeV}$, for a range of separation angles $9° < \theta < 22°$ (Mollerach et al. 2009). TA reported a largely isotropic distribution, with some hint of groupings of events at the highest energies, at a scale $20° < \theta < 30°$ (Abu-Zayyad et al. 2012). The statistical significance of these high energy results may not be sufficient.

In 2012, the PAO Collaboration conducted an analysis searching for self-clustering in the Auger data (Abreu et al. 2012), using three statistical methods

based on correlation functions. No strong evidence of non-uniformity was found based on the p-values obtained under the null hypothesis of no clustering.

Most of the previous work on these questions has been based on frequentist methodologies. These are known to have a number of limitations, however (see Chapter 2). Two recent studies have made the first steps to developing Bayesian methodologies for the analysis of the UHECR arrival directions: Watson et al. (2011) and Soiaporn et al. (2013). The authors develop multi-level models of cosmic ray injection, propagation and detection. The evidence for these models is then weighed against the evidence for the isotropic scenario.

Soiaporn et al. (2013) analyzed a sample of 69 UHECRs from the PAO, and reported evidence for a small but nonzero fraction of their UHECR sample to have originated at 17 nearby AGNs, of the order of a few percent to 20%. Watson et al. (2011) analyzed 27 PAO events, and, for fiducial values of their model parameters, reported that a fraction $F_{\rm AGN} = 0.15^{+0.10}_{-0.07}$ of their UHECR sample originate from AGNs in the Veron-Cetty & Veron (VCV) catalogue. These two studies and their results are discussed in greater detail in Chapter 4.

Chapter 2

Bayesian methods

2.1 Cox's theorem

In Bayesian probability theory, a very simple intuitive concept of probability is used. Probability is regarded as

"a measure of the degree of a belief in a proposition conditional on certain information being true."

Cox's theorem states that any framework for representing degrees of belief that satisfies certain criteria of reasonableness must be equivalent to probability theory. This remarkable result was first discussed by Cox (Cox 1946), and later developed by other authors such as Jaynes (Jaynes 2003). Here, we follow the analysis of Jaynes and sketch out an overview of Cox's theorem. We adopt a notation in which we refer to a degree of belief in proposition A conditional on proposition B being true as $(A|B)$. We also refer to $(A|B)$ as "plausibility of A given B". We denote the plausiblity of propositions A and B both being true given that C is true as $(AB|C)$. The Cox-Jaynes argument is based on three desiderata, which can be listed as

(i) A degree of belief is represented as a single real number. This reflects the fact that degrees of belief have a natural ordering: $(A|D) > (B|D)$ and $(B|D) > (C|D)$ implies that $(A|D) > (C|D)$.

(ii) The degrees of belief are in qualitative correspondence with common sense. This implies that for infinitesimal changes of A, $(A|B)$ changes infinitesimally. It also implies that if C is updated to C' in such a way that $(A|C') > (A|C)$ and $(B|AC') = (B|AC)$, then $(AB|C') > (AB|C)$. Also, if we define $\bar{A}$ as the negation of proposition A, then our degree of belief in $\bar{A}$ is decreased as our degree of belief in A is increased: $(\bar{A}|C') < (\bar{A}|C)$. Consistency with common sense also means that in the limiting case where the truth or falsehood of propositions is known with certainty, a system for reasoning about uncertainty should be equivalent to Aristotelian deductive logic.

(iii) Reasoning about degrees of belief is consistent. This means that (a) If there are more than one way to reach a conclusion, then all of those ways must reach the same result. (b) All of the available information has to be taken into account. (c) Equivalent states of knowledge must be represented by the same plausibility assignments.

We will now sketch out the proof that plausibilities that accord with these desiderata follow the product and sum rules of probability theory. The product rule of probability can be derived by considering the plausibility of two propositions A

and B both being true, given proposition C. Correspondence with common sense requires that this plausibility depends only on the plausibility of A being true, and the plausibility of B being true given that A is true:

$$(AB|C) = F[(A|C), (B|AC)], \tag{2.1}$$

for some function F. To satisfy the desideratum of consistency, F must satisfy the associativity equation:

$$F[F(a, b), c] = F[a, F(b, c)]. \tag{2.2}$$

It can be shown (see, e.g., Jaynes (2003)) that for this relation to be correct, there must be a function w such that

$$w(AB|C) = w(A|C)w(B|AC), \tag{2.3}$$

and that this function can be chosen without loss of generality to be a monotonic function that varies from 0 to 1. Thus, by rescaling the plausibility measure to $w(A|B)$ we have shown that the plausibilities follow the product rule of probability.

To derive the sum rule, we first consider that consistency with common sense requires that the conditional plausibility of the negation of A given C needs to be some function of the conditional plausibility of A given C:

$$w(\bar{A}|C) = S[w(A|C)], \tag{2.4}$$

for some function S. Jaynes shows that this equation, together with the product rule and the desideratum of consistency, leads to the equation

$$x \times S\left[\frac{S(y)}{x}\right] = y \times S\left[\frac{S(x)}{y}\right], \tag{2.5}$$

for $x \equiv w(A|C)$ and $y \equiv w(B|C)$. The unique solution to this equation is given by

$$S(x) = (1 - x^m)^{1/m}, \tag{2.6}$$

where $0 \leq x \leq 1$ and $0 < m < \infty$. From this expression, it follows that

$$w^m(A|B) + w^m(\bar{A}|B) = 1. \tag{2.7}$$

We can rescale our plausibility measure to a new function,

$$\Pr(A,\ B|C) \equiv w^m(AB|C), \tag{2.8}$$

so that we obtain the sum rule of probability,

$$\Pr(A|B) + \Pr(\bar{A}|B) = 1. \tag{2.9}$$

From Equation 2.3, it follows that

$$\Pr(A,\ B|C) = \Pr(A|C)\Pr(B|A,\ C) \tag{2.10}$$

Thus, any form of plausible reasoning that accords with the Cox-Jaynes desiderata can be shown to adhere to the product and sum rules of probability theory.

2.2 Bayes' theorem

Bayes' theorem stands at the core of Bayesian inference. The expression is derived from the axioms of probability theory, and is uncontroversial in itself. More contentious is the question of whether it can be used as a prescription for reasoning about uncertainty.

Let us use the notation we have established in Section 2.1, and let us consider propositions A, B. We can use the fact that

$$\Pr(A, B) = \Pr(B, A) \tag{2.11}$$

to rewrite the product rule of Equation 2.10:

$$\Pr(A)\Pr(B|A) = \Pr(B)\Pr(A|B) \tag{2.12}$$

$$\Pr(B|A) = \frac{\Pr(A|B)\Pr(B)}{\Pr(A)}. \tag{2.13}$$

This is Bayes' theorem. This simple result becomes very insightful if we replace A by the data $\boldsymbol{d}$, replace B by the hypothesis H the probability of which we are assessing, and incorporate the background information I into the equation:

$$\Pr(H|\boldsymbol{d}, I) = \frac{\Pr(\boldsymbol{d}|H, I)\Pr(H|I)}{\Pr(\boldsymbol{d}|I)}. \tag{2.14}$$

This expression connects the prior probability $\Pr(H|I)$ of the hypothesis with the

posterior probability $\Pr(H|\boldsymbol{d}, \mathrm{I})$. It is a prescription for how to correctly update one's state of knowledge in light of new evidence. $\Pr(\boldsymbol{d}|H, I)$ is the sampling distribution, and, when it is considered as a function of the hypothesis (and therefore not as a probability distribution) is called the likelihood function $\mathcal{L}(H)$. $\Pr(\boldsymbol{d}|I)$ is the normalization factor or marginal likelihood:

$$\Pr(\boldsymbol{d}|I) = \sum_H \Pr(\boldsymbol{d}|H, I)\Pr(H|I), \tag{2.15}$$

where the sum is over all possible hypotheses. The marginal likelihood plays a crucial role in Bayesian model comparison, and is discussed further in Section 2.5.

2.3 Priors

Prior information is a core feature of Bayesian statistics. A fundamental principle of Bayesian probability is that inference is impossible without assumptions. The choice or prior probability reflects the initial assumptions and the state of knowledge of the scientist about the problem at hand, before analysis of the data. While two scientists can start with different information and assumptions, and therefore different priors, as long as the priors have non-zero values in regions where the likelihood is high ("diffuse" priors), repeated application of Bayes' theorem will lead the posteriors of the scientists to converge with sufficient data. This phenomenon is referred to as "stability" or "robustness".

Often the prior that is used is grounded in information from previous experiments, so that the prior of one experiment is the posterior of another. If no obvious information is available, a non-informative prior is used. Such a prior can be an ignorance prior, which reflects indifference with respect to the symmetries of the problem in question. When that is not possible, often a simple diffuse reference prior is chosen. While the priors that are used need not necessarily be normalizable ("proper") themselves, it is important that they are chosen in such a way that the posterior resulting from the Bayesian inference is proper.

An example of a widely used prior is a uniform prior

$$\Pr(\theta) \propto \Theta(\theta - \theta_{\min}) \cdot \Theta(\theta_{\max} - \theta), \tag{2.16}$$

where Θ is the Heaviside step function and $\theta_{\max}$ and $\theta_{\min}$ are the upper and lower limits of the uniform prior over the parameter θ. Uniform priors reflect indifference with respect to all possible values of the parameter between the two limits. Another

example is Jeffrey's prior $\Pr(\theta) \propto \theta^{-1}$, which reflects indifference with respect to all orders of magnitude of the parameter that is known to be positive.

We can illustrate the robustness of Bayesian inference by using a coin example discussed by Jaynes (Jaynes 2003). Suppose that we are given a coin that is known to be biased, but not by how much and in what direction. Now let us define the parameter H, where a coin with H=1 denotes a coin that comes up heads on every toss, $H = 0$ comes up tails, and $H = 0.5$ is a fair coin. Different people may start out with different priors about the bias of the coin. One person might have reason to think that the coin is heavily biased in one direction, so they could start with a prior with two peaks at $H = 0$ and $H = 1$. Another person may start with a flat prior, giving equal probability to each possible value of H. Then these two people proceed tossing the coin, and update their priors. The likelihood that N tosses will produce R instances of heads can be shown to be

$$\Pr(\boldsymbol{d}|H, I) \propto H^R(1 - H)^{N-R}. \tag{2.17}$$

Figure 2.1 shows the evolution of the priors of the two people, and shows that through iterative Bayesian inference, they arrive at the same posterior probability distribution.

2.4 Bayesian parameter inference

Bayesian parameter inference is a methodology for arriving at estimates of a parameter given data obtained through a measurement.

First, a model M is specified. The model contains a vector of parameters, θ, that correspond to a set of hypotheses that are part of the model. For each of the parameters, a prior is specified that summarizes the initial state of knowledge about the parameters. For the measurement in question, a likelihood function is constructed.

Thus, the posterior probability for θ is

$$\Pr(\theta|\boldsymbol{d}, M) = \frac{\Pr(\boldsymbol{d}|\theta, M)\Pr(\theta|M)}{\Pr(\boldsymbol{d}|M)}. \tag{2.18}$$

As $\Pr(\boldsymbol{d}|M)$ is a constant that here merely serves the role of a normalization parameter we can write

$$\Pr(\theta|\boldsymbol{d}, M) \propto \Pr(\boldsymbol{d}|\theta, M)\Pr(\theta|M). \tag{2.19}$$

Often in this kind of analysis, nuisance parameters are present that are not of

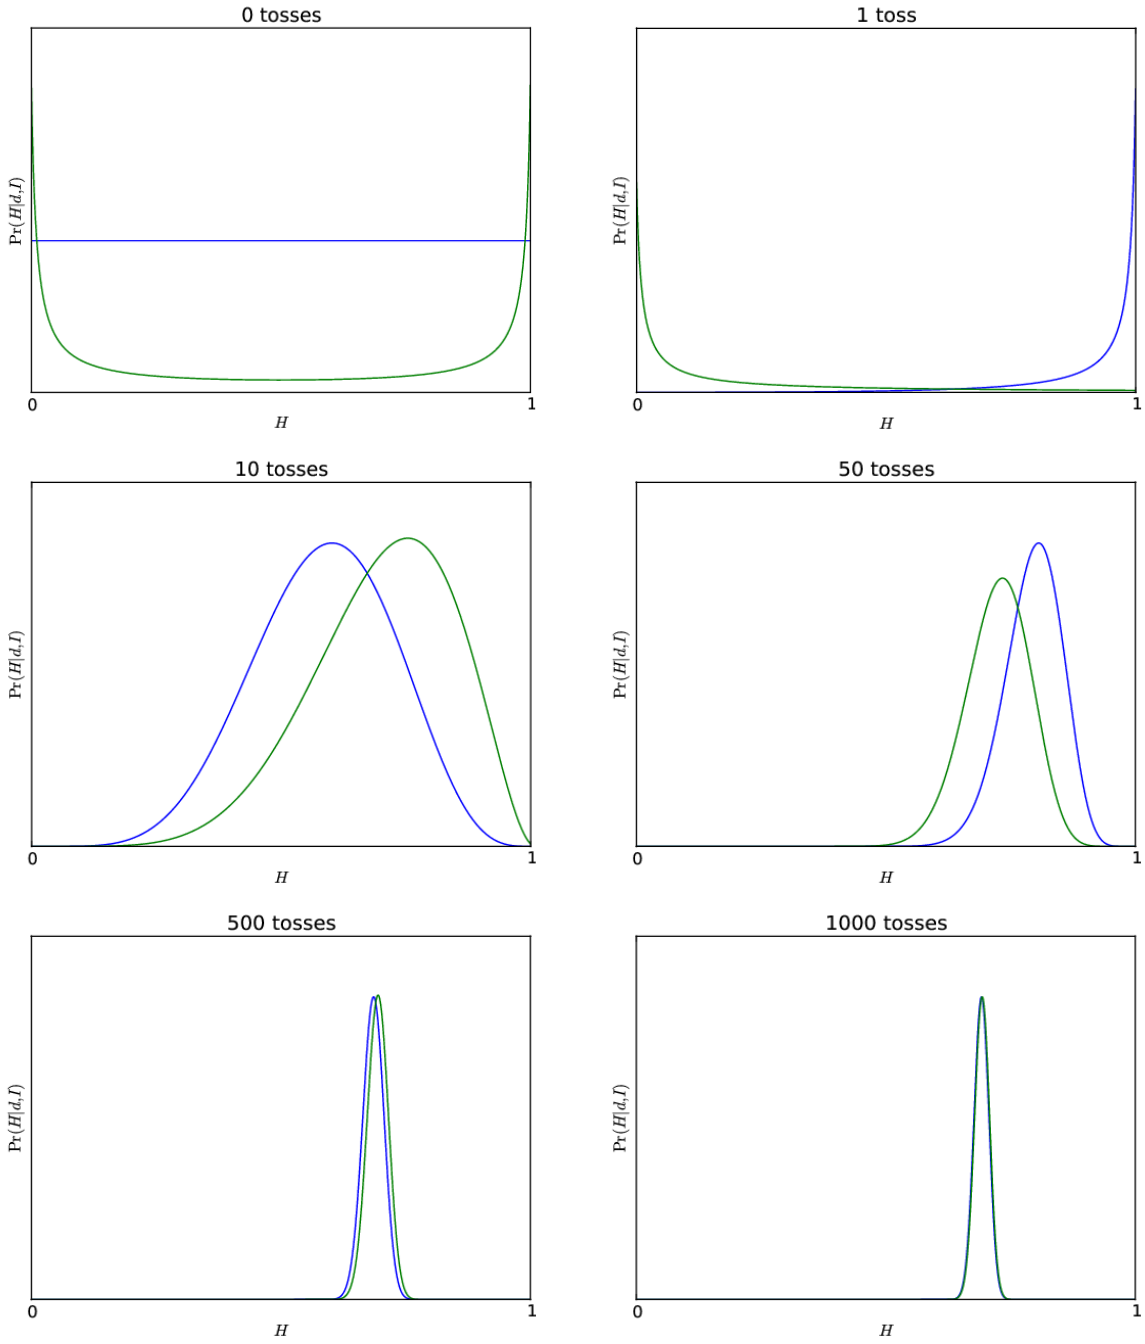


Figure 2.1: Robustness of Bayesian priors demonstrated with a coin example. Parameter H is a way to quantify the bias of the coin. Coins with $H = 0$ and $H = 1$ come up tails and heads each time, respectively, while $H = 0.5$ denotes a fair coin. In this example, two people start off with different priors: One believes that the coin is heavily biased in one direction or the other, so that their prior is peaked around $H = 0$ and $H = 1$. The other gives equal probability to each value of H, and starts with a flat prior. As they proceed tossing the coin, their posteriors begin to converge, and for a sufficient number of tosses become indistinguishable.

immediate interest, but which need to be accounted for. Nuisance parameters are generally handled by means of marginalization. If θ is divided into parameters that are of immediate interest ϕ, and the nuisance parameters ψ, the marginalized posterior can be written as

$$\Pr(\phi|\boldsymbol{d}, M) \propto \int_{\psi} \Pr(\boldsymbol{d}|\phi, \psi, M)\Pr(\phi, \psi|M)\mathrm{d}\psi. \tag{2.20}$$

Generally, the above procedure cannot be carried out analytically, so that numerical techniques are used to evaluate the likelihood.

We can obtain a point estimate of the parameter by summarizing the center of the posterior, typically using its mean or its mode. We can also obtain an interval estimate of the parameter. In Bayesian inference, a credible region is an interval that contains the true value of the parameter with a certain probability. For example, a 95% credible region of θ contains θ with 95% probability.

For further reading on Bayesian parameter inference see Sivia & Skilling (2006), Gregory (2005) and Gelman et al. (1995), for advanced application of parameter inference to astrophysical and cosmological problems, see Feigelson & Babu (2003), M. Hobson (2008).

2.5 Model comparison

Bayesian model comparison is a framework for choosing between two models based on some observed data $\boldsymbol{d}$. The process takes into account how well the models describe the data, as well as the complexity of the models. A highly complex model can always be devised to give a perfect fit to a set of data, so that model comparison attempts to strike an optimum balance between models that fit the data and models that are simple, thus satisfying the Occam's razor principle.

The models being weighed against each other contain a number of parameters and their prior distributions. The Occam's razor effect penalizes models both for the number of their parameters as well as the prior ranges of the parameters.

The crucial quantity in model comparison is the marginal likelihood $\Pr(\boldsymbol{d}|M)$, which in cosmology is often called the "Bayesian evidence". For a model M with some nuisance parameter θ, we can write

$$\Pr(\boldsymbol{d}|M) \equiv \int_{\theta} \Pr(\boldsymbol{d}|\theta, M)\Pr(\theta|M)\mathrm{d}\theta. \tag{2.21}$$

When there are two models M_0 and M_1, the quantity of interest is the ratio of

posterior probabilities:

$$\frac{\Pr(\mathrm{M}_0|\boldsymbol{d})}{\Pr(M_1|\boldsymbol{d})} = B_{01}\frac{\Pr(M_0)}{\Pr(M_1)}, \tag{2.22}$$

where B_{01} is the evidence ratio, or Bayes factor:

$$B_{01} = \frac{\Pr(\boldsymbol{d}|M_0)}{\Pr(\boldsymbol{d}|M_1)}. \tag{2.23}$$

The Bayes factor is the central quantity in Bayesian model comparison. The magnitude of the Bayes factor indicates how much our degree of belief about the relative viability of the two models has changed with the introduction of the new data.

2.5.1 Model comparison with poorly specified prior parameters

Bayesian model comparison encounters a difficulty when it deals with situations that lack strongly motivated priors, and where the natural uninformative priors are improper and non-normalizable. Such situations are encountered often in astronomy and cosmology. It has been argued that due to this problem, Bayesian model comparison is not applicable in cosmology (Efstathiou 2008 and Jenkins & Peacock 2011), which would have the implication that there is no self-consistent way of choosing between cosmological models.

However, methods have been developed to deal with such cases. One possibility is to try to motivate the prior from the data itself. At first glance, it may seem that this approach is not self-consistent, as it would lead to a circular analysis where the prior is derived from the data, and then that prior is used in the model comparison that analyzes the same data. This problem can be resolved by using a subset of the data as a training sample to derive a parameter prior. These parameter priors can then be used to calculate the marginal likelihood, and then the rest of the data can be used in a standard model comparison. The parameter prior thereby is the posterior derived by using the training sample:

$$\Pr(\boldsymbol{d}_2|\boldsymbol{d}_1, M) = \int_\theta \Pr(\boldsymbol{d}_2|\theta, M)\Pr(\theta|\boldsymbol{d}_1, M)\mathrm{d}\theta, \tag{2.24}$$

where

$$\Pr(\theta|\boldsymbol{d}_1, M) = \frac{\Pr(\theta|M)\Pr(\boldsymbol{d}_1|\theta, M)}{\int_\theta \Pr(\theta|M)\Pr(\boldsymbol{d}_1|\theta, M)\mathrm{d}\theta}, \tag{2.25}$$

and $\Pr(\theta|M)$ is a highly uninformative prior which doesn't need to be normalizable.

Methods of this kind have previously been discussed at great length in e.g.

Spiegelhalter & Smith (1982), Aitkin (1991), O'Hagan (1991), O'Hagan (1995) and Ghosh et al. (2006), and were shown to satisfy Cox' self-consistency requirements. One open question (discussed in e.g. Spiegelhalter & Smith 1982; O'Hagan 1995) is the relative merits of using a low or high fraction of the data to derive the priors (leaving, respectively, a high or low fraction for the model comparison).

2.5.2 Calculation of the evidence

The calculation of the evidence term requires taking a multi-dimensional integral over the entire parameter range, and generally is not analytically tractable. Additional difficulties arise when the likelihood is multimodal (a situation that often arises for mixture models), or is sharply peaked in parameter space. A number of numerical techniques have been developed for the computation of likelihoods, such as simulated annealing (see e.g. Gregory 2005; Press 2007; Clyde et al. 2007) and nested sampling (see e.g. Skilling 2006).

For some situations, approximations of the Bayes factor have been derived. One such case is relevant for the work that is discussed in the later chapters: In some situations, one of the models of the Bayesian model comparison is nested in the other. This means that the more complex model M_2 becomes identical to the simpler model M_1 for certain values of its parameters. Let us say that M_2 has parameters (ϕ,ψ) and M_1 has parameter ϕ, and that M_2 is identical to M_1 when $\psi = 0$. In such situations, the expression for the Bayes factor reduces to

$$B_{12} = \frac{\int_\phi \Pr(\phi|M_1)\Pr(\boldsymbol{d}|\phi, M_1)\,\mathrm{d}\phi}{\int_{\psi,\phi} \Pr(\psi,\phi|M_2)\Pr(\boldsymbol{d}|\psi,\phi,M_2)\,\mathrm{d}\psi\,\mathrm{d}\phi}. \tag{2.26}$$

It can be shown (Dickey 1971) that this expression can be simplified to yield

$$B_{12} = \frac{\Pr(\psi = 0|\boldsymbol{d}, M_2)}{\Pr(\psi = 0|M_2)}. \tag{2.27}$$

This expression is known as the Savage-Dickey Density Ratio, or SDDR. Qualitatively, this expression means that the nested model is preferred if, within the context of the more complicated model, the data result in an increased probability that $\psi = 0$.

2.5.3 Example case: On/Off measurements

Here we discuss the application of Bayesian model comparison methods to the On/Off problem of astrophysics. In high-energy astrophysics, when a measurement

is taken of the number of counts coming from a source of interest, often an auxiliary measurement is made by pointing the detector off-source. These are called the On and Off measurements, respectively. The counts that are detected in the Off measurement are thereby produced solely by the background rate R, while the counts in the On measurement are produced by both the source and the background rates Γ and R. From these two measurements, the source rate can then be estimated (e.g. Gregory 2010).

The likelihood for these kinds of measurements is the product of the Poisson likelihoods of the On and Off measurements:

$$\Pr(N_{\rm on}, N_{\rm off}|\Gamma, R) = \frac{(R\,T_{\rm off})^{N_{\rm off}}\exp(-R\,T_{\rm off})}{N_{\rm off}!} \times \frac{[(\Gamma+R)\,T_{\rm on}]^{N_{\rm on}}\exp[-(\Gamma+R)\,T_{\rm on}]}{N_{\rm on}!}, \tag{2.28}$$

where Γ and R and the source and background rates, $N_{\rm on}$ and $N_{\rm off}$ are the numbers of counts on and off source, and $T_{\rm on}$ and $T_{\rm off}$ times the detector spends on and off the source. We choose $N_{\rm on} = 54$ and $N_{\rm off} = 15$. Thus, $N_{\rm on} + N_{\rm off} = 69$, which is equivalent to the number of UHECR events from the PAO that are the focus of later chapters (Chapter 3 and 4).

We adopt a uniform prior over Γ and R, with $\Gamma \geq 0$, $R \geq 0$. We use a variable width parameter s to vary the width of the prior. The expression for the prior can be written as

$$\Pr(\Gamma, R) = \begin{cases} \frac{1}{s^2\Gamma_{\rm max}R_{\rm max}} & \text{if } R < sR_{\rm max},\ \Gamma < s\Gamma_{\rm max} \\ 0 & \text{otherwise} \end{cases}, \tag{2.29}$$

where $\Gamma_{\rm max}$ and $R_{\rm max}$ are chosen in such a way that when $s = 1$, the prior covers the 99.7% credible region implied by the combination of the likelihood and an infinitely broad uniform prior. This is illustrated in the bottom panel of Figure 2.2. The panel shows the On/Off likelihood, and shows prior regions for values of the hyperparameter $s = 0.1$, 1, 2. The scaling of $\Gamma_{\rm max}$ and $R_{\rm max}$ depends on the likelihood that is being analyzed, which enables the comparison of Bayes factors for different likelihoods, as will be seen in Section 2.5.4. The top three panels of Figure 2.2 show posteriors $\Pr(\Gamma, R|N_{\rm on}, N_{\rm off})$ for the same s values. As the priors are flat, the posteriors are equivalent to the likelihood in the prior region, normalized over the prior region.

Now we would like to conduct a model comparison: we would like to compare a model M_1, where all of the counts are derived purely from the background component, with model M_2, where the counts are derived from a combination between

background and source components. This is an example of nested models, because when the source rate is zero, M_2 becomes equivalent to M_1. According to Equation 2.27, the expression for the Bayes factor is

$$B_{12} = \frac{\Pr(\Gamma = 0|\boldsymbol{d}, M_2)}{\Pr(\Gamma = 0|M_2)}. \tag{2.30}$$

$\boldsymbol{d}$ here corresponds to $N_{\rm on}, N_{\rm off}$. We can use the posteriors shown in Figure 2.2 to illustrate the dependence of this Bayes factor on the hyperparameter s. For $s > 1$, the numerator $\Pr(\Gamma = 0|\boldsymbol{d}, M_2)$ remains constant. This is because when $s = 1$ the posterior $\Pr(\Gamma, R|\boldsymbol{d}, M_2)$ simply becomes the normalized likelihood, and remains like that for all higher values of s, as the figure shows. The denominator $\Pr(\Gamma = 0|M_2)$ falls linearly with s, because we chose a uniform prior. Thus, we expect that for $s > 1$, B_{12} increases linearly with s.

For lower values of s, the behaviour of B_{12} is more complicated. As can be seen in the top panel of Figure 2.2, most of the posterior becomes concentrated at the highest values of Γ and R. The expression for the likelihood, for very low values of R and Γ, reduces to

$$\Pr(N_{\rm on}, N_{\rm off}|\Gamma, R) = \frac{(R\,T_{\rm off})^{N_{\rm off}}}{N_{\rm off}!} \times \frac{[(\Gamma + R)\,T_{\rm on}]^{N_{\rm on}}}{N_{\rm on}!}. \tag{2.31}$$

For $N_{\rm on} = 54$ and $N_{\rm off} = 15$, this function becomes extremely steep in Γ and R, as Figure 2.2 shows. Thus, B_{12} tends to zero as $\Pr(\Gamma = 0|\boldsymbol{d}, M_2) \ll 1$. Figure 2.4 shows the behaviour of $B_{21} = 1/B_{12}$ as a function of s, and illustrates the behaviour of the Bayes factor for high and low values of the hyperparameter.

For the On/Off problem, a standard expression for the Bayes factor has been derived (Gregory 2010), and can be written as

$$\begin{aligned} B_{21} = {} & \frac{N_{\rm on}!}{\Gamma_{\rm max}\,T_{\rm on}\gamma[(N_{\rm on} + N_{\rm off} + 1), R_{\rm max}(T_{\rm on} + T_{\rm off})]} \\ & \times \sum_{i=0}^{N_{\rm on}} \frac{\gamma[(N_{\rm on} + N_{\rm off} - i), R_{\rm max}(T_{\rm on} + T_{\rm off})]}{i!(N_{\rm on} - i)!} \\ & \times \gamma(i + 1, \Gamma_{\rm max}\,T_{\rm on}) \left(1 + \frac{T_{\rm off}}{T_{\rm on}}\right)^{i}, \end{aligned} \tag{2.32}$$

where $\gamma(s, x)$ is the lower incomplete gamma function, defined here as

$$\gamma(s, x) = \int_0^x t^{s-1} e^{-t} {\rm d}t. \tag{2.33}$$

This expression produces a dependence of the Bayes factor on s that is exactly equivalent to the Bayes factor that is calculated by computing the SDDR.

2.5.4 Example case: Gaussian likelihood

We consider the case of a Gaussian likelihood given by

$$\Pr(\boldsymbol{d}|\Gamma, R) = \frac{1}{2\pi\sigma_\Gamma\sigma_R}\exp\left[-\frac{(\Gamma-\Gamma_\mu)^2}{2{\sigma_\Gamma}^2}\right]\left[-\frac{(R-R_\mu)^2}{2{\sigma_R}^2}\right], \tag{2.34}$$

where Γ_μ and R_μ are the coordinates of the likelihood mean, σ_Γ and σ_R are the standard deviations on the two parameters. The priors were selected in the same way as for the On/Off case. Figure 2.3 shows the likelihood, as well as prior regions for values of $s = 0.1$, 1, 2, and posteriors for those regions.

For high values of s, the posteriors looks very much like the posteriors for the On/Off case, so that we expect the same behaviour for the Bayes factor, including the linear behaviour for $s > 1$.

For lower values of s, the posteriors differ. In the Gaussian case, the posterior is not concentrated at the highest values of Γ and R. In the Gaussian case, the likelihood for low values of s becomes

$$\Pr(\boldsymbol{d}|\Gamma, R) = \frac{1}{2\pi\sigma_\Gamma\sigma_R}\mathrm{e}^{\frac{-{\Gamma_\mu}^2}{2{\sigma_\Gamma}^2}}\mathrm{e}^{\frac{-R_\mu^2}{2{\sigma_R}^2}}\left(1+\frac{\Gamma_\mu\Gamma}{2{\sigma_\Gamma}^2}+\frac{R_\mu R}{2{\sigma_R}^2}\right). \tag{2.35}$$

This means that the likelihood becomes linear and increasingly flat as $s \to 0$. The same happens to the posterior, so that the ratio in Equation 2.27 becomes a ratio of two normalized flat functions, so that qualitatively, we can expect it to approach unity. This can also be shown more rigorously, as for low values of Γ and R, Equation 2.27 reduces to

$$B_{12} = 1 - s\frac{\Gamma_\mu\Gamma_{\max}}{2{\sigma_\Gamma}^2}. \tag{2.36}$$

Figure 2.4 shows the dependence of $B_{21} = 1/B_{12}$ on s for the Gaussian case, and contrasts it with the behaviour of B_{21} for the On/Off case. The figure shows that the Bayes factors in these two cases agree at high values of s, but strongly diverge when s becomes small. B_{21} becomes extremely large for the On/Off case, and approaches unity for the Gaussian case.

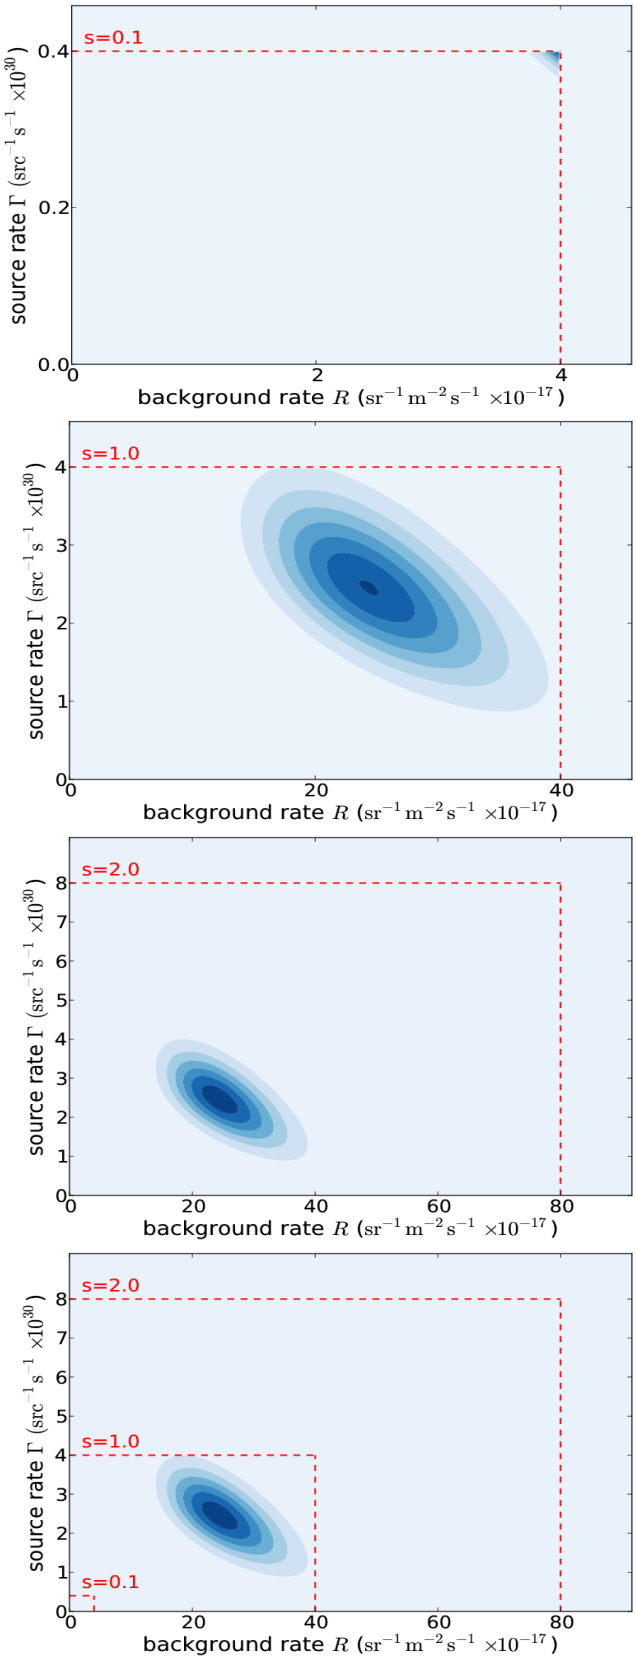


Figure 2.2: Lower panel: Example On/Off likelihood. The red lines denote prior regions for three different values of the hyperparameter s. Upper three panels: Posteriors for the same s values are displayed.

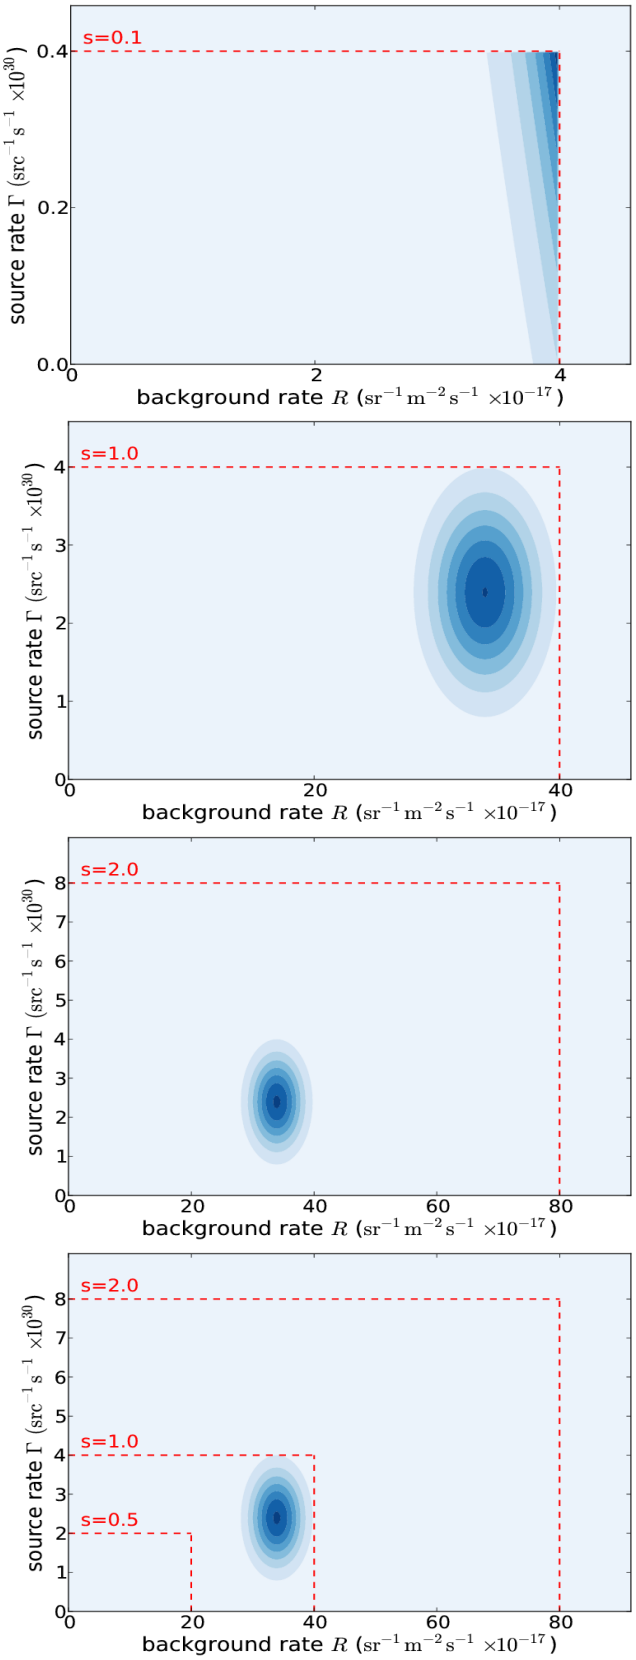


Figure 2.3: Lower panel: Example Gaussian likelihood. The red lines denote prior regions for three different values of the hyperparameter s. Upper three panels: Posteriors for the same s values are displayed.

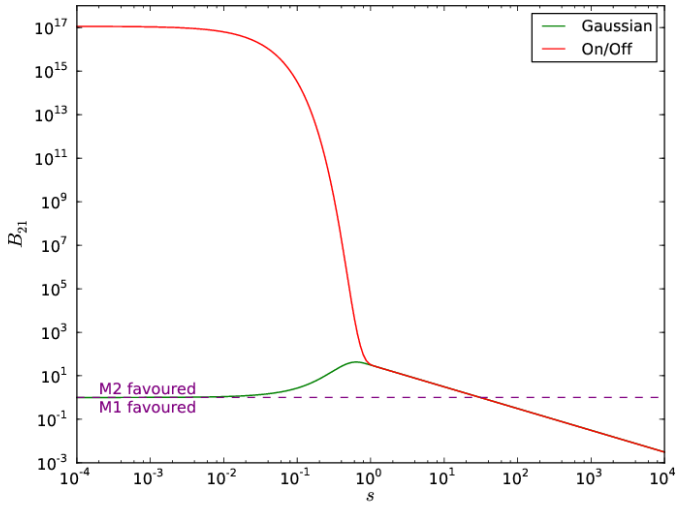


Figure 2.4: Dependence of the Bayes factor on hyperparameter s for the cases of the On/Off likelihood and the Gaussian likelihood.

2.6 Comparison with frequentist statistics

According to the frequentist school of thought, probability is defined as

"the limit of the relative frequency of an event in an infinite series of equiprobable trials."

This is fundamentally different from the Bayesian conception, as here the probability is not a description of the state of knowledge of the observer.

A classic example comparing these two conceptions of probability is that of a coin toss, which has been discussed by E. T. Jaynes among others (e.g. Jaynes 2003). The statement "The coin has a 1/2 chance of landing heads or tails." is interpreted differently by the two schools of thought:

Frequentist interpretation: *"If a coin is tossed a very large number of times, in the long run, the frequency of heads will approach 1/2."*
Bayesian interpretation: *"Based on the available information, we are completely unable to predict whether the tossed coin will land heads or tails."*

The frequentist interpretation is making a statement about the coin itself and the process involved in tossing it, while the Bayesian interpretation is only talking about the state of knowledge of the person tossing the coin.

If the Bayesian is given a coin that is known to be biased, but not by how much and in what way, then the Bayesian would give a probability to the coin coming up heads as 1/2. The frequentist, on the other hand, treats the probability as a

property of the coin itself. In the example of a biased coin of unknown bias, the frequentist would say that the probability of coming up heads can be anything *except* 1/2.

The conceptual foundations of frequentism have received considerable criticism. Jaynes argued for the Bayesian view, and labelled the association of probabilities with physical reality, rather than the mind, as the mind-projection fallacy. Jaynes argued that the outcome of a great number of coin-tosses is not solely dependent on the properties of the coin, but also on how the coin is tossed, so that a machine that controls the initial conditions of a coin toss could arbitrarily change the frequencies of heads or tails for any coin. The frequentist would need to argue that the tosses need to be "fair", but that leads to a circular argument, as it is difficult to define the "fairness" of the coin-toss except by saying that it should result in the correct frequency of heads.

Given the very different definition of probability in frequentism, the range of things to which probabilities can be assigned is greatly diminished. Frequentist methods have, for example, difficulties with dealing with single, unrepeatable events, which is a problem especially pertinent to cosmology. Most importantly, frequentists cannot assign probabilities to hypotheses. A hypothesis is either true or false, so that its probability, to the extent that it can be meaningfully discussed, is one or zero.

This is a significant problem, as the core of inductive, scientific reasoning is "Given these data, what is the probability that hypothesis H is true?" The discipline of statistics was developed to enable the use of classical frequentist probability theory for the assessment of the viability of hypotheses. It is based on the idea that data can be considered to be a random variable, while hypotheses cannot be.

Generally speaking, frequentist methodologies can be described as follows (see, e.g., Loredo (1992) for a similar exposition):
1. A procedure Π is specified for selecting a hypothesis based on some function of the data $\boldsymbol{d}$. The function is called a statistic, $S(\boldsymbol{d})$.
2. The data $\boldsymbol{d}$, and therefore statistic $S(\boldsymbol{d})$, are considered a random variable, and can therefore be assigned probabilities in the frequentist sense. The probability distribution of S given the hypothesis H, $\Pr(S|H)$, is called the sampling distribution.
3. Using the sampling distribution, it is determined what the result would be of applying Π to a large number of data samples that are predicted by H.
4. Π is applied to a real data sample.

Thus, frequentist statistics analyses hypotheses indirectly, by considering the frequencies of different data sets that might be observed if a given hypothesis is true. There are a number of flaws in frequentist methodologies that lead them to give results that are difficult to interpret, and that do not answer the questions that scientists are really interested in. Here, we review some of these problems and explain how they derive from the frequentist conception of probability.

2.6.1 Example 1: Confidence intervals

A widespread procedure that is used in frequentist parameter inference is the estimation of confidence intervals. Given a set of observations, a confidence interval is calculated, in such a way that the interval frequently includes the true value of the parameter of interest. The confidence level of the interval is the frequency with which the intervals contain the true value of the parameter in the limit of an infinite number of hypothetical samples. For example, a confidence level of 95% would mean that in the limit of an infinite number of confidence intervals calculated for samples of the kind that were observed, 95% would include the true value of the parameter.

This is very different from the Bayesian credible regions discussed in Section 2.4. The Bayesian approach fixes the 95% credible region as there being a 95% chance that the true value of the parameter is in this region. The frequentist approach fixes the parameter, and ensures that if confidence intervals are calculated for a great number of hypothetical data sets, 95% of the confidence intervals will contain the fixed parameter. This is no mere philosophical distinction, and can lead to important differences in results. Here, we consider two examples of simple physical problems, one of which leads to identical results for the Bayesian and the frequentist approaches, and another for which the results diverge significantly.

Problem 1: The mean of a Gaussian

Let us first examine a simple problem where the two approaches produce identical results: The estimation of the mean of a Gaussian. Suppose that we are observing an astrophysical object that is known to have a constant brightness. Let us assume that we have made a set of N observations of the brightness of this object, $\boldsymbol{d} = \{b_i\}_{i=1}^{N}$, and let us assume that our measurements are subject to Gaussian error with standard deviation σ_b. What can we learn about the brightness of the object from these observations?

Following the frequentist approach, we derive an unbiased estimate of the mean μ of the distribution. A statistic is said to be an unbiased estimator $\hat{\theta}$ of a parameter

θ if the mean of the sampling distribution of $\hat{\theta}$ is equal to the true value of θ:

$$\int \hat{\theta}(D)\mathrm{Pr}(\mathrm{D}|\theta_{\mathrm{true}})\mathrm{d}D = \theta_{\mathrm{true}}. \tag{2.37}$$

The unbiased estimate of the mean of the distribution is given by

$$\bar{b} = \frac{1}{N}\sum_{i=1}^{N} b_i. \tag{2.38}$$

Since we assumed Gaussian error, the sampling distribution for the mean is normal

$$\mathrm{Pr}(\bar{\mathrm{b}}|\mu) \propto \exp\Big[\frac{-(\bar{\mathrm{b}}-\mu)^2}{2\sigma_\mu^2}\Big], \tag{2.39}$$

where the standard error of the mean is

$$\sigma_\mu = \sigma_b/\sqrt{N}. \tag{2.40}$$

From this sampling distribution, we can derive a 95% confidence interval for the brightness of the object:

$$\mathrm{CI} = (\bar{\mathrm{b}} - 2\sigma_\mu, \bar{\mathrm{b}} + 2\sigma_\mu). \tag{2.41}$$

Now let us derive the corresponding credible region following the Bayesian approach. We start with a flat prior over the region of interest, and the likelihood

$$\mathrm{Pr}(\boldsymbol{d}|\mu) = \prod_{\mathrm{i}=1}^{\mathrm{N}} \frac{1}{\sqrt{2\pi\sigma_{\mathrm{b}}^2}}\exp\Big[\frac{(\mu - b_{\mathrm{i}})^2}{2\sigma_{\mathrm{b}}^2}\Big]. \tag{2.42}$$

We obtain the posterior

$$\mathrm{Pr}(\mu|\boldsymbol{d}) \propto \exp\Big[\frac{-(\mu-\bar{\mathrm{b}})^2}{2\sigma_\mu^2}\Big]. \tag{2.43}$$

We thus see that in this case, the Bayesian posterior on μ is equivalent to the frequentist sampling distribution, and the 95% credible region is equivalent to the frequentist confidence interval:

$$\mathrm{CR} = (\bar{\mathrm{b}} - 2\sigma_\mu, \bar{\mathrm{b}} + 2\sigma_\mu). \tag{2.44}$$

Problem 2: The truncated exponential

Let us now consider a case where the Bayesian and frequentist results are not equivalent, and that highlights some of the limitations of frequentist methodology. E.T. Jaynes discussed the problem of the 'truncated exponential' (Jaynes 1976). Jaynes described a device that operates without failure for a time t_0 because of a chemical inhibitor that has been injected into it. At t_0, the supply of the chemical is exhausted, and failures begin occurring in the device. The times at which the failures occur follow an exponential failure law:

$$\Pr(t|t_0) = \begin{cases} \exp(t_0 - t) & \text{if } t > t_0 \\ 0 & \text{if } t < t_0. \end{cases} \tag{2.45}$$

The shape of the function is illustrated in Figure 2.5.

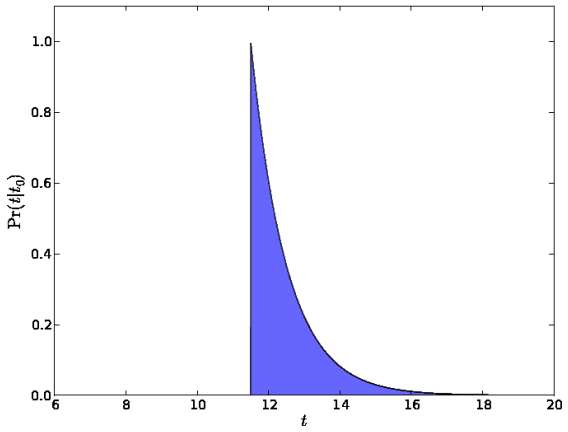


Figure 2.5: The truncated exponential.

The task is to use data on failure times to estimate t_0, the time of safe operation of the device[1].

Let us suppose that the failure times are given by $t = 12, 14, 16$, and let us attempt to derive a confidence interval for t_0.

We first note that the population mean is given by

$$\langle t \rangle = \int_0^\infty t \Pr(t) \mathrm{d}t = t_0 + 1. \tag{2.46}$$

[1] An astrophysical problem of the same form would be the observation of neutrinos emitted by a supernova. The emission of the neutrinos can be modelled by thermal emission from a sphere with exponentially decaying temperature. The arrival times of the neutrinos on Earth have the role of the measured failure times, while the time of the supernova is equivalent to the time when the inhibitor runs out. The problem is discussed in this context in Loredo (1992).

Thus, we can write down an unbiased estimator for t_0,

$$\hat{t_0} = \frac{1}{N}\sum_{i=1}^{N} t_i - 1. \tag{2.47}$$

The sampling distribution for $\hat{t_0}$ can be shown to be (Jaynes 1976):

$$\Pr(\hat{t_0}|t_0) = (\hat{t_0} - t_0 + 1)^{N-1}\exp[-N(\hat{t_0} - t_0 + 1)]. \tag{2.48}$$

From these expressions, and with the failure times given, we would estimate that $t_0 = 13$, and we derive the shortest 90% confidence interval for t_0 as being $12.15 < t_0 < 13.83$. This is absurd: The earliest event was observed at $t = 12$, implying an upper limit on t_0, and yet the entire interval is above that value.

Let us now conduct a Bayesian analysis of the same problem. We use a uniform prior, and calculate the likelihood as a product of N truncated exponentials:

$$\Pr(\{t_i\}|t_0, I) = \begin{cases} \exp(Nt_0)\exp(-\sum_{i=1}^{N} t_i) & \text{if } t_0 < t_1 \\ 0 & \text{if } t_0 > t_1. \end{cases} \tag{2.49}$$

We obtain the posterior

$$\Pr(t_0|\{t_i\}, I) = \begin{cases} N\exp[N(t_0 - t_1)] & \text{if } t_0 < t_1 \\ 0 & \text{if } t_0 > t_1. \end{cases} \tag{2.50}$$

The 90% credible interval is $11.23 < t_0 < 12$, which agrees with our intuitive sense.

So how can we explain the discrepancy between the Bayesian and frequentist results, and the seemingly absurd nature of the frequentist result?

Let us examine again the definitions of confidence intervals and credible regions:

90% credible region *"Given the data at hand, there is a 90% probability that the true value of the parameter lies within the credible region"*
90% confidence interval: *"If we compute confidence intervals from many hypothetical samples of data of this sort, in the long run 90% of these intervals will contain the true value of the parameter."*

Both of the results that we have obtained for the problem are correct answers, but they are answering very different questions. That the frequentist result seems absurd is because it gives us the correct answer to the wrong question. What the

scientist truly cares about is what can be learned about the hypothesis from the three observations at hand. The hypothetical space of "data of this sort" is not the primary concern. The frequentist approach can put no constraint on the value of the parameter given the observed data. It aims for good long-term behaviour averaged over many hypothetical data sets, but provides a nonsensical answer in the individual case at hand.

2.6.2 Example 2: Significance tests

One of the most often used procedures in frequentist statistics is significance testing. In significance testing, there is a default "null hypothesis" H_0, and the scientist is trying to determine whether or not it should be rejected. The scientist chooses a test-statistic S under null hypothesis H_0, and then for data $\boldsymbol{d}$ determines the probability of observing a value $S(\boldsymbol{d})$, or a value more extreme than $S(\boldsymbol{d})$. This probability is called the p-value. If the p-value falls below a particular threshold, normally taken as 5%, the null hypothesis is rejected. This threshold is called the significance level, and upper-bounds the probability that H_0 is rejected when it is correct.

This methodology has been very controversial. The core of the controversy centers around the fact that the hypothesis test rejects null hypothesis H_0 based on the fact that $Pr(\boldsymbol{d}|H_0)$ is very small. This is, however a very different statement from the statement that $Pr(H_0|\boldsymbol{d})$ is small. The p-value is very often confused with a probabilistic statement about the hypothesis.

For example, if H_0 has been rejected at a significance level of 5%, it is extremely tempting to interpret this as saying that H_0 has a $< 5\%$ probability of being correct, as the latter is what scientists really are interested in. Such an interpretation of the p-value is, however, wrong. A frequentist cannot make a probabilistic statement about a hypothesis at all.

This confusion can be regarded as a misapplication of deductive reasoning. Consider the following syllogism:

1. If hypothesis H_0 is true, data $\boldsymbol{d}$ cannot be observed.
2. Data $\boldsymbol{d}$ have been observed.
3. Therefore, H_0 is false.

This is a valid deductive argument. Now, let us examine the following argument:

1. If hypothesis H_0 is true, data $\boldsymbol{d}$ are highly unlikely to be observed.
2. Data $\boldsymbol{d}$ have been observed
3. Therefore, hypothesis H_0 is highly unlikely to be true.

This syllogism is not correct. For example, consider the two hypotheses shown in Table 2.1:

Table 2.1: Two example hypotheses.

Hypothesis	X = 1	X = 2	Pr(H)
H_0	0.010	0.990	0.3
H_1	0.011	0.989	0.7

If we observe X=1, we would reject H_0 at a 0.01 significance level. However, if we do a Bayesian calculation of $\Pr(H_0|\boldsymbol{d})$, we obtain

$$\Pr(H_0|\mathrm{X}=1) = \frac{\Pr(\mathrm{X}=1|H_0)\Pr(H_0)}{\Pr(H_0)\Pr(\mathrm{X}=1|H_0) + \Pr(H_1)\Pr(\mathrm{X}=1|H_1)}. \tag{2.51}$$

$$\Pr(H_0|\mathrm{X}=1) = \frac{0.3 \times 0.01}{0.3 \times 0.01 + 0.7 \times 0.011} \approx 0.28. \tag{2.52}$$

So that the probability of the hypothesis being false in fact is 0.28, not 0.01. The p-value does not necessarily tell the scientist anything about the question whether or not the null hypothesis is a correct description of the data.

P-values are not unique, and change due to arbitrary factors like the test statistic that is used. We can illustrate the non-uniqueness of p-values using a simple example, taken from Murphy (2012): imagine that data $\boldsymbol{d}$ constitutes 1000 coin tosses, 474 of which are tails. If frequentist methods are used to determine whether or not the coin is biased, the result depends on what experiment was carried out, and what other data sets might have been observed.

One possibility is that the coin was tossed 1000 times, and 474 tails were observed. In that case, we obtain the p-value

$$\sum_{k=0}^{474} \binom{1000}{k} \left(\frac{1}{2}\right)^{1000} = 0.05337. \tag{2.53}$$

So at a significance level of 0.05 we do not reject the null hypothesis of an unbiased

coin.

Another possibility is that the coin was tossed until 474 heads were obtained, and this process took 1000 tosses. In this case, we obtain the p-value

$$\sum_{n=1000}^{\infty} \binom{n-1}{473} \left(\frac{1}{2}\right)^n = 0.04994. \tag{2.54}$$

and reject the null hypothesis of an unbiased coin at a 0.05 significance level.

Thus, two different results are obtained depending on the "stopping rule" governing the observations, which lead to different ideas about what data might have been observed.

There are many more examples of this kind. These examples illustrate that the differences between the Bayesian and the frequentist conceptions of probability are not merely philosophical. While in many situations both methodologies will lead to similar results, there are situations where they will greatly diverge, and where the frequentist methods give results that are absurd. As frequentist methods cannot assign probabilities to hypotheses, they fail to answer the questions that scientists are truly interested in.

Chapter 3

A Bayesian self-clustering analysis

In this chapter, we discuss a fully Bayesian approach to analyzing the self-clustering of points on the sphere, which we apply to a sample of 69 UHECRs detected by the Pierre Auger Observatory. The analysis is based on a multi-step approach that enables the application of Bayesian model comparison to cases with weak prior information. Our results are consistent with both a uniform and clustered model of UHECRs. Data-sets of far greater magnitude are expected to be produced by future experiments such as the Japanese Experiment Module Extreme Universe Space Observatory (JEM-EUSO), which will enable this method to distinguish between the two models.

3.1 Introduction

As discussed in Chapter 1, one of the approaches to studying UHECR arrival directions is the search for catalogue-independent clustering of the directions. This approach is advantageous in cases of large magnetic deflections, transient sources and incomplete catalogues. A number of studies have previously investigated this question (e.g. Abreu et al. 2012), but no consensus has been reached. The application of Bayesian methods to this question faces a difficulty: We would like to conduct a Bayesian model comparison between a model of uniform UHECRs and a model in which the UHECRs are non-uniform and clustered. However, there exists an infinite variety of possible non-uniform models that might explain the distribution of UHECR arrival directions. This is a significant conceptual problem: it is difficult to decide which alternative clustered model should be used. This is in many ways similar to the cases of poorly specified prior parameters discussed in Chapter 2. The difference here is that not only does this situation lack well-motivated parameter priors, but the model itself is initially unspecified. In this chapter, we describe a Bayesian method that we have developed to study this question (Section 3.2). This method is similar to the Bayesian methods used in cases of poorly specified prior parameters discussed in Section 2.5, in that it splits the data sample into several sub-samples, some of which are used as training samples from which the non-uniform model is derived, so that the rest of the data can then be used in a Bayesian model comparison. We describe the application of this method to mock data sets (Section 3.2), and finally we discuss the application of the method to data from the PAO (Section 3.4). Our conclusions are summarized in Section 3.5.

3.2 Statistical formalism

Our primary aim here is to assess whether there is evidence that the distribution of UHECR arrival directions is anisotropic. We do this by using Bayesian inference in the context of two models: a uniform model, M_{u}, which would be the null hypothesis in a classical hypothesis test; and a non-uniform model, M_{n}, as yet unspecified. The posterior probability of the non-uniform model, conditional on data in the form of N UHECR arrival directions $\{\boldsymbol{r}_i\}$ (where $i \in \{1, 2, \ldots, N\}$), is given by Bayes's theorem as

$$\begin{aligned} &\Pr(M_{\mathrm{n}}|\{\boldsymbol{r}_i\}) \\ &= \frac{\Pr(M_{\mathrm{n}})\Pr(\{\boldsymbol{r}_i\}|M_{\mathrm{n}})}{\Pr(M_{\mathrm{u}})\Pr(\{\boldsymbol{r}_i\}|M_{\mathrm{u}}) + \Pr(M_{\mathrm{n}})\Pr(\{\boldsymbol{r}_i\}|M_{\mathrm{n}})}, \end{aligned} \tag{3.1}$$

where $\Pr(M_\mathrm{u})$ and $\Pr(M_\mathrm{n})$ are the prior probabilities of the two models, and $\Pr(\{\boldsymbol{r}_i\}|M_\mathrm{u})$ and $\Pr(\{\boldsymbol{r}_i\}|M_\mathrm{n})$ are the probabilities of the data under each of the models (i.e. the likelihoods). The ratio of the posterior probabilities can be written as

$$\frac{\Pr(M_\mathrm{n}|\{\boldsymbol{r}_i\})}{\Pr(M_\mathrm{u}|\{\boldsymbol{r}_i\})} = \frac{\Pr(M_\mathrm{n})}{\Pr(M_\mathrm{u})} B, \tag{3.2}$$

where

$$B = \frac{\Pr(\{\boldsymbol{r}_i\}|M_\mathrm{n})}{\Pr(\{\boldsymbol{r}_i\}|M_\mathrm{u})} \tag{3.3}$$

is the Bayes factor. In the convention adopted here, models M_u and M_n are favoured by small and large values of B, respectively.

The next task is to specify the two models to be compared and to evaluate the marginal likelihoods for both. The null hypothesis represented by the uniform model (Section 3.2.1) is unambiguous and yields the marginal likelihood given in Equation 3.4; the alternative non-uniform model (Section 3.2.2) is more complicated and is derived from a subset of the data, eventually yielding the marginal likelihood given in Equation 3.8. This requirement means that both marginal likelihoods are evaluated only for the remaining data that was not used to obtain the non-uniform model.

3.2.1 Uniform model

In the uniform model, M_u, the probability that a UHECR arrives from direction $\boldsymbol{r}$ is constant at $\Pr(\boldsymbol{r}|M_\mathrm{u}) = 1/(4\pi)$. Hence, the marginal likelihood for a test sample of N_t UHECRs with arrival direction $\{\boldsymbol{r}_t\}$ (with $t \in \{1, 2, \ldots, N_\mathrm{t}\}$) is given by

$$\Pr(\{\boldsymbol{r}_t\}|M_\mathrm{u}) = \frac{1}{(4\pi)^{N_\mathrm{t}}}. \tag{3.4}$$

This simple expression is, however, valid only in the case of uniform exposure; if the exposure is non-uniform, as is always the case for real experiments, it must be modified as described in Section 3.2.3.

Figure 3.1: Full process of model creation for data sets of 69 UHECRs for three test cases: uniform arrival directions (left); three sources (middle); and AGN sources from a realistic mock catalogue (right). Three aspects of the analysis procedure are shown: (1) the full input UHECR data set; (2) partition of the data set into • generating points, • fitting points and ○ testing points; and (3) the resultant mixture distribution of vMF kernels centred on the generating points. The model is created for a range of κ values, but for each of the three test cases, only the maximum likelihood value of κ is displayed here. These highest likelihood values are $\kappa = 0$, 108 and 13 for the uniform sources, three sources, and AGN sources, respectively.

3.2.2 Non-uniform model

To derive the non-uniform model, we develop a multi-stage Bayesian approach by splitting the arrival directions $\{\boldsymbol{r}_i\}$ into three subsets:

(i) First, N_g generating points $\{\boldsymbol{r}_g\}$ are chosen as the centres of smooth, localized kernels which can be combined into a mixture distribution on the sphere (Section 3.2.2.1).

(ii) Then, N_f fitting points $\{\boldsymbol{r}_f\}$ (with $f \in \{1, 2, \ldots, N_\mathrm{f}\}$) are used to obtain a distribution for the unspecified width parameter of the kernels (Section 3.2.2.2).

(iii) Finally, the remaining N_t testing points $\{\boldsymbol{r}_t\}$ (with $t \in \{1, 2, \ldots, N_\mathrm{t}\}$) are used to evaluate the marginal likelihood under this non-uniform model (Section 3.2.2.3).

The partitions of the data are chosen at random and the generating points are not linked to the putative UHECR sources in any way. This method is hence independent of any source catalogue or propagation model and, indeed, could be applied to any sample of points on the sphere. The three steps of this approach are illustrated in Figure 3.1 for the three test cases described in Section 3.2.4.

The resultant model (and marginal likelihood) is fully specified, but the algorithm for generating it has two free parameters: N_g and N_f. While the relative merits of using a low or high fraction of the data to generate and fit the model are an open question, here we take the simplest approach by using a third of the data at each step, so $N_\mathrm{f} = N_\mathrm{g} = \mathrm{floor}(N/3)$, leaving $N_\mathrm{t} = N - (N_\mathrm{f} + N_\mathrm{g}) \simeq N/3$ testing points. The results of varying these divisions are deferred to future work.

3.2.2.1 Generating a clustered model from the data

The first step to specifying a non-uniform model is to use the N_g generating points $\{\boldsymbol{r}_g\}$ as the centres of smooth, localized kernels of an as yet unspecified angular size.

The specific kernel chosen was the von Mises Fisher (vMF) distribution, which resembles a Gaussian on the sphere and is defined by the density

$$\Pr(\boldsymbol{r}|\overline{\boldsymbol{r}}, \kappa) = \frac{\kappa}{4\pi \sinh(\kappa)} e^{\kappa \boldsymbol{r}\cdot\overline{\boldsymbol{r}}}, \tag{3.5}$$

where $\overline{\boldsymbol{r}}$ is the central direction and κ is the concentration parameter. This is inversely related to the width of the distribution: for large values of κ the distribution is peaked over an angular scale of $\sim 1/\sqrt{\kappa}$, while if κ tends to 0 the distribution becomes uniform on the sphere. The vMF distributions were centred on the generating

points to give the mixture model density

$$\Pr(\boldsymbol{r}|\{\boldsymbol{r}_g\},\kappa) = \frac{\kappa}{4\pi \mathrm{N_g}\sinh(\kappa)} \sum_{g=1}^{N_\mathrm{g}} e^{\kappa \boldsymbol{r}\cdot\boldsymbol{r}_g}. \tag{3.6}$$

3.2.2.2 Obtaining a concentration distribution

The last step to fully defining the non-uniform model is to specify a distribution for κ. This is done by using the fitting points to obtain a fully normalized posterior for κ that can be used as a parameter prior in the model comparison step. A uniform prior for $\kappa \geq 0$ is chosen in order to include models with $\kappa = 0$ (which would not be possible for, e.g. a logarithmic prior in κ). The posterior distribution that results from generating points $\{\boldsymbol{r}_g\}$ and fitting points $\{\boldsymbol{r}_f\}$ is

$$\begin{aligned}
\Pr(\kappa|\{\boldsymbol{r}_g\},\{\boldsymbol{r}_f\}) &= \frac{\Pr(\kappa)\Pr(\{\boldsymbol{r}_f\}\,|\,\{\boldsymbol{r}_g\},\kappa)}{\int_0^\infty \Pr(\kappa')\Pr(\{\boldsymbol{r}_f\}\,|\,\{\boldsymbol{r}_g\},\kappa')\,\mathrm{d}\kappa'} \\
&\propto \Theta(\kappa)\prod_{f=1}^{N_\mathrm{f}} \Pr(\boldsymbol{r}_f|\{\boldsymbol{r}_g\},\kappa) \\
&\propto \frac{\Theta(\kappa)\,\kappa^{N_\mathrm{f}}}{\sinh^{N_\mathrm{f}}(\kappa)} \prod_{f=1}^{N_\mathrm{f}} \left(\sum_{g=1}^{N_\mathrm{g}} e^{\kappa \boldsymbol{r}_f\cdot\boldsymbol{r}_g} \right),
\end{aligned} \tag{3.7}$$

where $\Theta(\kappa)$ is the Heaviside step function that encodes the fact that κ is non-negative. The posterior distribution is straightforward to normalize numerically as it is (generally) unimodal and as there is only one parameter.

The alternative, non-uniform model for the UHECR arrival directions is hence fully specified (in the sense of being usable in Bayesian model comparison). It is a sum of vMF distributions centred on the set of generating points, $\{\boldsymbol{r}_g\}$, and with the distribution of vMF concentration parameter κ given by Equation 3.7.

3.2.2.3 Evaluating the marginal likelihood

Having specified the non-uniform model, M_n, with the generating points, $\{\boldsymbol{r}_g\}$ and obtained the distribution $\Pr(\kappa|M_\mathrm{n})$ by using the fitting points, $\{\boldsymbol{r}_f\}$, it is now possible to use the remaining data, the testing points $\{\boldsymbol{r}_t\}$, to evaluate the marginal likelihood. From Equation 2.20 this is

$$\Pr(\{\boldsymbol{r}_t\}\,|M_\mathrm{n}) = \int_0^\infty \Pr(\kappa|M_\mathrm{n})\Pr(\{\boldsymbol{r}_t\}\,|\kappa,M_\mathrm{n})\,\mathrm{d}\kappa, \tag{3.8}$$

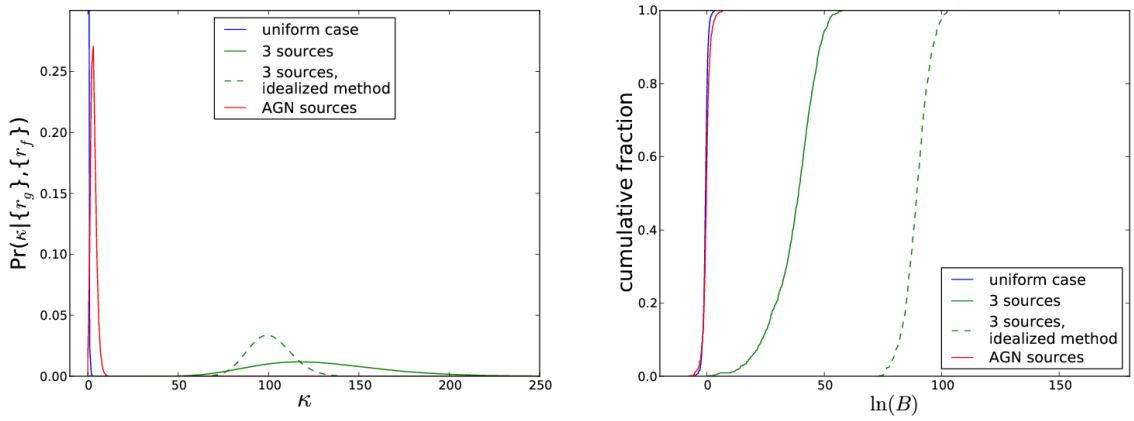


Figure 3.2: Left panel: κ posteriors. Right panel: cumulative fractions of Bayes factors, produced by the application of the multi-step method to test cases of 69 UHECR events. Three test cases are considered: uniform UHECRs; UHECRs generated from three sources; and UHECRs generated by AGNs from a realistic catalogue. For each of the test cases, the cumulative fractions are based on a total of 1,000 Bayes factors. In the case of three sources, in addition to the conventional application of the multi-step method, the results for an idealized method are displayed. In the idealized application of the method, the generating points are taken as the true centres of the vMF distributions that generate the UHECRs, rather than as a random subset of the data.

where $\Pr(\kappa|M_{\rm n}) = \Pr(\kappa|\{\boldsymbol{r}_g\},\{\boldsymbol{r}_f\})$ is given in Equation 3.7 and now plays the role of the prior distribution for κ, and the likelihood for the testing points is (cf. Equation 3.6)

$$\begin{aligned}\Pr(\{\boldsymbol{r}_t\}|\kappa,M_{\rm n}) &= \Pr(\{\boldsymbol{r}_t\}|\{\boldsymbol{r}_g\},\kappa)\\ &= \prod_{t=1}^{N_{\rm t}} \Pr(\boldsymbol{r}_t|\{\boldsymbol{r}_g\},\kappa)\\ &= \frac{\kappa^{N_{\rm t}}}{[4\pi N_{\rm g}\sinh(\kappa)]^{N_{\rm t}}} \prod_{t=1}^{N_{\rm t}} \left(\sum_{g=1}^{N_{\rm g}} e^{\kappa \boldsymbol{r}_t\cdot\boldsymbol{r}_g}\right). \end{aligned} \tag{3.9}$$

The one-dimensional integral in Equation 3.8 is, once again, straightforward to evaluate numerically. This then gets further modified by the non-uniform exposure, as described in Section 3.2.3.

3.2.3 Non-uniform exposure

When studying the measured arrival directions of CRs in a real experiment, the non-uniform exposure of the observatory (see Section 1.4.1) needs to be taken into

account. The distribution of arrival directions of detected CRs is then given by Bayes's theorem as

$$\Pr(\boldsymbol{r}|\mathrm{det}) \propto \Pr(\boldsymbol{r}) \Pr(\mathrm{det}|\boldsymbol{r}) \propto \Pr(\boldsymbol{r}) \frac{\mathrm{d}\epsilon}{\mathrm{d}\Omega}, \tag{3.10}$$

where $\Pr(\boldsymbol{r})$ is the distribution of arrival directions of all CRs, irrespective of whether they are detected.

For uniform UHECR arrival directions discussed in Section 3.2.1, $\Pr(\boldsymbol{r}|M_{\mathrm{u}}) = 1/(4\pi)$, so that $\Pr(\boldsymbol{r}, \mathrm{det}|E)$ simply becomes

$$\Pr(\boldsymbol{r}|\mathrm{det}, M_{\mathrm{u}}) = \frac{1}{\epsilon_{\mathrm{tot}}} \frac{\mathrm{d}\epsilon}{\mathrm{d}\Omega}. \tag{3.11}$$

For the non-uniform UHECR arrival directions discussed in Section 3.2.2, $\Pr(\boldsymbol{r}|\kappa, M_{\mathrm{n}})$ is given in Equation 3.6, so that

$$\Pr(\boldsymbol{r}|\mathrm{det}, \kappa, M_{\mathrm{n}}) \propto \frac{\mathrm{d}\epsilon}{\mathrm{d}\Omega} \sum_{g=1}^{N_{\mathrm{g}}} e^{\kappa \boldsymbol{r} \cdot \boldsymbol{r}_g}, \tag{3.12}$$

where the normalization depends on the position of the generating points, $\{\boldsymbol{r}_g\}$, the relative exposure and κ, and must be calculated numerically.

3.2.4 Illustration of the multi-step method

Figure 3.1 illustrates how the multi-step Bayesian method works for several simple test cases: a uniform source distribution; a model with three sources; and a model based on the AGN simulations described below in Section 3.3.2. The total number of UHECRs is 69 in all cases. The associated κ posteriors and the resultant distributions of Bayes factors are shown in Figure 3.2. The cumulative fractions of Bayes factors that are shown in Figure 3.2 are each based on a total of 1,000 Bayes factors.

The first test case was a very simple scenario: the UHECRs were simulated with isotropic arrival directions, for the case of uniform exposure. The κ posterior for the uniform case has its maximum very close to 0, and declines rapidly, because the vMF distributions that are fitted to the data are almost uniform. The Bayes factors for this case are small: the uniform model is favoured in 74.1% of the simulations.

The second test case is a simple model of non-uniform arrival directions: the UHECRs were sampled from three vMF distributions, representing three UHECR sources. The concentration parameter κ of the vMF distributions was taken as 90. The κ posterior for this case is systematically peaked at higher values, as can be

seen in Figure 3.2. It is peaked at a value higher than the input value of κ because each of the three original kernels is now accounted for by multiple narrower kernels that are slightly off-centre. The Bayes factors are very large: the non-uniform model is favoured in more than 99.9% of the simulations and the average Bayes factor is ~ 40.

For the case of three sources, it was also possible to apply an idealized form of the multi-step method: instead of using one third of the full data set as the generating points, the generating points were taken as the actual positions of the sources of the UHECRs. In this way, the idealized method does not share the catalogue-independence of the full three-step method described in Section 3.2.2. For this idealized case, the κ posterior is consistent with the input value, because the three original kernels are accounted for by three kernels located on the original kernel positions. This is also the reason why the Bayes factors are even larger than for the ordinary case. The idealized form of the multi-step method is useful to see the potentially strong impact the lack of knowledge about the source positions can have, although it hence cannot be used to analyze real data.

The third test case was the case for UHECRs generated by AGNs, simulated with the realistic model described in Section 3.3.2. The input value of $\kappa = 360$ was chosen to give the strongest plausible signal, but the resultant posterior is peaked close to $\kappa = 0$. The reason is that there are now so many sources compared to the number of UHECRs that the source distribution is undersampled. This is an indication that, given the weak (projected) clustering expected of nearby AGNs, a significantly larger UHECR sample would be needed for their self-clustering to be apparent. More realistic tests that are documented in Section 3.3 confirm this result.

3.3 Application to simulated UHECR samples

To investigate the effectiveness of the multi-stage Bayesian method described above, it was applied to realistic mock catalogues of UHECRs. Catalogues were created for two different UHECR scenarios: isotropic (Section 3.3.1) and AGN centred (Section 3.3.2). The samples of incoming UHECRs were then subjected to the PAO measurement process (Section 3.3.3). The distributions of Bayes factors for the resultant observed samples are analysed in Section 3.3.4.

3.3.1 Isotropic distribution of sources

The application of the multi-step method to uniform UHECR distributions acted as a false positive test. Computing large numbers of Bayes factors for uniform UHECR distributions can be used to establish how often the null hypothesis is wrongly rejected.

3.3.2 AGN sources

Simulated UHECR catalogues were created for the case of UHECRs originating in AGNs. To do so, we developed a model of UHECR injection, propagation and detection. Various forms of this model were used in our work in a number of contexts, so it is worth spelling it out in some detail. The model consisted of three components:

(i) a realistic model of the injection of the UHECRs at a source catalogue (Section 3.3.2.1);

(ii) a realistic model of UHECR propagation to Earth (Section 3.3.2.2);

(iii) a realistic model of UHECR detection (Section 3.3.3).

Mock catalogues of cosmic rays are available online, such as the Carmen and Consuelo mock catalogues (Center for Cosmology and Particle Physics 2013). However, we decided to create a program that produces custom UHECR mock catalogues, as that gave us control over the input parameters.

It should be noted that despite the fact that the arrival energy is not directly used in the method, it is important. This is because only UHECRs above a certain energy are measured, so that sources at different distances will contribute different numbers of events.

3.3.2.1 Injection at the sources

In order to produce a realistic UHECR catalogue, the galaxy catalogue that is used must be volume-limited to a greater distance than is available in current all-sky galaxy surveys. Therefore, a mock galaxy catalogue was used for that purpose. The AGN sources were drawn randomly from the simulated Las Damas "Consuelo" catalogues[1], following a similar procedure to Berlind et al. (2011). Galaxies were selected with a probability $p_i = w_i / \sum_j w_j$, where $w_i = r_i^{-2}$ and r_i is the distance to the ith galaxy, to account for flux conservation.

[1] http://lss.phy.vanderbilt.edu/lasdamas/

Two source densities were used: $10^{-3.5}\,\mathrm{Mpc}^{-3}$ and $10^{-4.5}\,\mathrm{Mpc}^{-3}$. These are the highest and lowest source densities available in the Consuelo catalogues, and represent a reasonable range of possible source densities.

The injection spectrum of the UHECRs at the sources was assumed to be a power-law of the form $Q(E) \propto E^{-\alpha}$, where $Q(E)\,\mathrm{d}E$ is the number of cosmic rays emitted with energy between E and $E + \mathrm{d}E$ per unit time, and α is the power law index. Simulations were conducted for three values of the index: 2.0, 2.3 and 2.7, spanning the range of values used in e.g. De Domenico & Insolia (2013), Abreu et al. (2013), Ahlers & Salvado (2011) and Decerprit & Allard (2011).

3.3.2.2 Propagation model

To model the propagation of protonic UHECRs, two effects need to be taken into account: The angular deflection of the UHECR arrival directions due to the extra-Galactic and Galactic magnetic fields, and the energy loss of the travelling UHECRs (see Section 1.3).

We combined the magnetic deflection that a UHECR experiences during propagation and the uncertainty in its detected arrival direction into a single kernel, which was chosen to be a vMF distribution. We refer to this as the buckshot model. We assumed a fiducial smearing angle of $\sigma \simeq 3$ deg ($\kappa = 360$), and also conducted investigations for smearing angles of $\sigma \simeq 6$ and 10 deg ($\kappa = 90$ and 30).

A limitation of the simple buckshot model is the fact that it does not account for the direction-dependence of the UHECR deflection in structured magnetic fields. Models that take into account such direction dependence would yield correlated UHECR arrival directions. An example of an extension of the buckshot model that takes into account the direction dependence of the deflection has been discussed in Soiaporn et al. (2013), where the vMF distribution of a given source is centred not on the source itself, but on a particular guide direction that is associated with that source. Other studies have discussed and conducted simulations with more refined models of the magnetic fields that account for the direction dependence (e.g. Medina Tanco et al. 1998; Nagar & Matulich 2010; Farrar et al. 2013; Keivani et al. 2015). A possible extension of our work would be to incorporate such models into our framework.

The energy loss during propagation can be modelled by taking into account the processes described in Section 1.3.1. When this research was conducted, the energy loss due to scattering off the infrared background light had not yet been incorporated into the framework, so that the loss processes that were considered were the two

scattering processes off the CMB radiation, and one continuous, adiabatic energy loss due to the expansion of the Universe[2].

The BH scattering process has a very short mean free path over the entire energy range, so that the stochastic nature of this process barely has any impact. Thus, the BH energy loss was regarded as continuous, like the adiabatic energy loss. The GZK effect, on the other hand, has very large mean free paths at lower energies, so that the stochastic nature of the process needed to be directly modelled.

For each process individually, it is clear how the given parameters could be used to construct a model of the energy loss. For the continuous processes, the total amount of energy that a UHECR loses as it travels to the Earth from a given distance is determined by solving the differential equation

$$\frac{\mathrm{d}E}{\mathrm{d}r} = -\frac{E}{L_{\mathrm{loss}}(E)}, \tag{3.13}$$

where L_{loss} is the loss length (see Equation 1.4). The stochastic GZK process can also easily be modelled as the mean free path and the mean energy loss during a collision are known.

While modelling each of the processes individually was straightforward, it was more challenging to combine them into a single model of energy loss. Specifically, the question was how to combine the stochastic nature of the photoproduction with the other two continuous losses. A realistic model needs to take into account the energy losses due the individual processes, but also account for the way in which the energy losses of the separate processes affect one another during propagation.

For example, it was found during the research that if all three processes are operating at once, rather than separately, then the total amount of energy lost is actually less. This is because the GZK process, which accounts for most of the energy loss, is weaker for lower energies. As the BH and adiabatic energy losses reduce the UHECR's energy, the UHECR enters into domains of "weak GZK" sooner, so that in the end less energy is lost overall.

The following approach has been used to combine the processes: The UHECR is moved through small, incremental steps, during each of which the stochastic process has an opportunity to occur, and the continuous processes subtract a fraction of the energy. If the steps are sufficiently small, this becomes equivalent to a realistic, simultaneous operation of the processes. A "small" step hereby is small compared

[2]Subsequent tests have shown that the inclusion of the scattering off the infrared background light does not considerably alter the results.

to the mean free path of the stochastic process. Thus, the full method consisted of the following steps:

(i) Start particle at a certain energy E0 and source distance L0.

(ii) Draw a random number from a uniform distribution between 0 and λ_{GZK}. This is required to determine whether or not a GZK scattering event occurs.

(iii) Take the incremental distance step dL to be $0.1 * \lambda_{\mathrm{GZK}}$. In the last step, when the remaining distance is less than this value, use the remaining distance as the value of dL.

(iv) Calculate the amount of energy that is lost due to the processes that are approximated as being continuous (the BH scattering and the adiabatic loss) as the ray travels through dL. Calculate the amount of energy lost through the GZK process only if the random number that was drawn in the second step is lower than dL. This accounts for the stochastic nature of the GZK scattering process.

(v) Update the energy by combining the energy losses of the previous step, update the distance by subtracting dL.

(vi) Iterate while $\mathrm{L} > 0$.

The full process is displayed on the flowchart in Figure 3.3.

3.3.3 Measurement

All of the simulations were done for a PAO-like experiment, three aspects of which were modelled explicitly:

(i) PAO's non-uniform exposure was taken into account by accepting arriving UHECRs with a probability proportional to the relative exposure $\mathrm{d}\epsilon/\mathrm{d}\Omega$ defined in Section 3.2.3.

(ii) The error in PAO's energy measurement is about 12% (see Section 3.2.3), and was included in the model. This is significant as only UHECRs that have an observed energy above a fixed threshold are included in the simulated samples.

(iii) We combined the magnetic deflection of the UHECRs (Section 1.3.2) with the angular uncertainty of the PAO (Section 1.4.1) into a single kernel, as discussed in Section 3.3.2.2.

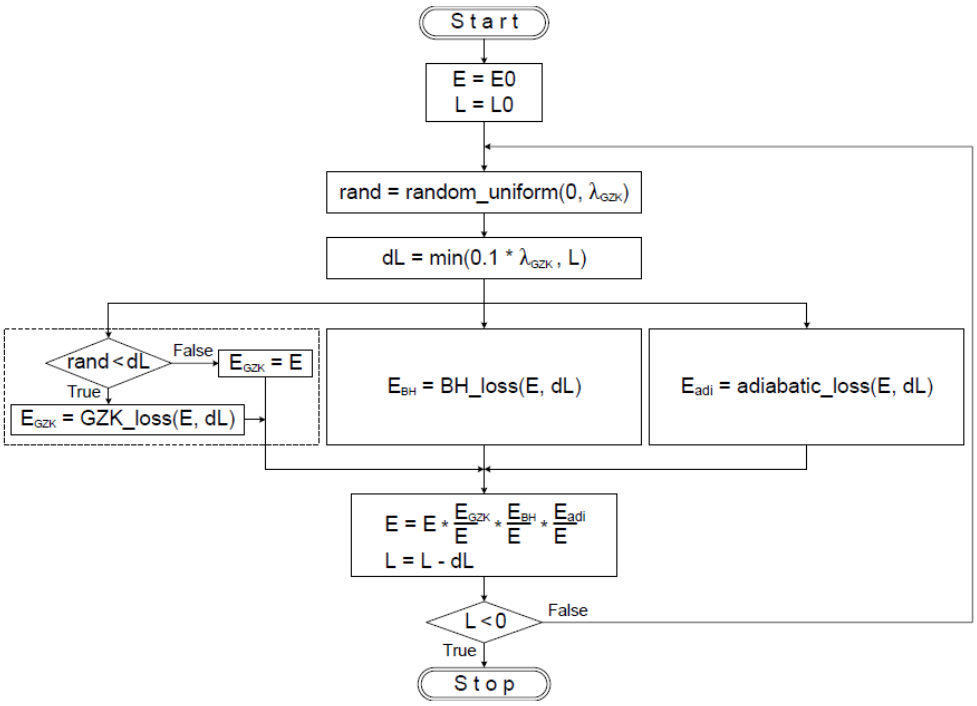


Figure 3.3: Flowchart visualising the propagation model. The flowchart shows how the three energy loss processes, two of which are continuous and one of which is stochastic, are integrated into a single propagation model.

3.3.4 Results of the simulations

Simulations were performed and Bayes factors evaluated for the isotropic model, and for the AGN-centred model with 18 combinations of the above parameters:

(i) source densities of $10^{-3.5}\,\mathrm{Mpc}^{-3}$ and $10^{-4.5}\,\mathrm{Mpc}^{-3}$;

(ii) injection parameters α of 2.0, 2.3, 2.7;

(iii) concentration parameters κ of 30, 90, 360.

For each of the 18 combinations of parameters, 1,000 samples of 69 UHECRs were created (matching the size of the PAO data set). For each sample, Bayes factors were computed for each of three energy thresholds: $5.5 \times 10^{19}\,\mathrm{eV}$, $8.0 \times 10^{19}\,\mathrm{eV}$, and $10 \times 10^{19}\,\mathrm{eV}$. Including the 1,000 realizations of the isotropic model, 55,000 Bayes factors were computed in total.

The results of these simulations are shown as cumulative distributions of Bayes factors in the top half of Figure 3.4. These are compared to similar cumulative distributions for the case of uniformly distributed UHECRs.

The Bayes factors tend to be larger for the source-centred case than for the uniform case. The difference between the results for uniform and non-uniform UHECRs

is greater for the case of low source density, as for higher source density the UHECR distribution would eventually tend to a uniform distribution.

Furthest away from the uniform case is the model with the lowest source density, highest κ and highest α. Higher κ means that the UHECR arrival directions are more closely correlated with the positions of the sources. High α reduces the GZK horizon, meaning fewer contributing AGN sources and hence more non-uniformity.

The threshold energy value does not have a substantial effect on the distribution of Bayes factors. It is difficult to predict the effect of the threshold energy qualitatively, because there are two competing effects: a lower threshold would increase the sample size, which makes the non-uniformity more apparent; a higher threshold decreases the effective GZK horizon, which would increase the non-uniformity signature. This means that there is some ideal threshold that gives the greatest chance of detecting whatever anisotropy is present.

While the results for the uniform and non-uniform cases are clearly different, the difference is not very significant. If we take a threshold value of $\ln(B) = 5$ to represent a decisive detection, then anisotropy is detected only for 0.002% and 5% of the samples for source densities of $10^{-3.5}\,\mathrm{Mpc}^{-3}$ and $10^{-4.5}\,\mathrm{Mpc}^{-3}$, respectively. The conclusion is that the clustering expected from a realistic model of AGN-sourced UHECRs is too weak to be detected from a sample of 69 events. This is consistent with the results of Abreu et al. (2012).

The simulations were repeated for 100 samples of $N = 690$ UHECRs (i.e. $10 \times$ our PAO sample). The results are shown in the bottom half of Figure 3.4. The difference between the uniform and non-uniform cases becomes very apparent for all combinations of parameters. For source densities of $10^{-3.5}\,\mathrm{Mpc}^{-3}$ and $10^{-4.5}\,\mathrm{Mpc}^{-3}$, 22% and 93% of the Bayes factors are above the threshold of $\ln(B) = 5$. UHECR samples of 690 events are sufficient to detect self-clustering for a realistic model.

We assumed a pure proton composition of UHECRs. However, as discussed in Section 1.5, it is possible that UHECRs have a more complex mixed nuclear composition, including heavier nuclei such as iron. For iron, the magnetic deflection angle would be increased by a factor of 26, leading to a deflection of ~ 50 to ~ 250 deg. This makes it more difficult to associate the UHECRs with specific sources. However, the detection of clustering is also made easier by the fact that heavier nuclei lose more energy through additional scattering processes, which reduces the GZK horizon and thus the number of candidate sources. The energy loss length for cosmic rays with $E \gtrsim 5 \times 10^{20}\mathrm{eV}$ is reduced from $\sim 10\,\mathrm{MeV}$ for protons to $\sim 2\,\mathrm{MeV}$ for iron, which reduces the GZK horizon by a factor of ~ 5 (Stanev 2009). The net effect of these two factors will need to be established through additional simulations.

3.4 Analysis of the PAO data

We have applied the multi-stage Bayesian method described in Section 3.2 to the PAO data set in order to assess the anisotropy of the measured UHECR arrival directions. This data set consists of 69 events observed from 1 January 2004 to 31 December 2009, as described in Section 1.4.1. As the results depend to some extent on the way the data are split into the three subsets, Bayes factors were calculated for 1,000 different random, but equal sized, partitions. The cumulative distribution of Bayes factors is shown in Figure 3.5.

The Bayes factors were calculated for different partitions of the same sample. Apart from the distribution for the PAO data, Figure 3.5 also shows the distribution for a uniform sample of 69 UHECRs, as well as the distribution for a UHECR sample generated from a realistic AGN catalogue (with a source density of $10^{-3.5}\,\mathrm{Mpc}^{-3}$, $\kappa = 30$ and $\alpha = 2.0$). The results shown here differ from those shown in Figure 3.4, insofar as they result from different random partitions of a single sample (i.e., PAO, uniform or AGN-sourced) rather than being drawn from completely independent samples. However, the distributions produced using these two methods are comparable and the main conclusions remain unchanged.

A sensible way of dealing with the range of Bayes factors is to characterize their distribution by the arithmetic or geometric mean. There is no compelling reason to choose one over the other (see e.g. O'Hagan 1997), but the fact that the logarithm of the Bayes factor is symmetric between the two models suggests that the geometric mean is more natural. The geometric mean was 0.57 and the arithmetic mean was 1.26. From Equation 3.1, if we assume a prior probability of 0.5 for both models, we calculate mean posterior probabilities for the clustered model of 0.37 and 0.56 for the respective means. Thus, there is no clear preference for either of the models, and the data are consistent with both. We do not detect evidence for self-clustering. Figure 3.5 shows that for data sets of this size, the distributions of Bayes factors for the uniform and AGN-centred cases cannot be clearly distinguished. This is consistent with the results of Abreu et al. (2012).

3.5 Conclusions

We have developed a Bayesian method for the analysis of the self-clustering of points on a sphere and have applied it to a sample of 69 UHECRs detected by PAO up until 31 December 2009.

The method is a three-step Bayesian approach, in which the data are divided

into three subsets: the first two subsets of the data are used to generate a model of self-clustered UHECRs; the third subset is used to perform Bayesian model comparison between this self-clustered model and a uniform model of UHECRs. This approach is an extension of the Bayesian model comparison methods that were developed by Spiegelhalter & Smith (1982), Aitkin (1991), O'Hagan (1991) and O'Hagan (1995). Like the multi-step method that is presented here, those approaches are aimed to evaluate the marginal likelihood in cases when there is weak prior information on the model parameters.

There is some ambiguity in the partitioning of the full data set. In the present implementation, the total data set is divided into three subsets of equal size. However, it is possible that a different partitioning, or perhaps an average over partitions could make this method more effective. These issues will be explored in future work.

We tested our model comparison method on mock catalogues of UHECRs. The results for uniform UHECR arrival directions were compared to the results for UHECRs originating in AGNs from a realistic mock catalogue. UHECR clustering in a realistic AGN centred model is too weak to be detected in a sample of 69 events, but would be detectable in samples of 690 events. This is consistent with the results of Abreu et al. (2012).

We assumed a pure proton composition of the cosmic rays, but there are some indications that heavier nuclei are also part of the composition (Unger 2008). The effect of including heavier nuclei will be investigated through additional simulations.

For the PAO data, Bayes factors were calculated for different random partitions of the data. The geometric and arithmetic means of the Bayes factors were 0.57 and 1.26 respectively, corresponding to posterior probabilities of 0.37 and 0.56 for the clustered model. Thus, we did not find strong evidence for clustering in the PAO data, although the data are also consistent with the AGN-centred simulations.

It will be of great interest to repeat this analysis for greater UHECR sample sizes, as we have shown that samples of 690 events are sufficient to distinguish between uniform and clustered models. Recently, the PAO Collaboration presented an analysis of an extended sample of 231 events with $E_{\rm obs} \geq 5.2 \times 10^{19}$ eV (Aab et al. 2015), for which this analysis can be repeated. Looking further ahead, the planned Japanese Experiment Module Extreme Universe Space Observatory (JEM-EUSO, Adams Jr. et al. 2013) on the International Space Station (ISS) is scheduled for launch in 2017 and is expected to detect ~ 200 UHECRs annually over its five year lifetime.

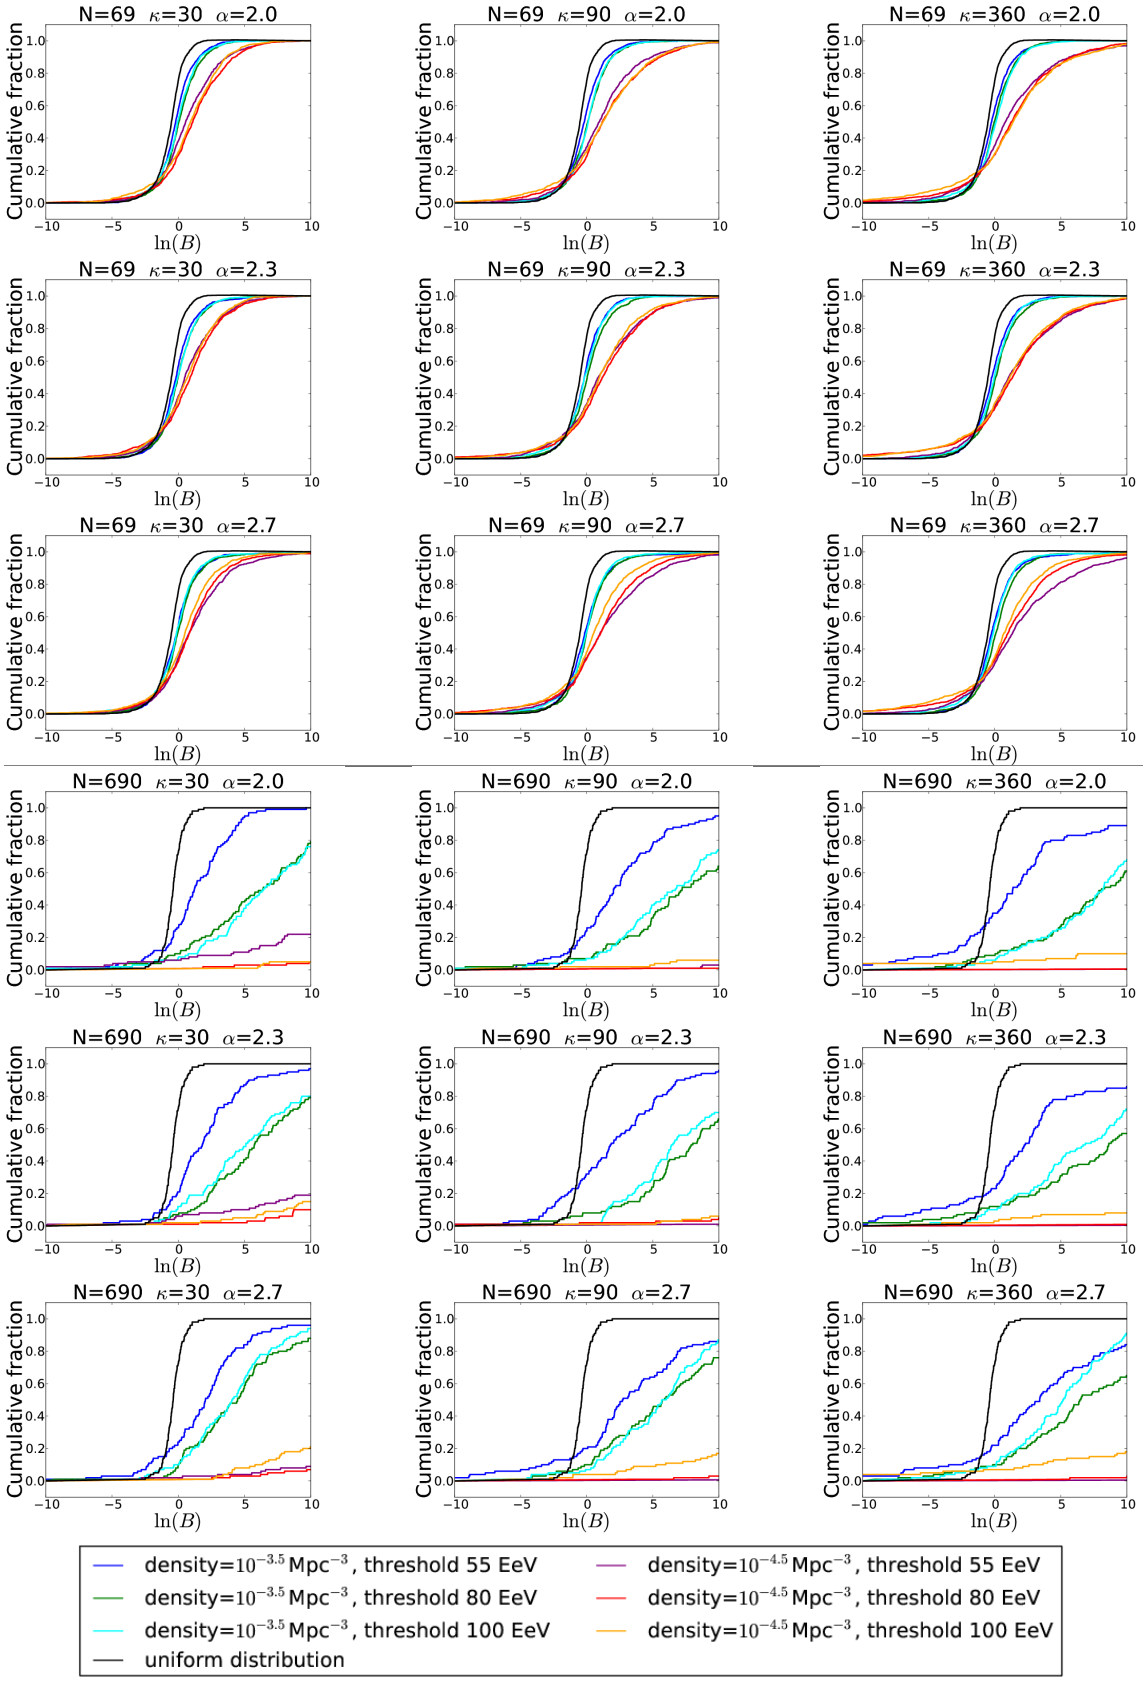


Figure 3.4: Results of the multi-step method applied to mock UHECR catalogues. Cumulative distributions of Bayes factors have been produced for three energy thresholds, two source densities, and for different values of the sample size N, the injection parameter α and the concentration parameter κ, as indicated above.

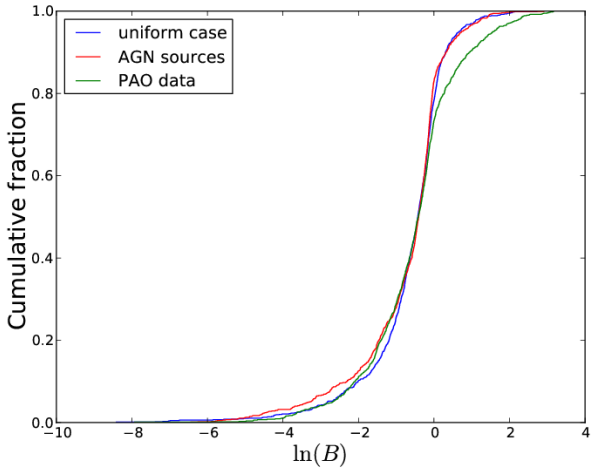


Figure 3.5: Cumulative fractions of Bayes factors, produced by the application of the multi-step method to 1,000 partitions of: (a) the PAO data; (b) 69 simulated UHECRs from uniform sources; and (c) 69 simulated UHECRs from a realistic mock catalogue of AGNs.

Chapter 4

A Bayesian analysis of UHECR source fractions

In this chapter we discuss a Bayesian method that we have developed to derive constraints on the source fraction of a given UHECR sample. The source fraction is defined as the fraction of rays in a UHECR sample that are expected to have originated at the sources in whichever source catalogue is under consideration. Our work can be regarded as a refinement and a generalization of previous studies on this subject. We have applied our method to a sample of 69 UHECRs from the PAO, and have determined the source fractions for AGNs in the Veron-Cety & Verson (VCV) catalogue, as well as AGNs detected with the Swift Burst Alert Telescope (Swift-BAT), galaxies from the 2MASS Redshift Survey (2MRS), and an additional volume-limited sample of 17 nearby AGNs. Conducting analyses for these source catalogues has enabled us to compare our results with several previous studies. For fiducial values of the model parameters, we report 68% credible intervals for the fraction of source originating UHECRs of $0.09^{+0.05}_{-0.04}$, $0.25^{+0.09}_{-0.08}$, $0.24^{+0.12}_{-0.10}$, and $0.08^{+0.04}_{-0.03}$ for the VCV, Swift-BAT and 2MRS catalogues, and the sample of 17 AGNs, respectively. For reasonable ranges of the prior parameters, the purely isotropic model is disfavoured.

4.1 Introduction

Chapter 1 discussed a range of previous studies of the correlations between UHECR arrival directions and catalogues of possible sources. No clear consensus on this question has been reached, in part due to the difficulty of analyzing such small sample sizes. Given the small size of the UHECR data sets, it is important to utilize as much of the available information as possible. This can be achieved by adopting a Bayesian methodology, that involves models of the relevant physical processes. The first steps to such a comprehensive Bayesian work have been made in the recent work of Watson et al. (2011) and Soiaporn et al. (2013).

Watson et al. (2011) analysed the 27 events that were previously investigated in Abraham et al. (2007), and derived a posterior for the fraction that originated from AGNs in the Veron-Cetty & Veron (VCV) catalogue (Véron-Cetty & Véron 2006). To do so, they used a two-component parametric model characterized by a source rate Γ and a background UHECR rate R. The model assumed that the UHECR arrival directions are points drawn from a Poisson intensity distribution on the celestial sphere. The intensity distribution was obtained with a computational UHECR model. Watson et al. (2011) report strong evidence of a UHECR signal from the VCV AGNs. They find a low AGN fraction that is consistent with Abreu et al. (2010). For fiducial values of the model parameters, they report a 68% credible interval for the AGN fraction of $F_{\rm AGN} = 0.09^{+0.05}_{-0.04}$.

Soiaporn et al. (2013) developed a multi-level Bayesian framework to attempt to associate the 69 UHECRs that were recorded at the PAO in the period 2004-2009 with 17 nearby AGNs catalogued by Goulding et al. (2010) (hereafter G10). They report evidence for a small but nonzero fraction of the UHECRs to have originated at the AGNs from G10, of the order of a few percent to 20%.

We extend the formalism of Watson et al. (2011) with a more refined UHECR model, and generalize it in such a way that it can be applied to a greater variety of source catalogues. With also apply the study to a greater data set, analyzing the 69 events that were recorded by the PAO in the period 2004-2009. As source catalogues, we consider AGNs from the VCV, Swift-BAT, and G10 catalogues, and galaxies from the 2MRS catalogue. This allows us to compare our results with the results of Watson et al. (2011), Abreu et al. (2010) and Soiaporn et al. (2013).

After discussing the UHECR and source data sets in Section 4.2, we explain our UHECR model in Section 4.3, discuss the statistical formalism of our Bayesian model comparison in Section 4.4, and the analysis of mock data sets in Section 4.5. The results of applying the formalism to the PAO data are discussed in Section 4.6.

4.2 Data

4.2.1 UHECR sample

The sample of UHECR events that was used in this analysis were the 69 UHECRs recorded at the PAO between January 2004 and November 2009 (see Section 1.4.1). This sample is displayed on Figure 4.1 on top of PAO's relative exposure.

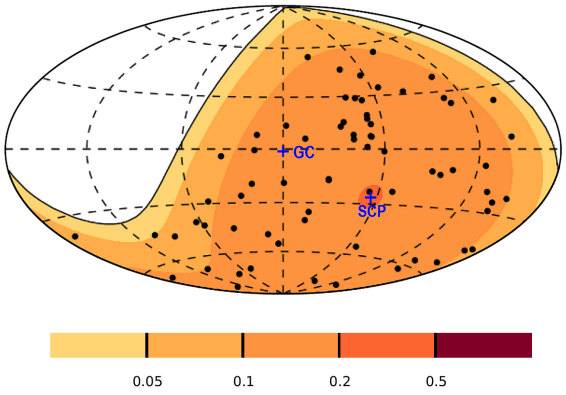


Figure 4.1: Relative PAO exposure in Galactic coordinates. The arrival directions of the 69 UHECRs are shown as black points. The Galactic centre (GC) and south celestial pole (SCP) are indicated.

4.2.2 Source catalogues

As potential source catalogues, we consider AGNs from the VCV, Swift-BAT and G10 catalogues, and galaxies from the 2MRS catalogue. This allows us to compare our analysis for the Swift-BAT and 2MRS sources with the analysis from Abreu et al. (2010), our analysis for the VCV sources with the analyses from both Abreu et al. (2010) and Watson et al. (2011), and our analysis of the G10 sources with Soiaporn et al. (2013).

We use the 12th edition of the VCV catalogue, selecting sources with $z_{\rm obs} \leq 0.03$, as AGNs with higher redshift are too far away to be plausible UHECR sources, and can be shown to have a negligible effect on the results. We omit sources for which absolute magnitudes are not stated. The total number of VCV AGNs that meet those requirements is $N_{\rm VCV} = 921$. This is the same sample of sources that was used in Abraham et al. (2007), Abreu et al. (2010) and Watson et al. (2011), and in PAO's more recent analysis Aab et al. (2015). While the VCV catalogue is heterogeneous and thus not ideal for statistical studies, it is close to complete for the low-redshift AGNs that are of relevance here.

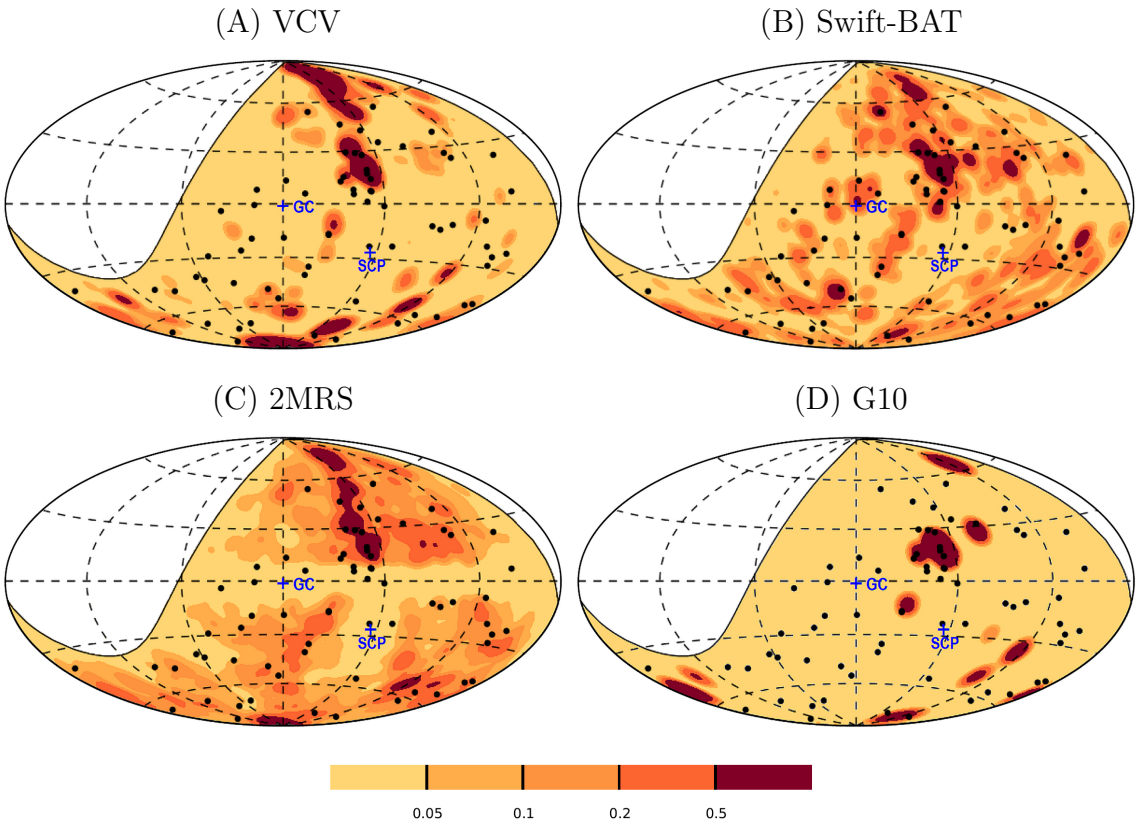


Figure 4.2: Positional dependence of the expected number of source originating events, for the VCV, Swift-BAT, 2MRS, and G10 catalogues. A fiducial value of the smearing parameter $\sigma = 3\,\mathrm{deg}$ is assumed. The arrival directions of the 69 UHECRs are shown as black points. Galactic coordinates are used, and the Galactic centre (GC) and south celestial pole (SCP) are indicated.

For the Swift-BAT catalogue, we use the 58 month version, which includes a total of $N_{\mathrm{BAT}} = 1092$ sources. In the case of the 2MRS catalogue, we use the catalogue version 2.4, 2011 Dec 16. We exclude events that are within 10° of the Galactic plane, to avoid biases due to the incompleteness of the catalogue in the region of the Galactic plane. This leaves a total of $N_{\mathrm{2MRS}} = 20,702$ galaxies. These samples of Swift-BAT and 2MRS sources are the same as those used by Abreu et al. (2010).

The G10 catalogue is a well-characterized volume-limited sample of AGNs. The 17 AGNs contained in it constitute all infrared-bright AGNs within 15 Mpc. This is the same sample that was used by Soiaporn et al. (2013).

For each of the catalogues, Figure 4.2 shows the simulated observed UHECR flux that was calculated by our models (see Section 4.3).

4.3 UHECR model

This analysis required a realistic model of UHECR injection, propagation, and detection, both to compute the likelihoods in our statistical formalism (Section 4.4), and to create simulated mock catalogues of UHECRs to test our methods (Section 4.5). The creation of mock catalogues was carried out using a very similar procedure to the one described in Section 3.3, with the sole change being the inclusion of the additional scattering process off the infrared background radiation. That additional process was incorporated in the same manner as the Bethe-Heitler scattering. For the computation of the likelihoods, the model was adapted somewhat. Instead of using the semi-stochastic approach that was used to generate mock-data sets, the likelihoods were computed by using a fully continuous model of the loss process. This is described in greater detail in Section 4.3.2.

4.3.1 Injection

We adopt a model in which any given UHECR source emits UHECRs with an emission spectrum given by

$$\mathrm{d}N_{\mathrm{emit}}/\mathrm{d}E_{\mathrm{emit}} \propto E_{\mathrm{emit}}^{-\alpha-1}, \tag{4.1}$$

where the logarithmic slope α was taken to be 3.6 (Abraham et al. 2010; Watson et al. 2011). The spectrum is normalized in such a way that the total emission rate of UHECRs with energy greater than E_{emit} is given by

$$\frac{\mathrm{d}N_{\mathrm{emit}}(>\ E_{\mathrm{emit}})}{\mathrm{d}t} = \Gamma \left(\frac{E_{\mathrm{emit}}}{E_{\mathrm{min}}} \right)^{-\alpha}, \tag{4.2}$$

where $E_{\mathrm{min}} = 5.5 \times 10^{19}\,\mathrm{eV}$ is the minimum UHECR emission energy and Γ is the rate at which the source emits UHECRs with $E_{\mathrm{emit}} > E_{\mathrm{min}}$.

4.3.2 Energy loss during propagation

As discussed in Section 1.3, the energy loss processes experienced by UHECRs can be characterized in terms of the loss length $L_{\mathrm{loss}} = -E(\mathrm{d}E/\mathrm{d}r)^{-1}$. If a continuous energy loss is assumed, then it is possible to calculate the total amount of energy that a UHECR loses as it travels to the Earth from a given distance by solving the differential equation

$$\frac{\mathrm{d}E}{\mathrm{d}r} = -\frac{E}{L_{\mathrm{loss}}(E)}. \tag{4.3}$$

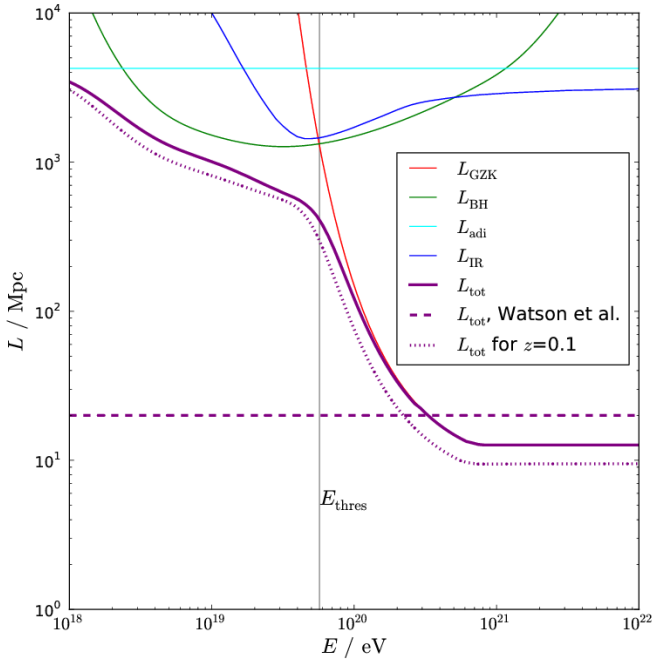


Figure 4.3: Loss lengths from the four energy loss processes, compared to the constant loss length used by Watson et al. (2011), as described in Section 4.3.2.

Our model takes into account all of the energy loss processes experienced by protonic UHECRs that are discussed in Section 1.3. The loss lengths corresponding to the processes are displayed in Figure 4.3. The energy dependence of $L_{\rm loss}(E)$ is one of the main improvements of this propagation model over the model used in Watson et al. (2011), where $L_{\rm loss}$ was taken to be a constant. The constant value of $L_{\rm loss}$ used by Watson et al. (2011) is also displayed in Figure 4.3 for comparison.

4.3.3 Effective smearing

To model that magnetic smearing of the UHECRs as well as the uncertainty in the measurement of their arrival direction, we used the same buckshot model that was described in Section 3.3.2.2. Once again, we conducted investigations for values of the smearing angle of $\sigma \simeq 3$, 6 and 10 deg, corresponding to values of the smearing parameter κ of 360, 90, and 30, respectively.

4.3.4 Observed UHECR flux

The number of UHECRs from source s above a threshold energy $E_{\rm thres}$ observed on Earth per unit area per unit time, $\mathrm{d}N_s(E_{\rm obs} \geq E_{\rm thres})/\mathrm{d}t\mathrm{d}A$, is a quantity that

is important in this statistical analysis. This rate is proportional to the rate of UHECRs emitted by the source, Γ_s, but it also depends on the distance-dependence of the UHECR energy loss, and on the UHECR injection spectrum. We use the UHECR propagation model described in Section 4.3.2 to determine the injection energy corresponding to the threshold energy E_{thres} and to the source distance D_s. Combining this value with Equation 4.2 and with the source distance D_s, we obtain

$$\frac{\mathrm{d}N_s(E_{\mathrm{obs}} \geq E_{\mathrm{thres}})}{\mathrm{d}t\,\mathrm{d}A} = \frac{\Gamma_s}{4\pi D_s^2}\left[\frac{E_{\mathrm{emit}}(E_{\mathrm{thres}})}{E_{\mathrm{min}}}\right]^{-\gamma}. \tag{4.4}$$

This expression assumes that the observed energy E_{obs} is equivalent to the arrival energy of the UHECR, E_{arr}. Thus, for the purposes of the calculation, the 12% energy uncertainty of the PAO measurements is neglected. The variation in source rates Γ_s among the sources that we are considering is not negligible. We use the source rate of Centaurus A as the reference value Γ. The source rate of a source s is obtained by weighing the flux F_s of that source in a particular band against the flux F_{Cen} of Centaurus A in that same band. The wave band of the flux thereby is different depending on the source catalogue. For VCV, the flux of the source in the V-band is used, for Swift-BAT the X-ray flux, for 2MRS the IR flux, for G10 the K-band flux. The fluxes are thus used as weights, so that sources with higher flux contribute more UHECRs. This approach is very similar to the approach used in Abreu et al. (2010), where fluxes were used to weigh the sources from the Swift-BAT and 2MRS catalogues in the same way. It is also an improvement over the methods used in Watson et al. (2011), where the sources were weighed by the inverse square of their distances, rather than the flux. Our approach is more general, as it allows us to compare sources of different luminosities.

Incorporating the fluxes into the formalism, we obtain the expression

$$\frac{\mathrm{d}N_s(E_{\mathrm{obs}} \geq E_{\mathrm{thres}})}{\mathrm{d}t\,\mathrm{d}A} = \frac{\Gamma}{4\pi D_{\mathrm{Cen}}^2}\frac{F_s}{F_{\mathrm{Cen}}}\left[\frac{E_{\mathrm{emit}}(E_{\mathrm{thres}})}{E_{\mathrm{min}}}\right]^{-\gamma}, \tag{4.5}$$

where D_{Cen} is the distance to Centaurus A.

4.4 Statistical formalism

Given a sample of UHECRs arrival directions, we would like to determine the fraction of these rays that have come from a set of sources under consideration. To do so, we use a two-component parametric model characterized by two rates: The source rate Γ and the isotropic background rate R. As elaborated in Section 4.3.4, we use the

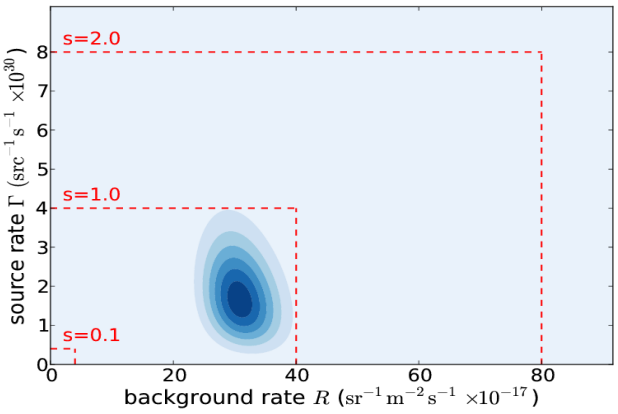


Figure 4.4: Example likelihood and priors for three values of s.

source rate of Centaurus A as the reference value of Γ. We obtain a joint posterior distribution for the two rates:

$$\Pr(\Gamma, R|\boldsymbol{d}) = \frac{\Pr(\Gamma, R)\Pr(\boldsymbol{d}|\Gamma, R)}{\int_0^\infty \int_0^\infty \Pr(\Gamma, R)\Pr(\boldsymbol{d}|\Gamma, R)\,\mathrm{d}\Gamma\,\mathrm{d}R}, \tag{4.6}$$

where $\Pr(\Gamma, R)$ is the prior distribution for Γ and R, and $\Pr(\boldsymbol{d}|\Gamma, R)$ is the likelihood (i.e., the probability of obtaining the data set $\boldsymbol{d}$ given values of Γ and R).

4.4.1 Prior

We adopt a uniform prior over Γ and R, with $\Gamma \geq 0$, $R \geq 0$. This plausibly encodes our ignorance of the two parameters, and, unlike maximum entropy priors, includes a possible value of 0 for both parameters. We have conducted our analysis for flat priors of varying width, using a variable width parameter s. The expression for the prior can be written as

$$\Pr(\Gamma, R) = \begin{cases} \frac{1}{s^2 \Gamma_{\max} R_{\max}} & \text{if } R < sR_{\max},\ \Gamma < s\Gamma_{\max} \\ 0 & \text{otherwise} \end{cases}. \tag{4.7}$$

The values of $\Gamma_{\max}$ and $R_{\max}$ were chosen differently depending on the likelihood, so that Bayes factors for different likelihoods could be compared. Figure 4.4 shows an example likelihood, and shows that $\Gamma_{\max}$ and $R_{\max}$ are chosen in such a way that when $s = 1$, the prior covers the 99.7% credible region implied by the likelihood and an infinitely broad uniform prior. This gives a data driven scaling for the rates, and is the same approach that we used in the example cases discussed in Chapter 2.

4.4.2 The likelihood

To compute the likelihood, we use a 'counts in cells' approach, in which the sky is divided into $1800 \times 3600 = 6{,}480{,}000$ pixels, which are distributed uniformly in right ascension and declination. Thus, the data set $\boldsymbol{d}$ can be rewritten as a set of counts in each pixel $\{N_{\mathrm{c},p}\}$.

The likelihood $\Pr(\boldsymbol{d}|\,\Gamma, R)$ is then given by a product of the individual Poisson likelihoods in each pixel,

$$\begin{aligned}&\Pr(\boldsymbol{d}|\Gamma, R)\\ &= \prod_{p=1}^{N_\mathrm{p}} \frac{(\overline{N}_{\mathrm{src},p} + \overline{N}_{\mathrm{bkg},p})^{N_{\mathrm{c},p}} \exp[-(\overline{N}_{\mathrm{src},p} + \overline{N}_{\mathrm{bkg},p})]}{N_{\mathrm{c},p}!},\end{aligned} \tag{4.8}$$

where $\overline{N}_{\mathrm{src},p}$ and $\overline{N}_{\mathrm{bkg},p}$ are the expected counts in pixel p due to sources and background, respectively.

The expected number of counts in pixel p that are contributed by the background is

$$\overline{N}_{\mathrm{bkg},p} = R \int_p \frac{\mathrm{d}\epsilon}{\mathrm{d}\Omega}\,\mathrm{d}\Omega_{\mathrm{obs}}, \tag{4.9}$$

where the integral is over the pixel p, and $\mathrm{d}\epsilon/\mathrm{d}\Omega$ is the relative exposure (Section 4.2). The expected number of source originating events in pixel p

$$\overline{N}_{\mathrm{src},p} = \sum_{s=1}^{N_\mathrm{s}} \frac{\mathrm{d}N_s(E_{\mathrm{obs}} \geq E_{\mathrm{thres}})}{\mathrm{d}t\,\mathrm{d}A} \int_p \frac{\mathrm{d}\epsilon}{\mathrm{d}\Omega} \Pr(\tilde{r}_{\mathrm{obs}}|\tilde{r}_s)\,\mathrm{d}\Omega_{\mathrm{obs}}, \tag{4.10}$$

where the sum is over the sources, $\Pr(\tilde{r}_{\mathrm{obs}}|\tilde{r}_s)$ is a vMF distribution, and $\mathrm{d}N_s(E_{\mathrm{obs}} \geq E_{\mathrm{thres}})/\mathrm{d}t\mathrm{d}A$ is the observed UHECR flux discussed in Section 4.3.4. Inserting Equations 4.9 and 4.10 into Equation 4.8, we arrive at the full likelihood.

The positional dependence of $\overline{N}_{\mathrm{bkg},p}$ follows the relative exposure of PAO, as shown in Figure 4.1. The positional dependence of $\overline{N}_{\mathrm{src},p}$ depends both on the PAO exposure and on the distribution of sources in the given catalogue. Figure 4.2 shows the dependence for the four catalogues that are used in this study. The dependence is dominated by the distribution of local AGNs, by far the strongest source being Centaurus A ($l = 309.5°$, $b = 19.4°$), which previously studies (e.g. Abraham et al. 2007) have suggested as the dominant UHECR source.

4.4.2.1 Rearranging the expression for the likelihood

The expression for the likelihood proved to be inefficient for use, as it required a great number of computations: The total number of pixels was $N_{\rm p} = 1800 \times 3600 = 6{,}480{,}000$. If a $\Gamma \times R$ grid of 100×100 is used, a total of 64,800,000,000 calculations would be required.

The total number of calculations can be greatly reduced by rearranging the expression. For a given data set, we can separate the product of Equation 4.8 into a product over those pixels that include an event, $\{q\}$, and pixels that do not, $\{r\}$. Using the fact that $N_q = 1$ for all $\{q\}$ and $N_r = 0$ for all $\{r\}$, we can write

$$\Pr(\boldsymbol{d}|\Gamma, R) = \prod_{r=1}^{N_{\rm r}} \exp[-(\overline{N}_{{\rm src},r} + \overline{N}_{{\rm bkg},r})] \times \prod_{q=1}^{N_{\rm q}} (\overline{N}_{{\rm src},q} + \overline{N}_{{\rm bkg},q}) \exp[-(\overline{N}_{{\rm src},q} + \overline{N}_{{\rm bkg},q})] \tag{4.11}$$

$$= \exp[-(\Gamma\Sigma_{\rm src} + R\Sigma_{\rm bkg})] \times \prod_{q=1}^{N_{\rm q}} (\overline{N}_{{\rm src},q} + \overline{N}_{{\rm bkg},q}) \exp[-(\overline{N}_{{\rm src},q} + \overline{N}_{{\rm bkg},q})]. \tag{4.12}$$

where $\Sigma_{\rm src} = \sum_{r=1}^{N_{\rm r}} m_{{\rm src},r}$ and $\Sigma_{\rm bkg} = \sum_{r=1}^{N_{\rm r}} m_{{\rm bkg},r}$, and $m_{{\rm src},p}$ and $m_{{\rm bkg},p}$ are two pixelized maps obeying the equations

$$\overline{N}_{{\rm src},p} = \Gamma m_{{\rm src},p} \tag{4.13}$$

$$\overline{N}_{{\rm bkg},p} = R m_{{\rm bkg},p}. \tag{4.14}$$

Thus, the initial expression has been rearranged in such a way that the vast majority of Poisson calculations is contained within the sums $\Sigma_{\rm src}$ and $\Sigma_{\rm bkg}$. These sums can be calculated in advance for the entire grid of Γ and R. This greatly reduces the total number of calculations required for Equation 4.8, and speeds up the full calculation by a factor of $\sim 10^5$.

4.4.2.2 Comparison of the likelihood to the On/Off example

Figure 4.5 displays the Poisson product likelihood for the PAO data, and displays prior regions for values of $s = 0.1, 1, 2$, and posteriors for those regions. The posteriors behave in a similar way to the On/Off case discussed in Section 2.5.3, in that for small values of s, the entire posterior becomes concentrated at the highest values of R and Γ. The product of Equation 4.8, for low values of Γ and R, reduces

to

$$\Pr(\boldsymbol{d}|\Gamma, R) = \prod_{q=1}^{N_{\rm q}} (\Gamma m_{{\rm src},q} + R m_{{\rm bkg},q}). \tag{4.15}$$

As $N_{\rm q} = 69$, the function becomes extremely steep in Γ and R. Mathematically, the On/Off likelihood can be regarded as a special case of the Poisson product of Equation 4.8.

4.4.3 The source fraction

The source fraction[1] is defined as the fraction of the UHECRs expected to have originated at the sources in whichever catalogue is under consideration and is given by

$$F_{\rm src}(\Gamma, R) = \frac{\sum_{p=1}^{N_{\rm p}} \overline{N}_{{\rm src},p}}{\sum_{p=1}^{N_{\rm p}} \overline{N}_{{\rm src},p} + \overline{N}_{{\rm bkg}.p}}. \tag{4.16}$$

The posterior for $F_{\rm src}$ can be calculated from the posterior over the rates as

$$\begin{aligned} &\Pr(F_{\rm src}|\boldsymbol{d}) \\ &= \int_0^{\Gamma_{\rm max}} \int_0^{R_{\rm max}} \Pr(\Gamma, R|\boldsymbol{d})\, \delta_{\rm D}[F_{\rm src} - F_{\rm src}(\Gamma, R)]\, {\rm d}\Gamma\, {\rm d}R. \end{aligned} \tag{4.17}$$

$\Pr(F_{\rm src}|\boldsymbol{d})$ is insensitive to $R_{\rm max}$ and $\Gamma_{\rm max}$ provided they are sufficiently large.

4.4.4 Model comparison

We would like to compare model M_1 where all the UHECRs are drawn from a uniform distribution with model M_2 where the UHECRs are derived from a combination of a background and a source originating component. To do this, we conduct a Bayesian model comparison. As described in Section 2.5, for a data set $\boldsymbol{d}$, and two models M_1 and M_2 the Bayes factor is

$$B_{12} = \frac{\Pr(\boldsymbol{d}|M_1)}{\Pr(\boldsymbol{d}|M_2)}. \tag{4.18}$$

In the specific case that is considered here, the models are nested: When $\Gamma = 0$, model M_2 reduces to model M_1. This means that we can use the expression for the

[1]The source fraction $F_{\rm src}$ is equivalent to the AGN fraction $F_{\rm AGN}$ used in Watson et al. (2011) but now generalized to allow for non-AGN progenitors.

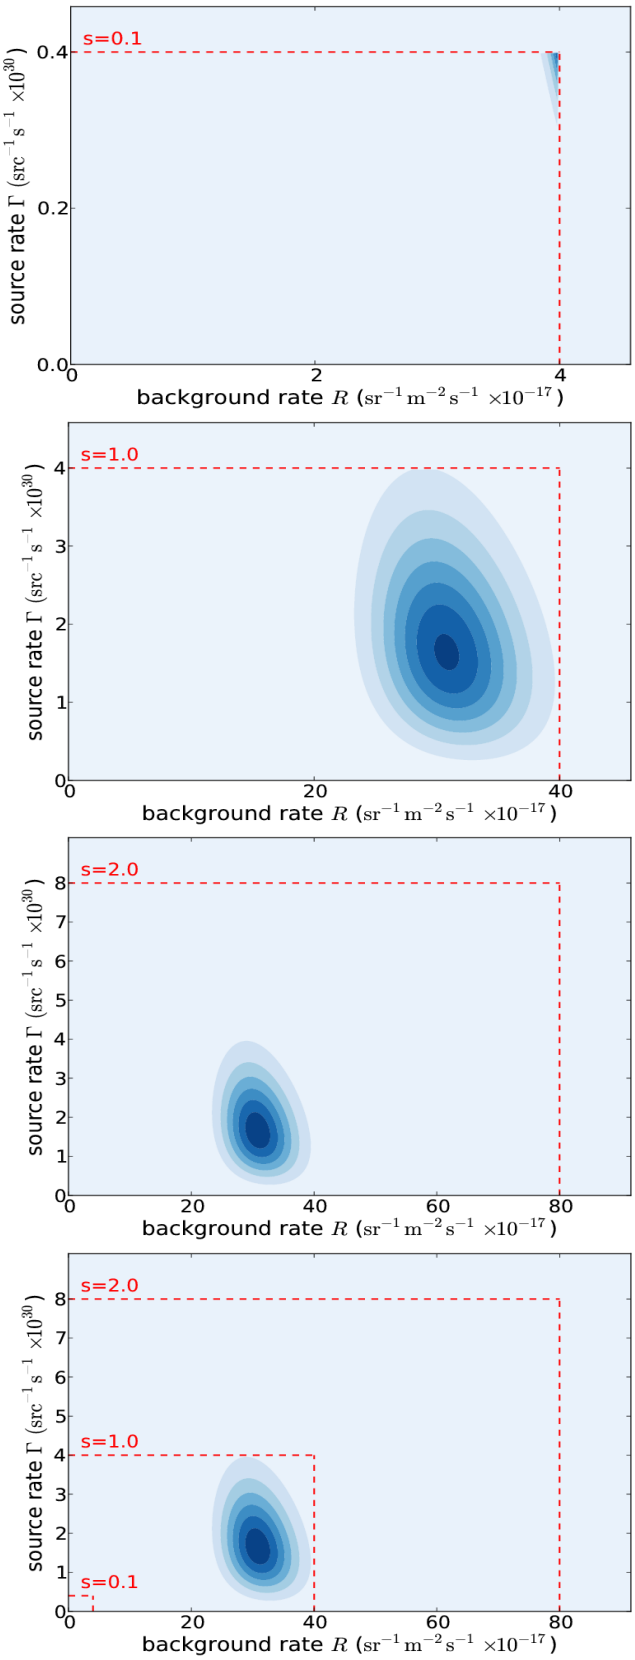


Figure 4.5: Lower panel: Poisson product likelihood for the PAO data. The red lines denote prior regions for three different values of the hyperparameter s. Upper three panels: Posteriors for the same s values are displayed.

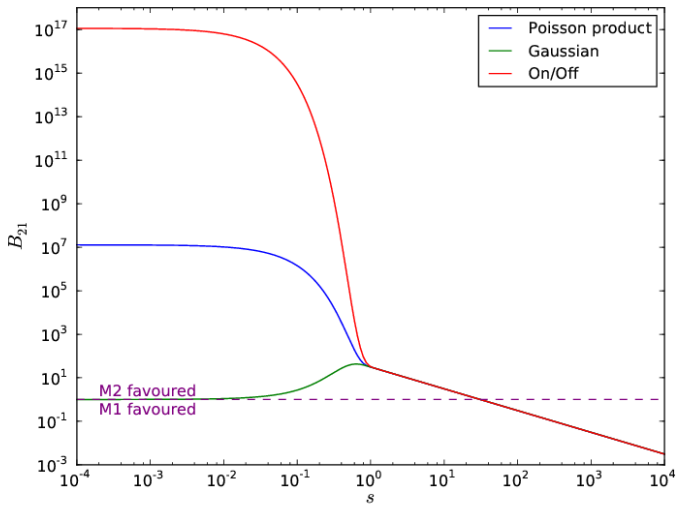


Figure 4.6: Dependence of the Bayes factor on hyperparameter s for the cases of Gaussian likelihood, the Poisson product likelihood, and the On/Off likelihood.

SDDR discussed in Section 2.5.2, and can write

$$B_{12} = \frac{\Pr(\Gamma = 0|\boldsymbol{d}, M_2)}{\Pr(\Gamma = 0|M_2)}. \tag{4.19}$$

Figure 4.6 shows a plot of Bayes factors that have been calculated for the Poisson product likelihood, for different values of the hyperparameter s. This plot is compared with similar plots for the On-Off likelihood and the Gaussian likelihood. The Bayes factors that are shown in the figure are the Bayes factors favouring the complex model, $B_{21} = 1/B_{12}$. For all three cases, B_{21} falls linearly for $s > 1$. For lower values of s, the Bayes factor for the Gaussian case approaches 1, while for the PAO and On/Off cases B_{12} becomes $\gg 1$, as the uniform model is extremely disfavoured. The Bayes factors for the On/Off case behave very similarly to the Poisson product case, as the former can be regarded as a special case of the latter.

4.5 Simulations

In order to investigate the constraining power of a data set of 69 events, we apply the method to simulated data sets. We use two extreme cases:

(i) Uniform arrival directions. These rays were drawn from a probability distribution that followed the PAO exposure.

(ii) UHECRs originating at sources from a catalogue. We conducted simulations

for all four of the catalogues. In each catalogue, the sources were weighted by their fluxes and the PAO exposure. Random sources were then selected, and the semi-stochastic propagation model of Section 4.3 was used to propagate rays from the sources to the Earth.

The posteriors for the source and background rates, as well as the posteriors for the source fraction, are summarized in Figure 4.7. The posteriors for the uniform and source centred cases are completely disjoint, which demonstrates that in extreme scenarios where all UHECRs originate either from a uniform background or from a source catalogue, a data set of 69 events should be sufficient to distinguish between the two models. Figure 4.7 also shows the Bayes factors as functions of s for the two cases. The Bayes factors B_{21} that are displayed are the inverses of the SDDR given in Equation 4.19, and favour the more complex model for Bayes factors > 1.

To assess the results of the Bayes factor simulations, we can derive a rough range of plausible values of s from physical models, and then look at the behaviour of the Bayes factors at those physically plausible values. Plausible models of UHECR injection predict that the UHECR luminosity of a source like Centaurus A is of the order of $2.9 \times 10^{39}\,\mathrm{erg\,s^{-1}} \simeq 1.81 \times 10^{51}\,\mathrm{eV\,s^{-1}}$ (Fraija et al. 2012). If this is taken as the typical UHECR luminosity of a source, then for a UHECR energy range of $(5.5 - 100) \times 10^{19}\,\mathrm{eV}$, the range of source rates can be calculated by dividing the UHECR luminosity by the limiting values of this range. The result of this calculation is a range of source rates Γ of roughly $(2 - 33) \times 10^{30}\,\mathrm{s^{-1}}$. The values of s corresponding to this range have been marked on Figure 4.7. (The values are slightly different for each of the simulations. For the sake of clarity, only the values for the uniform simulation are displayed, the others being broadly similar.) For the sourced case, model M_2 is strongly favoured for all physically plausible values of s, while for the uniform case, the simple uniform model M_1 is favoured for the physically plausible values.

4.6 Results

The results of the application of the statistical methods described in Section 4.4 to the data described in Section 4.2 are shown in Figures 4.8 and 4.9. Figure 4.8 contrasts the results from our analysis with the equivalent results from Watson et al. (2011), and with the results for an intermediate case. The use of a more refined propagation model leads to a higher posterior probability for lower source rates. The reason for that is that in Watson's propagation model, the energy loss length is constant and very small (Figure 4.3). UHECRs experience more drastic energy loss

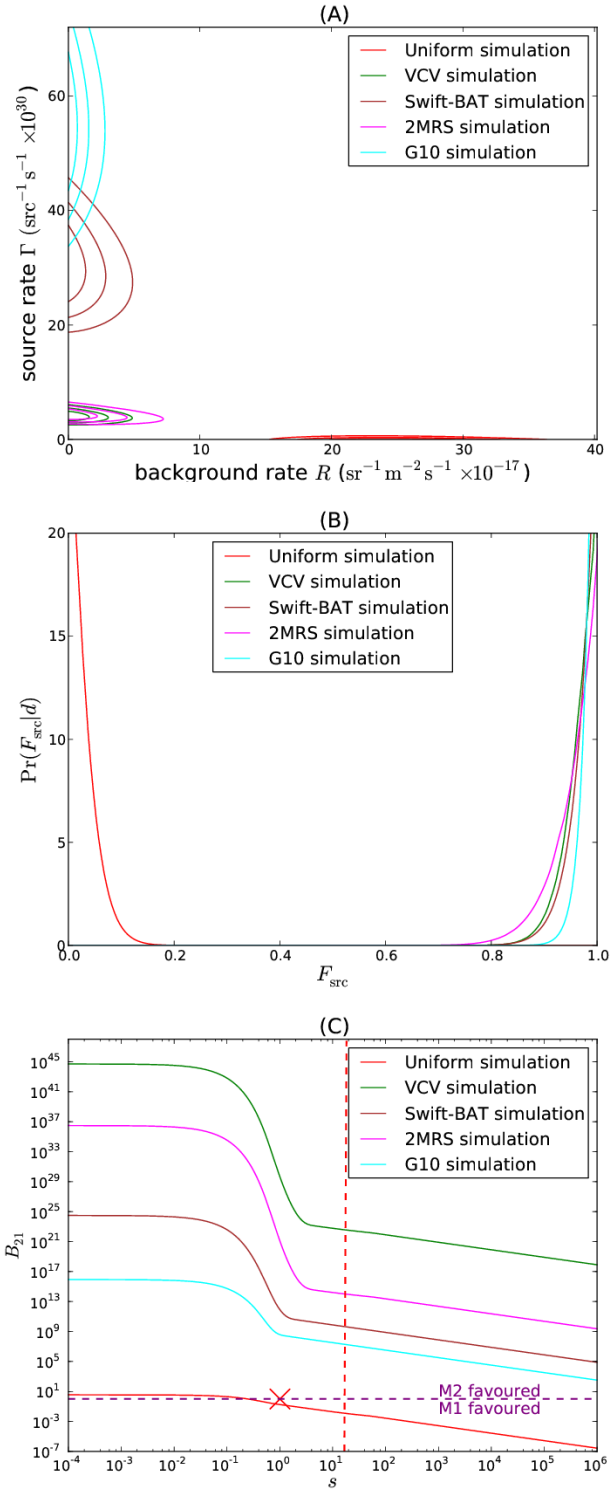


Figure 4.7: Results from simulations: Uniform UHECRs, and UHECRs originating at sources from the 4 catalogues. In all cases, 69 events are used. (A) Posteriors for Γ and R. The contours are the 68.3%, 95.4% and 99.7% highest posterior density credible regions. (B) Posteriors for the source fraction. (C) Plot of Bayes factors B_{21} as a function of the hyperparameter s. In (C), the $\times$-mark and the vertical line signify the minimum and the maximum values of the physically plausible range of s. The minimum and maximum values that are displayed correspond to the uniform simulation.

than in the more realistic model, which leads to more distant AGNs being excluded as plausible source candidates. As fewer sources are included, a higher source rate is required to generate the same sample of UHECRs.

The inclusion of 69 events reduces the extent to which the non-uniform model is favoured. This is evident from the posterior of the source fraction, and also from the behaviour of B_{21}. This result agrees with the results of Abreu et al. (2010), which reported that the full 69 events yield lower evidence of anisotropy than the earlier study Abraham et al. (2007), which analysed 27 events (see Section 1.5.3.3).

Figure 4.9 shows results for all four of the source catalogues, and for all values of the smearing parameter. Displayed are the posteriors for the source fraction, as well as plots of B_{21} against s. The constraints on the source fraction for all cases are shown in Table 4.1. The figures and table show that for greater smearing, the range of plausible values of $F_{\rm src}$ is increased, and the most probable value of the source fraction is higher than for the fiducial model of $\sigma = 3\,\deg$. The reason is that for greater magnetic deflection, the UHECR intensity distribution becomes more uniform, so that the uniform and mixed models become more difficult to distinguish, and a greater range of $F_{\rm src}$ values become viable.

The plots of B_{21} demonstrate that for all physically plausible prior ranges of the model parameters, the fully isotropic model is disfavoured.

These results for the VCV, Swift-BAT, and 2MRS catalogues can be compared with the results of Abreu et al. (2010), who used a correlation-based analysis on the VCV catalogue that mirrored the analysis in Abraham et al. (2007). Abreu et al. (2010) reported a correlation of $(38^{+7}_{-6})\%$ between UHECRs and sources from the VCV catalogue, which was considerably lower than the $(69^{+11}_{-13})\%$ correlation that was reported in Abraham et al. (2007). This reduction in the correlation is consistent with our findings that the source fraction is reduced as we increase the data set from 27 to 69 events. In addition to these correlation based methods, Abreu et al. (2010) conducted a likelihood based study similar to the analysis presented here, where the likelihood was taken as a probability map of arrival directions of UHECRs, parametrized by a magnetic smoothing angle σ and a fraction of isotropic rays $f_{\rm iso}$, which is equivalent to $1 - F_{\rm src}$. These likelihood-based studies were conducted for the Swift-BAT and 2MRS catalogues. For the 2MRS case, the maximum likelihood values of $f_{\rm iso}$ and σ are reported as 0.56 and 7.8^o, respectively. The σ value lies between our chosen smearing angles 6 deg and 10 deg. The value for $f_{\rm iso}$ corresponds to a value of $F_{\rm src}$ of 0.44, which is consistent with our $F_{\rm src}$ credible intervals for these chosen smearing angles. For the case of Swift-BAT, the maximum likelihood value of $f_{\rm iso}$ is given as 0.64, which corresponds to a source fraction of 0.36. The maximum

Table 4.1: Maximum a posteriori estimates and 68% credible intervals for $F_{\rm src}$.

Catalogue	$\sigma = 3\,\mathrm{deg}$	$\sigma = 6\,\mathrm{deg}$	$\sigma = 10\,\mathrm{deg}$
VCV	$0.09^{+0.05}_{-0.04}$	$0.14^{+0.07}_{-0.06}$	$0.22^{+0.09}_{-0.08}$
Swift-BAT	$0.25^{+0.09}_{-0.08}$	$0.37^{+0.11}_{-0.10}$	$0.46^{+0.13}_{-0.12}$
2MRS	$0.24^{+0.12}_{-0.10}$	$0.33^{+0.14}_{-0.14}$	$0.40^{+0.15}_{-0.15}$
G10	$0.08^{+0.04}_{-0.03}$	$0.14^{+0.06}_{-0.05}$	$0.22^{+0.07}_{-0.07}$

likelihood estimate of the smearing angle is reported as 1.5^o, which is lower than our minimum chosen value of 3 deg. Despite the difference between the angles, a $F_{\rm src}$ value of 0.36 can still be considered broadly consistent with the 68% credible interval for 3 deg, $0.25^{+0.09}_{-0.08}$.

Our results for the G10 catalogue can be compared with the work of Soiaporn et al. (2013). That analysis involved the full data set of 69 events, and found evidence for small but nonzero values of $F_{\rm src}$, of the order of a few percent to 20%, ruling out values of $F_{\rm src} > 0.3$. This is broadly consistent with our results, which suggest that values of $F_{\rm src} < 0.3$ are the most probable for all values of the smearing parameter.

These results assume a pure proton composition of UHECRs. As discussed in Section 1.3.1, the energy loss for heavier nuclei is similar for most of the relevant energy range, but becomes significantly stronger at the highest energies. To take into acount the energy loss for heavier nuclei, our UHECR model would need to be changed, as the current model does not account for the changes in composition that heavier UHECRs experience during spallation. A possible avenue for future research could be the incorporation of a computational framework such as GALPROP (Moskalenko & Seo 2014) into our statistical formalism, which does account for the spallation processes. Rough calculations within the existing model have shown that the differences in the energy loss at the highest energies do not affect the main result of our analysis, as the fully uniform model is still disfavoured for all plausible values of s.

One possible avenue for future research is the inclusion of the UHECR arrival energies in the analysis. At the moment, the arrival energies are only taken into account in the sense that only events above a certain energy threshold are registered. However, some previous studies included them directly in their models (e.g. Soiaporn et al. 2013). It is expected that more energetic UHECRs arrive from closer sources, so that the arrival energies are an important piece of information that is ignored in a method that only focuses on the arrival directions.

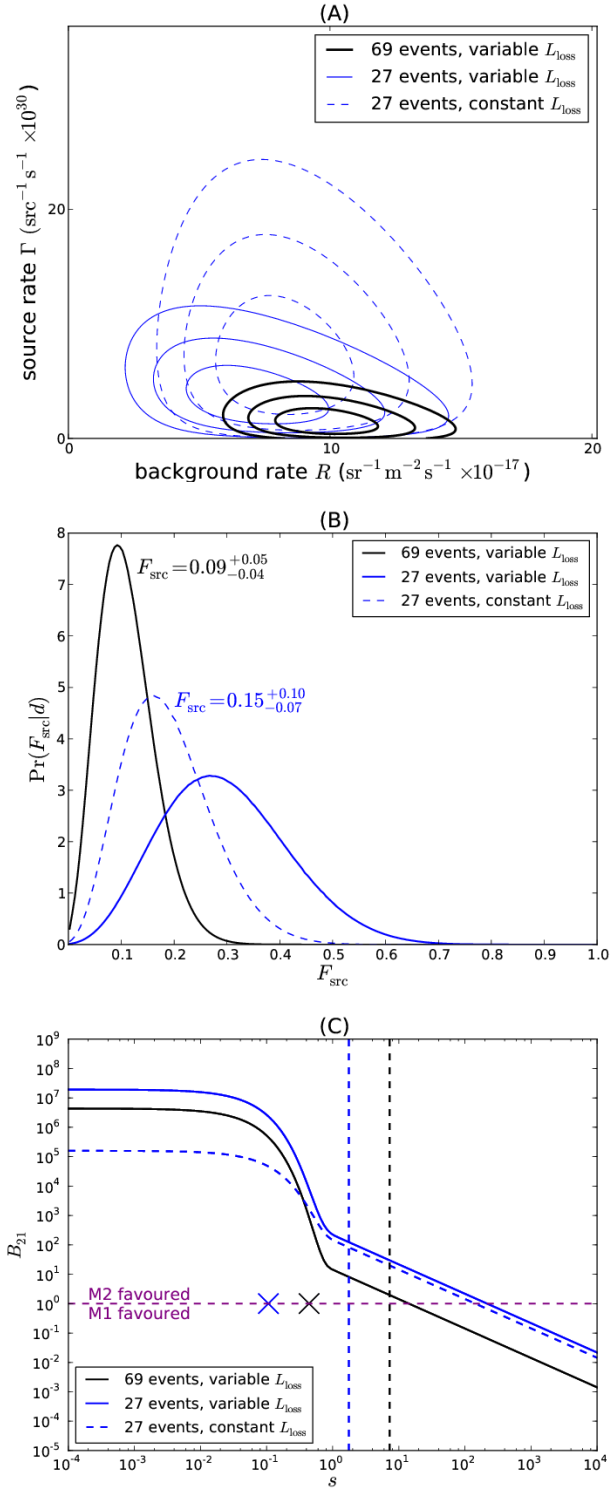


Figure 4.8: Results for $\sigma = 3\,\text{deg}$, and the sources from the VCV catalogue. Results for 27 and 69 events, and for constant and variable loss lengths are displayed. (A) Posteriors for the source and background rates. The contours are the 68.3%, 95.4% and 99.7% highest posterior density credible regions. (B) Posterior for the source fraction. (C) Plot of Bayes factors B_{21} as a function of the hyper-parameter s. In (C), physically plausible ranges of s are shown for the cases of 27 events (blue) and 69 events (black), with a variable loss length. The $\times$-marks and the vertical lines signify the minimum and the maximum values of the physically plausible ranges of s.

4.7 Conclusions

We have developed a Bayesian method that determines constraints on the source fraction of a UHECR sample, where the source fraction is defined as the fraction of the sample that is expected to have originated at the sources in a given source catalogue. Our method can be regarded as an extension of the method discussed in Watson et al. (2011). Our method uses a more refined UHECR model than Watson et al. (2011), and has been generalized in such a way that it can be applied to a greater range of possible source catalogues.

We have applied this method to the 69 UHECRs detected by the PAO with energies $E_{\rm obs} > 5.5 \times 10^{19}\,{\rm eV}$ in the period 2004-2009 to determine the fraction of these UHECRs that originated from AGNs from the VCV, Swift-BAT, and G10 catalogues, and galaxies from the 2MRS catalogue. Applying our method to a number of source catalogues has enabled us to compare our results with several previous studies, Watson et al. (2011), Abreu et al. (2010) and Soiaporn et al. (2013).

For the fiducial magnetic smearing parameter of $\sigma = 3$ deg, we report 68% credible intervals for the source fraction of $0.09^{+0.05}_{-0.04}$, $0.25^{+0.09}_{-0.08}$, $0.08^{+0.04}_{-0.03}$ and $0.24^{+0.12}_{-0.10}$ for the VCV, Swift-BAT, G10 and 2MRS catalogues, respectively. For all physically plausible values of the model parameters, the fully uniform model is disfavoured. The results of our study are in broad agreement with the previous work.

Future extensions of this method could include further refinements of the UHECR model, by incorporating a more realistic treatment of the magnetic fields, incorporating a way of dealing with heavier nuclei, and taking into account the UHECR arrival energies. Our analysis can be extended to greater UHECR catalogues.

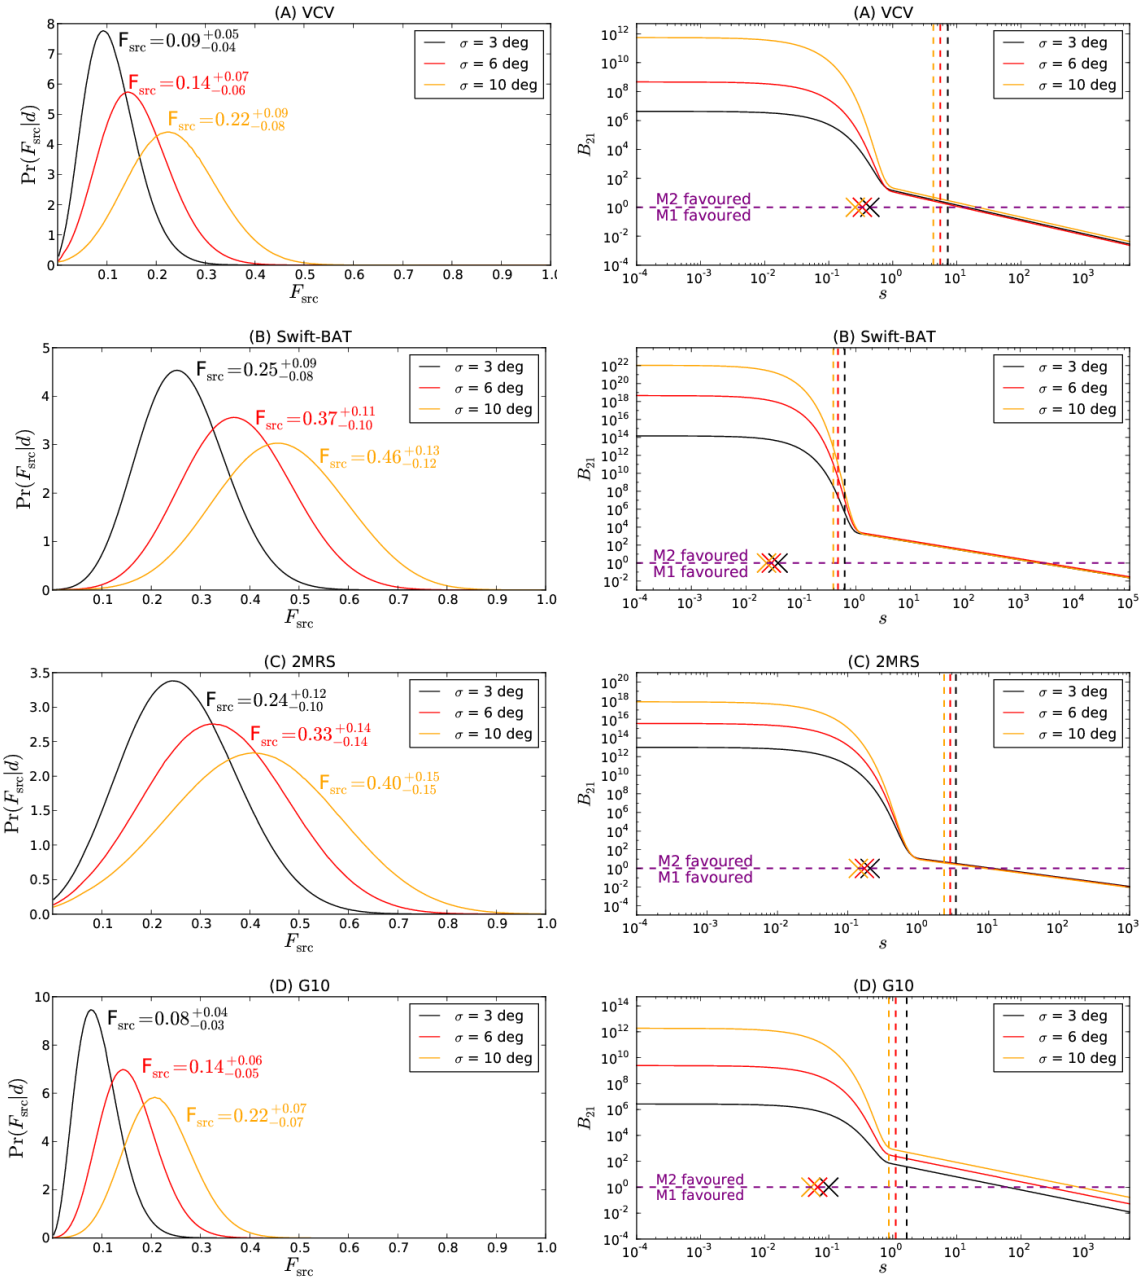


Figure 4.9: Posteriors of the source fraction, and plots of B_{21} against the hyper-parameter s, for the three smearing angles $\sigma = 3$ deg, 6 deg and 10 deg, and for the three source catalogues (A) VCV, (B) Swift-BAT, and (C) 2MRS. The plots of B_{21} show physically plausible ranges of s: The $\times$-marks and the vertical lines signify the minimum and the maximum values of these ranges.

Chapter 5

Thesis Conclusions

In this work, we have developed Bayesian methods for the analysis of UHECRs, and have applied these methods to the sample of 69 events that were recorded by the PAO in the period 2004-2009. Our methods have explored both of the approaches to the study of UHECR arrival directions: the analysis of catalogue-independent clustering of the particles, and the study of associations between the particles and source catalogues, which in our case was done by determining constraints on the source fraction of the UHECR sample.

The application of Bayesian methods to the study of the catalogue-independent clustering of UHECRs is challenging, as this problem lacks a well-motivated anisotropic model that could be weighed against the isotropic model. To solve this problem, we have developed a multi-step Bayesian method that derives the anisotropic model from a subset of the full UHECR data. This method is similar to two-step approaches that have been used in the past for situations with poorly specified prior parameters (e.g. Spiegelhalter & Smith (1982), Aitkin (1991), O'Hagan (1991) and O'Hagan (1995)). We have developed a model of UHECR injection, propagation, and detection, and have used it to create mock data-sets on which the multi-step method could be tested. Our analysis showed that a UHECR sample of 69 is not sufficient for our method to distinguish between a uniform and an anisotropic model, so that the PAO data set was found to be consistent with both uniformity and anisotropy. We calculated Bayes factors for different partitions of the data set. The geometric and arithmetic means of the Bayes factors were 0.57 and 1.26, respectively, so that we do not report strong evidence of anisotropy.

Our approach to the study of associations between UHECRs and source catalogues was similar to the work of Watson et al. (2011), the main differences being our use of a more refined propagation model, and a more general methodology that can be used for catalogues other than AGNs. We have analyzed associations with AGNs from the VCV, Swift-BAT, and G10 catalogues, and galaxies from the 2MRS

catalogue, which enabled us to compare our results with the work of Abreu et al. (2010) and Soiaporn et al. (2013) as well as Watson et al. (2011). We applied our method to the 69 events recorded by the PAO between 2004 and 2009, rather than the smaller sample of 27 events analyzed by Watson et al. For fiducial values of the model parameters, we report 68% credible intervals for the source fraction of $0.09^{+0.05}_{-0.04}$, $0.25^{+0.09}_{-0.08}$, $0.08^{+0.04}_{-0.03}$ and $0.24^{+0.12}_{-0.10}$ for these respective catalogues. For all physically plausible values of the model parameters, the fully isotropic models were disfavoured. These values are broadly consistent with the results of the previous studies.

There are several ways this work could be continued and extended. The multi-step method that was used in Chapter 3 is not specific to the UHECR problem and could have broad applications for anisotropy searches in other areas of astronomy, such as the search for angular anisotropies in the distribution of gamma-ray bursts described by e.g. Balazs et al. (1998) and Magliocchetti et al. (2003). An open question with regard to the multi-step method is whether there is an ideal way to partition the data, rather than splitting it into subsets of equal size.

Our UHECR models included a number of simplifications. We modelled the magnetic deflection of the UHECRs as a simple Gaussian smearing, ignoring possible direction-dependent UHECR deflections due to structured magnetic fields. We focused on protonic UHECRs, and our model cannot simulate the spallation processes of heavier nuclei. Future work could incorporate more sophisticated models of these processes into our formalism. More refined models of the relevant magnetic fields have been discussed by e.g. Medina Tanco et al. (1998), Nagar & Matulich (2010), Farrar et al. (2013) and Keivani et al. (2015). An example of a propagation model that takes into account spallation is GALPROP (Moskalenko & Seo 2014).

The search for associations between UHECR arrival directions and catalogues of potential sources could be extended by taking into account the arrival energies of the UHECRs. It is expected that more energetic UHECRs arrive from closer sources, so that the arrival energies are an important piece of information that is not taken into account by the present method. A Bayesian framework taking the energies into account has been discussed in Soiaporn et al. (2013).

It will be very interesting to repeat our analyses for samples of greater size, such as the 241 events with with $E_{\rm obs} \geq 5.2 \times 10^{19}\,{\rm eV}$ that were recently analyzed in Aab et al. (2015), and especially to the samples of ~ 1000 events that are expected to be produced by the JEM-EUSO mission (Adams Jr. et al. 2013). By applying the methods developed here to the large UHECR samples that will soon be available it should be possible to determine the origin of these particles.

Bibliography

Aab A., et al., 2015, Astrophys. J., 804, 15

Abbasi R. U., et al., 2008a, Astropart. Phys., 30, 175

Abbasi R. U., et al., 2008b, Phys. Rev. Lett., 100, 101101

Abbasi R. U., et al., 2010a, Phys. Rev. Lett., 104, 199902

Abbasi R., et al., 2010b, Astrophys. J. Lett., 713, L64

Abraham J., et al., 2004, Nucl. Instrum. Meth. in Phys. Res. A, 523, 50

Abraham J., et al., 2007, Science, 318, 938

Abraham J., et al., 2008a, Astropart. Phys., 29, 188

Abraham J., et al., 2008b, Phys. Rev. Lett., 101, 061101

Abraham J., et al., 2009, preprint, (`arXiv:0906.2347`)

Abraham J., et al., 2010, Physics Letters B, 685, 239

Abreu P., et al., 2010, Astropart. Phys., 34, 314

Abreu P., et al., 2011, preprint, (`arXiv:1107.4809`)

Abreu P., et al., 2012, J. Cosmology & Astro-Part. Phys., 4, 40

Abreu P., et al., 2013, preprint, (`arXiv:1305.1576`)

Abu-Zayyad T., et al., 2012, Astrophys. J., 757, 26

Abu-Zayyad T., et al., 2013, Astroph. J., 768, L1

Achterberg A., Gallant Y. A., Kirk J. G., Guthmann A. W., 2001, MNRAS, 328, 393

Adams Jr. J. H., et al., 2013, preprint, (`arXiv:1307.7071`)

Ahlers M., Salvado J., 2011, Physical Review D, 84, 085019

Aitkin M., 1991, J. R. Statist. Soc. B, 53, 111

Allard D., Parizot E., Olinto A. V., Khan E., Goriely S., 2005, Astron. Astrophys., 443, L29

Aloisio R., 2008, Nuovo Cimento B Serie, 123, 915

Alvarez-Muñiz J., Engel R., Stanev T., 2002, Astrophys. J., 572, 185

Anderson C. D., 1933, Phys. Rev., 43, 491

Axford W. I., Leer E., Skadron G., 1977, International Cosmic Ray Conference, 11, 132

Balazs L. G., Meszaros A., Horvath I., 1998, Astronomy & Astrophysics, 339, 1

Barcikowski E., et al., 2013, in European Physical Journal Web of Conferences. p. 01006 (`arXiv:1306.4430`), doi:10.1051/epjconf/20135301006

Berezhko E., Knurenko S., Ksenofontov L., 2012, Astropart. Phys., 36, 31

Berlind A., Busca N., Farrar G. R., Roberts J. P., 2011, arXiv:1112.4188,

Bertone G., Isola C., Lemoine M., Sigl G., 2002, Physical Review D, 66, 103003

Bird D. J., et al., 1993, Phys. Rev. Lett., 71, 3401

Blasi P., Epstein R. I., Olinto A. V., 2000, Astrophys. J., 533, L123

Boyer J. H., et al., 2002, Nucl. Instrum. Meth. in Phys. Res. A, 482, 457

Center for Cosmology and Particle Physics 2013, Mock Catalogs for Ultrahigh Energy Cosmic Rays, Available at http://cosmo.nyu.edu/mockUHECR.html

Chiba N., Hayashida N., et al., 1992, Astropart. Phys., 1, 27

Clay J., Berlage H. P., 1932, Naturwissenschaften, 20, 687

Clyde M. A., Berger J. O., Bullard F., Ford E. B., Jefferys W. H., Luo R., Paulo R., Loredo T., 2007, in Babu G. J., Feigelson E. D., eds, Astronomical Society of the Pacific Conference Series Vol. 371, Statistical Challenges in Modern Astronomy IV. p. 224

Cox R. T., 1946, Am. Jour. Phys., 14, 1

De Domenico M., Insolia A., 2013, J. Phys. G, 40, 015201

Decerprit G., Allard D., 2011, Astron. Astrophys., 535, A66

Dickey J., 1971, Ann. Math. Stat., 42, 204

Dolag K., Grasso D., Springel V., Tkachev I., 2004, Journal of Korean Astronomical Society, 37, 427

Efstathiou G., 2008, MNRAS, 388, 1314

Egorova V., et al., 2004, Nucl. Phys. Proc. Suppl., 136, 3

Epele L. N., Roulet E., 1998, J. High Energy Phys., 10, 009

Farrar G. R., Jansson R., Feain I. J., Gaensler B. M., 2013, J. Cosmol. Astropart. Phys., 1, 023

Feigelson E. D., Babu G. J., eds, 2003, Statistical Challenges in Astronomy. Springer, New York

Fermi E., 1949, Physical Review, 75, 1169

Fraija N., et al., 2012, Astrophys. J., 753, 40

García-Gámez D., 2013, in European Physical Journal Web of Conferences. p. 04008, doi:10.1051/epjconf/20135304008

Gelman A., Robert C., Chopin N., Rousseau J., 1995, Bayesian Data Analysis

George M. R., Fabian A. C., Baumgartner W. H., Mushotzky R. F., Tueller J., 2008, MNRAS, 388, L59

Ghosh J., Delampady M., Samanta T., 2006, An Introduction To Bayesian Analysis: Theory And Methods. Springer, Berlin

Goulding A. D., Alexander D. M., Lehmer B. D., Mullaney J. R., 2010, MNRAS, 406, 597

Gregory P., 2005, Bayesian Logical Data Analysis for the Physical Sciences. Cambridge University Press, New York, NY, USA

Gregory P. C., 2010, Bayesian Logical Data Analysis for the Physical Sciences. Cambridge University Press, pp 376–388

Han J. L., 2008, Nucl. Phys. B (Proc. Suppl.), 175, 62

Han J. L., Manchester R. N., Qiao G. J., 1999, MNRAS, 306, 371

Harari D., Mollerach S., Roulet E., 1999, J. High Energy Phys., 8, 22

Harari D., Mollerach S., Roulet E., 2002, J. High Energy Phys., 7, 006

Hayashida N., Honda K., et al., 1996, Phys. Rev. Lett., 77, 1000

Hess V. F., 1912, Z. Phys., 13

Hillas A. M., 1984, Ann. Rev. Astron. Astrophys, 22, 425

Hillas A. M., 2006, Astrophysics preprint (http://arxiv.org/abs/astro-ph/0608074),

Hooper D., Sarkar S., Taylor A. M., 2008, Physical Review D, 77, 103007

Ivanov A. A., 2009, Nucl. Phys. B (Proc. Suppl.), 190, 204

Jansson R., Farrar G. R., Waelkens A. H., Enßlin T. A., 2009, Journal of Cosmology and Astro-Particle Physics, 7, 021

Jaynes E. T., 1976, in Harper W. L., Hooker C. A., eds, , Vol. 2, Foundations of Probability theory, Statistical Inference, and Statistical Theories of Science. D. Reidel Pub. Co., Dordrecht, pp 175–257

Jaynes E. T., 2003, Probability Theory: The Logic of Science. Cambridge University Press

Jenkins C. R., Peacock J. A., 2011, MNRAS, 413, 2895

Jui C. C. H., et al., 2012, JPCS, 404, 012037

Keilhauer B., 2006, in Magill S. R., Yoshida R., eds, American Institute of Physics Conference Series Vol. 867, Calorimetry in High Energy Physics: XII. pp 175–182 (`arXiv:astro-ph/0608074`), doi:10.1063/1.2396952

Keivani A., Farrar G. R., Sutherland M., 2015, Astropart. Phys., 61, 47

Kotera K., 2011, Physical Review D, 84, 023002

Kotera K., Olinto A. V., 2011, Ann. Rev. Astron. Astrophys, 49, 119

Letessier-Selvon A., Stanev T., 2011, Rev. Mod. Phys., 83, 907

Letessier-Selvon A., et al., 2014, Braz. J. Phys., 44, 560

Loredo T. J., 1992, in Feigelson E. D., Babu G. J., eds, Statistical Challenges in Modern Astronomy. Springer-Verlag, New York, pp 275–297

M. Hobson A. Jaffe A. L. P., ed. 2008, Bayesian Methods in Cosmology. Cambridge University Press, Cambridge, UK

Magliocchetti M., Ghirlanda G., Celotti A., 2003, MNRAS, 343, 255

Medina Tanco G. A., de Gouveia Dal Pino E. M., Horvath J. E., 1998, Astrophys. J., 492, 200

Mollerach S., et al., 2009, Nucl. Phys. B (Proc. Suppl.), 190, 198

Moskalenko I. V., Seo E., 2014, Adv. Space Res., 53, 1377

Murphy K. P., 2012, Machine Learning: A Probabilistic Perspective. MIT Press

Nagar N. M., Matulich J., 2010, Astron. Astrophys., 523, A49

Neddermeyer S. H., Anderson C. D., 1937, Phys. Rev., 51, 884

Nonaka T., et al., 2009, Nucl. Phys. Proc. Suppl., 190, 26

O'Hagan A., 1991, J. R. Statist. Soc. B, 53, 136

O'Hagan A., 1995, J. R. Statist. Soc. B, 57, 99

O'Hagan A., 1997, Test, 6, 101

Pacini D., 1912, Nuovo Cimento

Pe'Er A., Murase K., Mészáros P., 2009, Phys. Rev. D, 80, 123018

Press W., 2007, Numerical Recipes 3rd Edition: The Art of Scientific Computing. Cambridge University Press, `https://books.google.co.uk/books?id=1aA0dzK3FegC`

Puget J. L., Stecker F. W., Bredekamp J. H., 1976, Astrophys. J., 205, 638

Ryu D., Kang H., Cho J., Das S., 2008, Science, 320, 909

Shapiro S. L., Teukolsky S. A., 1983, Black Holes, White Dwarfs, and Neutron Stars: The Physics of Compact Objects, pp 335–369

Sigl G., Miniati F., Enßlin T. A., 2004, Physical Review D, 70, 043007

Sivia D. S., Skilling J., 2006, Data Analysis: A Bayesian Tutorial. Oxford science publications, Oxford University Press, Oxford, New York, `http://opac.inria.fr/record=b1133948`

Skilling J., 2006, Bayesian Anal., 1, 833

Soiaporn K., Chernoff D., Loredo T., Ruppert D., Wasserman I., 2013, Ann. Appl. Stat., 7, 1249

Spiegelhalter D. J., Smith A. F. M., 1982, J. R. Statist. Soc. B, 44, 377

Stanev T., 2004, Astrophysics e-prints (arxiv: http://arxiv.org/abs/astro-ph/0411113)

Stanev T., 2009, New J. Phys., 11, 13

Stecker F. W., Salamon M. H., 1999, Astrophys. J., 512, 521

Takeda M., Hayashida N., et al., 1998, Phys. Rev. Lett., 81, 1163

Takeda M., et al., 2003, Astropart. Phys., 19, 447

Tinyakov P. G., Tkachev I. I., 2001, J. Exp. Theor. Phys. Lett., 74, 445

Tinyakov P. G., Tkachev I. I., 2002, Astropart. Phys., 18, 165

Tinyakov P. G., Tkachev I. I., 2005, Astropart. Phys., 24, 32

Troitsky S. V., 2013, Physics Uspekhi, 56, 304

Unger M., 2008, International Cosmic Ray Conference, 4, 373

Véron-Cetty M.-P., Véron P., 2006, Astron. Astrophys., 455, 773

Watson L. J., Mortlock D. J., Jaffe A. H., 2011, MNRAS, 418, 206