

Impact of using the ultrahigh-energy cosmic ray arrival energies to constrain source associations

Francesca Capel^{1,2★} and Daniel J. Mortlock^{3,4,5}

¹The Oskar Klein Centre for Cosmoparticle Physics, SE-106 91 Stockholm, Sweden

²Department of Physics, KTH Royal Institute of Technology, AlbaNova, SE-106 91 Stockholm, Sweden

³Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK

⁴Statistics Section, Department of Mathematics, Imperial College London, London SW7 2AZ, UK

⁵Department of Astronomy, Stockholm University, AlbaNova, SE-106 91 Stockholm, Sweden

Accepted 2018 December 29. in original form 2018 November 15

ABSTRACT

We present a Bayesian hierarchical model which enables a joint fit of the ultrahigh-energy cosmic ray (UHECR) energy spectrum and arrival directions within the context of a physical model for the UHECR phenomenology. In this way, possible associations with astrophysical source populations can be assessed in a physically and statistically principled manner. The importance of including the UHECR energy data and detection effects is demonstrated through simulation studies, showing that the effective GZK horizon is significantly extended for typical reconstruction uncertainties. We also verify the ability of the model to fit and recover physical parameters from CRPROP 3 simulations. Finally, the model is used to assess the fraction of the publicly available data set of 231 UHECRs detected by the Pierre Auger Observatory which are associated with the *Fermi*-LAT 2FHL catalogue, a set of starburst galaxies, and *Swift*-BAT hard X-ray sources. We find association fractions of $9.5^{+2.4}_{-5.9}$, $22.7^{+6.6}_{-12.4}$, and $22.8^{+6.6}_{-8.0}$ per cent for the 2FHL, starburst galaxies, and *Swift*-BAT catalogues respectively.

Key words: methods: data analysis – methods: statistical – cosmic rays.

1 INTRODUCTION

Ultrahigh-energy cosmic rays (UHECRs) are the most energetic particles ever detected with energies of above $\sim 10^{18}$ eV (1 EeV) and their origin remains an open question (see Kotera & Olinto 2011; Aloisio 2017; Anchordoqui 2018 for recent reviews). UHECRs are extremely rare, with a flux of around 1 particle km^{-2} century⁻¹ above 50 EeV, but can be detected via air showers resulting from their interaction with the Earth's atmosphere. The development of large ground-based air shower detectors has allowed for measurements of the UHECR energy spectrum, arrival directions, and mass composition. The two largest experiments, the Pierre Auger Observatory (PAO, Aab et al. 2015a) and the Telescope Array Project (TA, Abu-Zayyad et al. 2012; Tokuno et al. 2012), have now detected several thousands of UHECRs.

Despite the increase in available data, the study of UHECRs remains challenging due to the complexity of the physical processes involved in their acceleration, propagation, and detection. A range of extra-Galactic astrophysical sources are postulated, including non-relativistic shocks in galaxy clusters (Norman, Melrose & A 1995; Ryu et al. 2003), relativistic shocks in active galactic nuclei

(AGNs, Biermann & Strittmatter 1987; Dermer et al. 2009), γ -ray bursts (Waxman 1995; Vietri 1995), and strong electric fields present in pulsars (Blasi, Epstein & Olinto 2000; Fang, Kotera & Olinto 2012). As the UHECRs propagate from their sources to Earth, they experience considerable deflections due to the effects of Galactic and extra-Galactic magnetic fields, as well as interactions with photons of the cosmic microwave background (CMB) and the extra-Galactic background light (EBL). The energy losses through interactions with the CMB give rise to the GZK effect (Greisen 1966; Zatsepin & Kuzmin 1966), implying that UHECRs detected at Earth with energies greater than 50 EeV must come from within a ~ 250 Mpc horizon.

As UHECRs are expected to come from relatively nearby, and their high energies mean that they should experience small deflections of less than $\sim 10^\circ$, many attempts have been made to find angular associations between UHECR arrival directions and potential sources (see e.g. Stanev et al. 1995; Tinyakov & Tkachev 2001; Alvarez-Muñiz, Engel & Stanev 2002; Abraham et al. 2007; Abreu et al. 2010; Kim & Kim 2011; Oikonomou et al. 2013; Aab et al. 2015b, and references therein). Most previous work has taken a hypothesis testing approach based on purely spatial correlations. A general consensus has not been reached, with differing conclusions on possible associations based on different UHECR data sets, source catalogues, and statistical methods. Most recently, the PAO

* E-mail: capel@kth.se

Collaboration reported a correlation between the arrival directions of UHECRs with arrival energies greater than 39 EeV and the positions of 23 starburst galaxies (SBGs) within 250 Mpc (Aab et al. 2018a, hereafter A18). The TA Collaboration followed up on this result by repeating the analysis on their UHECR data set, finding a result that is both consistent with isotropy and the SBG catalogue (Abbasi et al. 2018a).

While studying the UHECR arrival directions in isolation can provide some insight, it is crucial to consider the complementary information provided by the UHECR arrival energies in order to draw physical conclusions. We do not expect the highest energy UHECRs to come from distant sources or to be highly deflected. Within the context of a physical model for UHECR acceleration, propagation, and detection, the UHECR energies provide an additional constraint that avoids unphysical source-UHECR angular associations and allow us to assess whether a population of sources can realistically represent the observed data set. Including the energy data also allows us to directly constrain the underlying physical processes, such as the injection spectrum of the source population and the magnetic fields responsible for the energy-dependent deflections of UHECRs.

Building a comprehensive statistical model for the UHECR observations is non-trivial due to the high numbers of parameters and uncertainties involved, but can be achieved within the framework of Bayesian hierarchical modelling. Previous work has made headway in this direction, Watson et al. (2011) and Khanin & Mortlock (2016) developed a two-component parametric model characterized by source and background rates, modelling UHECR arrival directions as drawn from a binned Poisson intensity on the celestial sphere. Their model also accounts for UHECR energy losses during propagation in terms of weighting the sources, but does not include the UHECR energies explicitly in the likelihood. In Khanin & Mortlock (2016), the model is also applied to the public PAO data set of 69 UHECR events to a variety of catalogues, finding the largest association fractions of $0.25^{+0.09}_{-0.08}$ and $0.24^{+0.12}_{-0.10}$ for the *Swift*-BAT hard X-ray catalogue and the 2MASS Redshift Survey, respectively. Further, Soiaporn et al. (2013, hereafter S13) present an extendable hierarchical framework in which they derive an expression for the posterior distribution by modelling UHECR arrival directions as an inhomogeneous Poisson process on the celestial sphere. Their model was applied to the public data set of 69 UHECR events detected by PAO and a volume-limited sample of 17 nearby AGNs (Goulding et al. 2010), finding small but non-zero association fractions of up to ~ 0.2 .

Motivated by the desire to include the extra information provided by the UHECR energies, we present a joint model for the UHECR energies and arrival directions. This builds on the formalism of S13 and includes the ideas of Watson et al. (2011) and Khanin & Mortlock (2016). We derive the model in Sections 2 and 3, then validate the model through simulations in Section 4. Following this, we apply the model to the most recent public PAO data set. A range of source catalogues are tested against this data set and the results are presented in Section 5. The code used in this work is available online.¹

2 UHECR MODEL

For a given set of detected UHECRs with measured energies and arrival directions, a catalogue of potential sources and some




Figure 1. Directed acyclic graph showing the hierarchical model for UHECR energies and arrival directions. Data and parameters are shown using shaded and unshaded circles, respectively. The two boxes gather together the source and UHECR populations. Parameters outside of the boxes are global and are not associated with either population. The notation is explained in the text.

knowledge on the physics of UHECR acceleration and propagation, the statistical problem is to quantify the probability of association between the UHECRs and the sources. This can be done by breaking the problem down into different levels and building a hierarchy based on the underlying physical processes. We consider the case of K sources, indexed by k , and I UHECRs, indexed by i , as shown in Fig. 1. Following S13, the model has four levels: the source population (Section 2.1), and the UHECR acceleration, propagation, and detection (Sections 2.2–2.5). We also highlight some key implications of the model in more detail in Section 2.6

2.1 Source population

We assume a population of steady-state sources, such that the UHECR flux at Earth does not change over the time-scales of the available observations. In this way, the production rate from source k can be modelled as a homogeneous Poisson point process in time, characterized by its rate, L_k in units of time^{-1} . The k th source is defined by a direction on the celestial sphere, ϖ_k , and a distance from the Earth, D_k . We consider equal luminosities for all sources, such that $F_k = L/(4\pi D_k^2)$, for $k = 1$ to K .

We also include an isotropic background component to represent UHECRs from a population of distant, unresolved sources, or sources not present in the chosen catalogue. This component is labelled by $k = 0$ and is described by F_0 , the total flux due to this background population in units of $\text{area}^{-1} \text{time}^{-1}$.

2.2 Acceleration

The complex acceleration physics that produces UHECRs is summarized by modelling the emission spectrum of UHECRs with a power law. The number of UHECRs at a certain energy is described by $dN/dE \propto E^{-\alpha}$, where the spectral index α is expected to be ~ 2 from *Fermi* acceleration arguments (Bell 1978; Baring 2001). The

¹https://github.com/cescalara/uhecr_model (to be made available upon publication).

spectrum is normalized such that the emission rate per unit energy of UHECRs from source k is given by

$$\frac{dN_k}{dE dt} = L \frac{\alpha - 1}{E_{\min}} \left(\frac{E}{E_{\min}} \right)^{-\alpha}, \quad (1)$$

where $E_{\min}$ is the minimum UHECR emission energy and L is the rate at which source k emits UHECRs with $E > E_{\min}$. We use $E_{\min} = 52$ EeV when considering the PAO data set in Section 5 and fits to the simulated data in Section 4 as it is the minimum energy of the sample data. However, when generating simulated data sets in Section 4, we use $E_{\min} = 20$ EeV in order to take into account the effect of the uncertainty of the energy reconstruction, as described in Section 2.5.

2.3 Composition

The UHECR composition is challenging to reconstruct from air shower measurements, with the depth of the shower maximum, $X_{\max}$ and its fluctuations $\sigma(X_{\max})$ most often used to measure the hadronic component of the shower. The most recent observations are inconclusive, with TA favouring a lighter, proton composition (Abbasi et al. 2018b) and the PAO a mixed composition (Aab et al. 2016) at the highest energies. However, these results are actually consistent within the uncertainties present (Hanlon 2018).

We assume that all UHECRs are protons, allowing us to implement a relatively simple model for the UHECR propagation to demonstrate our framework. The model presented here could be extended to include a mixed UHECR composition as described in Section 7.

2.4 Propagation

During their propagation, UHECRs lose energy (Section 2.4.1) and are deflected by magnetic fields (Section 2.4.2).

2.4.1 Energy losses

We make use of the continuous energy loss approximation (Berezinsky et al. 1988; Chodorowski, Zdziarski & Sikora 1992; Anchordoqui et al. 1997; De Domenico & Insolia 2012) for energy losses during propagation. Under this approximation, the arrival energy of UHECRs can be calculated by solving

$$\frac{dE}{dz} = -\frac{E}{L_{\text{loss}}(E, z)}, \quad (2)$$

where z is the redshift of the UHECRs and $L_{\text{loss}}(E, z)$ is the loss length. There are three distinct processes which contribute significantly to the energy losses of protons for the energies considered here. Therefore, $L_{\text{loss}}(E, z)$ can be expressed as

$$L_{\text{loss}}^{-1}(E, z) = \frac{dt}{dz} [\beta_{\pi}(E, z) + \beta_{e^{\pm}}(E, z) + \beta_{\text{adi}}(z)], \quad (3)$$

where the β terms are the losses per unit time corresponding to the following processes:

- (i) photomeson production (e.g. $p + \gamma \rightarrow n + \pi^+$) due to high-energy protons interacting with the CMB (the GZK effect);
- (ii) Bethe–Heitler pair production ($p + \gamma \rightarrow p + e^+ + e^-$), also due to CMB interactions;
- (iii) the adiabatic expansion of the Universe.

Using the parametrization described in De Domenico & Insolia (2012), L_{loss} is as shown as a function of energy in Fig. 2. For




Figure 2. The total L_{loss} is shown as a function of energy for the cases of redshift $z = 0$ and 0.1 , corresponding to ~ 400 Mpc and thus covering the GZK horizon. The three different processes are also shown separately for $z = 0$ (cf. fig. 2 in De Domenico & Insolia 2012). The dotted line shows $E = 52 \times 10^{18}$ eV, the minimum energy of the PAO UHECR data set used in this work.

the typical energies considered here ($E > 52$ EeV), the dominant process is photomeson production. While continuous energy losses are a reasonable approximation for protons in this case, they do not account for spread in the resulting energy spectrum due to the kinematics of the interaction and the Poisson noise in the number of interactions (Hill & Schramm 1985; Achterberg et al. 1999).

For the energy loss calculations used throughout this work, we assume $z = DH_0/c$, where D is the luminosity distance in Mpc. This approximation is valid for the small redshifts ($z < 0.06$) considered. For all calculations involving cosmological parameters, we assume $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, and a standard Λ cold dark matter model with $\Omega_m = 0.3$ and $\Omega_\Lambda = 0.7$.

2.4.2 Magnetic deflections

UHECRs undergo deflections due to the various magnetic fields encountered during their propagation. The magnitude of these deflections is difficult to quantify, given the lack of knowledge about the strength and structure of the Galactic and extra-Galactic magnetic fields (see e.g. Vallée 2004; Kulsrud & Zweibel 2008 for an overview). Progress has been made in modelling the Galactic magnetic field (Jansson & Farrar 2012). The regular component of the Galactic magnetic field can act as a lens, leading to correlated deflections of UHECRs which are a function of the position on the sky (Waxman & Miralda-Escude 1996). Broadly speaking, the local Galactic magnetic field can be described with a regular component of roughly $3 \mu\text{G}$ and a Kolmogorov-type turbulent component of similar magnitude on length-scales of ~ 100 kpc (Durrer & Neronov 2013). Extra-Galactic magnetic fields are less constrained by observations and can vary over many orders of magnitude in both strength and coherence length, from smaller than $\sim \text{nG}$ in voids over 1 Mpc scales to up to $\sim \mu\text{G}$ in 100 kpc scale galaxy clusters (Kotera & Lemoine 2008; Aharonian, Kelner & Prosekin 2010; Durrer & Neronov 2013).

Following Harari et al. (2002), we model the total magnetic field as a random Gaussian field with zero mean and root mean square field strength $\bar{B}$ and do not separate the Galactic and extra-

Galactic contributions. In the limit of small deflections and the distance travelled being much larger than the coherence length ($D \gg l_c$), the deflections of UHECRs follow a probability distribution characterized by zero mean and root mean square angular scale, given by

$$\bar{\theta} \approx 2.3^\circ \left(\frac{E}{50 \text{ EeV}} \right)^{-1} \left(\frac{\bar{B}}{1 \text{ nG}} \right) \left(\frac{D}{10 \text{ Mpc}} \right)^{1/2} \left(\frac{l_c}{1 \text{ Mpc}} \right)^{1/2}, \quad (4)$$

where the quantities have been scaled to typical extra-Galactic values and a pure proton composition is assumed. Given the uncertainties in modelling the magnetic fields, we leave $\bar{B}$ as a free parameter in the model, but choose a fixed $l_c = 1 \text{ Mpc}$ to avoid degeneracy and represent extra-Galactic scales.

2.5 Detection

The exposure of a ground-based UHECR observatory such as PAO or TA can be calculated by considering the geometry of the system. For an observatory at latitude a_0 that can detect UHECRs arriving with zenith angles up to θ_m , the exposure as a function of position on the sky, $\epsilon(\omega)$ is given by

$$\epsilon(\omega) = \frac{\alpha_T}{M} m(\omega), \quad (5)$$

where ω is a unit vector describing the UHECR arrival direction and α_T is the total exposure of the observatory in units of area $\times$ time $\times$ solid angle and $M = \int d\omega m(\omega)$. The arrival direction, ω can also be expressed in terms of right ascension and declination. The time-averaged exposure over 1 sidereal day is purely a function of declination, δ , and is given by

$$m(\omega) = m(\delta) = \cos(a_0) \cos(\delta) \sin(\alpha_m) + \alpha_m \sin(a_0) \sin(\delta), \quad (6)$$

where α_m and the θ_m dependence are as described in Sommers (2001). $\epsilon(\omega)$ is proportional to the probability that a UHECR is detected given that it has arrival direction ω . For the PAO, we have $a_0 = 35.2^\circ$, $\theta_m = 80^\circ$, and $\alpha_T = 66,000 \text{ km}^2 \text{ sr yr}$, as described in Aab et al. (2015b).

In addition to the exposure effects, the reconstruction uncertainties in the measured UHECR energies and arrival directions are also important. The uncertainty in the energy reconstruction is reported as a fraction of the arrival energy of the UHECR, $\sigma_E = f_E E$ where f_E is between 0 and 1 (Aab et al. 2015a). We model the detected energies as being Gaussian distributed

$$P(\hat{E}|E, \sigma_E) \sim \mathcal{N}(E, f_E^2 E^2), \quad (7)$$

where $\hat{E}$ is the detected energy and E is the true arrival energy. Similarly, the uncertainty in the reconstruction of the arrival direction is modelled using a von Mises-Fisher (vMF, Fisher 1953) distribution centred on the true arrival direction and with the spread characterized by the dimensionless parameter κ_d as

$$P(\hat{\omega}|\omega, \kappa_d) = \frac{\kappa_d}{4\pi \sinh(\kappa_d)} e^{\kappa_d \hat{\omega} \cdot \omega}, \quad (8)$$

where $\hat{\omega}$ is the detected arrival direction, ω is the true arrival direction, and κ_d is inversely related to the width of the distribution. If $\kappa_d \rightarrow 0$, the vMF distribution becomes uniform on the sphere; if $\kappa_d \gg 1$, the vMF distribution approaches a bivariate Gaussian distribution near the mode. As the angular reconstruction uncertainty, σ_ω , is often reported in units of degrees, we can convert this to κ_d in analogy with the ‘1 σ ’ region of a Gaussian distribution. As highlighted in S13, we have

$$\kappa_d \approx 7552 \left(\frac{\sigma_\omega}{1^\circ} \right)^{-2}. \quad (9)$$

The reconstruction uncertainties of the PAO are reported to be $f_E \leq 0.12$ for the energies (Aab et al. 2015a,b) and $\sigma_\omega \leq 0.9^\circ$ for the arrival directions (Bonifazi 2009). As we do not have heteroscedastic event-by-event uncertainties, these maximum values are adopted in our analysis. We do not consider the additional systematic uncertainty of ~ 14 per cent on the absolute energy scale.

2.6 Impact of detection uncertainties

The number of UHECRs detected per unit area per unit time from a particular source, N_k , is proportional to the rate of UHECR emitted by that source, L , but must also take into account the source energy spectrum and the effects of UHECR energy losses. The energy loss model described in Section 2.4.1 can be used to calculate the corresponding energy of a UHECR at the source, given its arrival energy. In this way, we can define a threshold source energy, $\tilde{E}_{th}$ that corresponds to a threshold arrival energy at Earth, E_{th} for a source at a given distance, D_k . Along with equation (1), this allows us to define the number of UHECR per unit time and area as

$$\frac{dN_k(E > E_{th})}{dr dA} = \frac{L}{4\pi D_k^2} \left(\frac{\tilde{E}_{th}^k}{E_{min}} \right)^{1-\alpha}. \quad (10)$$

In order to apply equation (10) to a sample that has been selected according to the detected UHECR energies, as will be necessary when working with the available UHECR data, we must assume that a hard threshold on the detected energies corresponds to a hard threshold on the source energies. This is not true, due to the energy detection uncertainties (equation 7). We approximate these effects by truncating the Gaussian distribution in equation (7) at E_{th} , which allows for the use of equation (10). The validity of the approximation is determined by the ability of the model to fit data simulated with realistic detection effects, as discussed in Section 4. Elsewhere in the model, we do not ‘invert’ the detected UHECR energies in order to infer their energies at the source, but introduce latent initial energy parameters, $\tilde{E}_i$, to fit for the UHECR source energies, including the detection effects.

The uncertainty on the arrival energies means that some events in a UHECR sample will be lower energy events which are reconstructed with energies above the threshold, resulting in an effective extension of the GZK horizon. To quantify the impact of this, we consider the probability that a UHECR comes from within a distance D , given that it is detected with $\hat{E} > E_{th}$. This can be expressed as

$$\begin{aligned} P(< D | \hat{E} > E_{th}) &= \int_0^D dD' P(D' | \hat{E} > E_{th}) \\ &= \int_0^D dD' P(D') P(\hat{E} > E_{th} | D'), \end{aligned} \quad (11)$$

where $P(\hat{E} > E_{th} | D)$ can be estimated through simulations and $P(D) \propto D^2$ for the case of homogeneously distributed sources in the nearby universe. We evaluate $P(< D | \hat{E} > E_{th})$, both with and without detection uncertainties and the results are shown in Fig. 3. For an uncertainty of 12 per cent on the reconstructed energies and $E_{th} = 52 \text{ EeV}$, we expect around 30 per cent of UHECR in the sample to have come from beyond 250 Mpc. Therefore, if one only considers sources within 250 Mpc, we expect to find a background contribution of 30 per cent to the observed UHECR flux due to the contribution of sources beyond the GZK horizon, which would be negligible if assuming that we observe E instead of $\hat{E}$.




Figure 3. The probability that a UHECR is detected from within a distance D , given that it is detected with an energy $\hat{E} > E_{\text{th}}$. The solid and dashed lines show the results with and without detection uncertainties respectively for the case of $E_{\text{th}} = 52$ and $\alpha = 2$. The grey dotted line shows $P(< D | \hat{E} > E_{\text{th}}) = 0.9$, a value which can be used to define the GZK horizon (Harari, Mollerach & Roulet 2006).

To summarize the implications of the model for potential UHECR sources, we also consider the joint probability of the UHECR deflections and the source distances, conditioned on different detected energies. To make this easier to visualize, we express the result in terms of Cartesian coordinates with the z -axis representing the arrival direction of a UHECR and the xy -plane the surface on to which it arrives. Using the same assumptions on the source population as mentioned above, this can be expressed as

$$P(x, y = 0, z | \hat{E}, \hat{\omega} = 0, S) = \frac{P(D, \theta, \phi, \hat{E} | \hat{E}, \hat{\omega} = 0, S)}{D^2 \sin(\theta)} \quad (12)$$

$$= \frac{P(D, \theta, | \hat{E}, \hat{\omega} = 0, S)}{D^2 \sin(\theta)},$$

where S signifies that the UHECR is detected at Earth, $\hat{\omega}$ are the detected UHECR arrival directions, θ are the deflections of the detected UHECR from the source location, and D and ϕ make up the spherical coordinate system. In the last step, we have marginalized over ϕ and $\hat{E}$. If we ignore the detection uncertainties, thus conditioning on (E, ω) instead of $(\hat{E}, \hat{\omega})$, this can be expressed analytically as

$$P(x, y = 0, z | E, \omega = 0, S) \propto D^{-3} \tilde{E}(E, D)^{2-\alpha} \frac{e^{\theta_0^{-2} \tilde{E}(E, D)^2 D^{-1} \cos(\theta)}}{\sinh[\theta_0^{-2} \tilde{E}(E, D)^2]} \left| \frac{dE(\tilde{E}', D)}{d\tilde{E}'} \right|_{\tilde{E}' = \tilde{E}(E, D)}, \quad (13)$$

where θ_0 is the constant part of the $\bar{\theta}$ dependence given in equation (4) and we have used the vMF distribution to represent the UHECR deflections with $\bar{B} = 1$ nG. We plot this distribution in the top panel of Fig. 4 for the case of $E = 50$ and 70 EeV. We see that higher energy UHECRs must originate from closer sources in addition to having smaller deflections. In this way, the model reflects that there is a particle-specific GZK horizon for each UHECR that is smaller for higher energy particles. The contours show two peaks in the distribution, indicating that UHECRs are most likely to come from either $D = 0$ or towards the location of the effective GZK horizon for the energy considered. This can be understood by the fact that due to the energy losses under the




Figure 4. The joint probability that a UHECR originates from (x, z) given that it is detected with arrival energy E , with rotational symmetry around the z -axis. The top panel shows the result for the expression in equation (13) for the case of $\bar{B} = 1$ nG and $\alpha = 2$. The lower panel shows the same probability evaluated by using simulations and selecting arrival energies in the ranges of $[50, 55]$ and $[70, 75]$ EeV.

continuous loss approximation, protons produced at the horizon will ‘pile up’ at the arrival energy, whereas for nearby sources, the source spectrum is equal to the arrival spectrum and so follows the power-law form. Additionally, we expect a larger contribution from near the effective horizon. This effect is less prominent at higher energies ($E \gg E_{\text{GZK}}$), as the energy losses become effectively exponential and the Jacobian term in equation (13) tends to unity.

We evaluate $P(x, y = 0, z | \hat{E}, \hat{\omega} = 0, S)$ using simulations both with and without detection uncertainties on the UHECR energies and arrival directions. For the case without detection uncertainties, we compare with result with the calculation described above in Fig. 4 and see that we find a good match between the two, with some small differences to be expected due to the selection of a range of arrival energies in the simulation. We then show the comparison with and without detection uncertainties in Fig. 5. These distributions represent a slice through the 3D ‘bubble’ of possible UHECR origins given that it is detected with energy $E/\hat{E}$ and direction $\omega/\hat{\omega}$ (aligned with the z -axis). Again, we note the importance of including the detection effects, with the distribution of possible UHECR origins extending further in both x and z for detection uncertainties of $\sigma_\omega = 0.9^\circ$ and $f_E = 0.12$, which correspond to the case of the PAO data set used in this work. Increasing the horizon distance from, for example, 250 to 360 Mpc more than doubles the volume and hence the number of potential sources that should be considered.

3 STATISTICAL FORMALISM

Using the above UHECR model, we derive an expression for the likelihood, which gives the probability of the observed data given a set of model parameters (Section 3.1). For clarity, we highlight the main points here; the full derivation can be found in Appendix A. We then show how the model parameters can be used to obtain an expression for the associated fraction of UHECRs in Section 3.2. We describe the prior distributions in Section 3.3 before bringing this together with the likelihood in Section 3.4 to obtain an expression




Figure 5. The joint probability that a UHECR originates from (x, z) given that it is detected with energy E or $\hat{E}$, evaluated using simulations. The distribution has rotational symmetry around the z -axis. As in Fig. 3, the dashed and solid lines show the results with and without detection uncertainties respectively for the case of $\bar{B} = 1$ nG and $\alpha = 2$. The different panels from top to bottom show the case for increasing detected energy. At higher energies, the double peak visible in the distribution disappears, as discussed in the text. Note that both the x - and z -axes ranges vary in order to show the structure of the distributions.

for the posterior distribution which can be evaluated numerically. Finally, we also derive a simpler model for only the UHECR arrival directions in Section 3.5 for comparison.

3.1 Likelihood

Whilst the production of UHECRs can be modelled as constant in time, their detection cannot, due to the rotation of the Earth and the corresponding time dependence of the instantaneous exposure of an observatory. This gives a detection rate which is a function of both UHECR arrival directions and time, $r(\omega, t)$. The resulting likelihood function is that of an inhomogeneous Poisson point process (Loredo 2004; Streit 2014) and has the form

$$P(\hat{E}, \hat{\omega} | L, \alpha, \bar{E}, \bar{B}) \propto e^{-\bar{N}} \prod_{i=0}^I r(\omega_i, t_i), \quad (14)$$

where $\bar{N} = \int \int d\omega dt r(\omega, t)$ is the total expected number of UHECRs in the observation period. The likelihood which we derive here has the same overall structure, but with modifications due to the fact that we model the UHECRs as being produced with labels, λ_i , which identify their sources, and initial energies, $\tilde{E}_i$. This introduces associated distributions for the labels and energies such that UHECR production can be described by a marked Poisson point process.

In order to write down an expression for the Poisson rate, it is necessary to consider contributions to the likelihood from all possible source-UHECR combinations. This can be achieved by marginalizing over the latent UHECR labels, resulting in what is effectively a mixture model over different source contributions, including the background component (see Section 2.1). Building on equation (14), we have

$$P(\hat{E}, \hat{\omega} | L, \alpha, \bar{E}, \bar{B}) = e^{-\bar{N}} \prod_{i=0}^I \sum_{k=0}^K \left[F_k P(\hat{\omega}_i | \varpi_k, D_k, \tilde{E}_i, \bar{B}) P(\hat{E}_i | E_i) P(E_i | \tilde{E}_i, D_k) P(\tilde{E}_i | \alpha) \right]. \quad (15)$$

Here, the F_k act as weights, with $P(\hat{\omega}_i | \varpi_k, D_k, \tilde{E}_i, \bar{B})$ quantifying the probability of observing a UHECR with arrival direction $\hat{\omega}_i$, given that it comes from source k , with initial energy $\tilde{E}_i$ and travels through a magnetic field of mean strength $\bar{B}$. The latent ω_i have been marginalized over. The other terms in the likelihood, $P(\hat{E}_i | E_i)$, $P(E_i | \tilde{E}_i, D_k)$, and $P(\tilde{E}_i | \alpha)$ give the probability of observing a UHECR with energy $\hat{E}_i$ given that the production spectrum is a power law with spectral index α and that the energy losses during propagation are as described in Section 2.4.1.

The background component with $k = 0$ assumes no energy losses (equivalent to $D_0 = 0$ Mpc). This means that if a low associated fraction is found, the model essentially fits an isotropic distribution of arrival directions and α represents the arrival energy spectrum and not that of the sources.

To calculate the expected number of events, $\bar{N}$, we must sum over the contribution of all possible sources, whilst taking into account the effect of the GZK horizon and the exposure of the detector. Energy losses are considered in the derivation of equation (10), which defines an expression for the flux of UHECRs from each source. We continue by defining an effective exposure factor for each source, ε_k , which is the convolution of the source vMF distribution and the observatory exposure. This depends on both the direction of the source on the celestial sphere and the spread of the vMF distribution associated with that source. The effective exposure is hence indirectly dependent on the source UHECR energy distribution, the source distance and the magnetic field strength. This gives

$$\bar{N} = \sum_{k=0}^K \frac{dN_k(E > E_{th}^k)}{dt dA} \varepsilon_k(D_k, \alpha, \tilde{E}_{th}^k, \bar{B}), \quad (16)$$

which can be then substituted into equation (15) to obtain the full expression of the likelihood.

3.2 The associated fraction

The fraction of UHECRs associated with the source catalogue, f , can be calculated in this framework as a derived parameter:

$$f = \frac{\sum_{k=1}^K F_k}{F_0 + \sum_{k=1}^K F_k}. \quad (17)$$

This is the expected association fraction, given the observations, and not the exact fraction of associated UHECRs in the sample. We

emphasize that in addition to this summary parameter, it is also possible to calculate the explicit probability for each source-UHECR association (or equivalently, the marginal posterior distribution for the UHECR labels λ_i), as detailed in Appendix B. By estimating f , we perform a form of embedded model comparison between the isotropic and source models within the framework of parameter estimation (Kamary et al. 2014).

3.3 Priors

All parameters discussed in this section and shown in Fig. 1 are left free when fitting the model to data. The so-called hyperparameters, or highest level parameters, α , $\bar{B}$, L , and F_0 , have associated hyperpriors. Motivated by the desire to include our knowledge of UHECR phenomenology into the analysis and avoid giving weight to unphysical regions of the parameter space, we choose weakly informative hyperpriors for the these parameters (Gelman 2006; Simpson et al. 2017; Gelman, Simpson & Betancourt 2017). Wide Gaussian priors on α and $\bar{B}$ reflect the lack of knowledge on these parameters and similarly for the case of L and F_0 , they reflect that we do not expect these parameters to tend to infinity. We have

$$\begin{aligned} P(\alpha) &\sim \mathcal{N}(\mu_\alpha, \sigma_\alpha), \\ P(\bar{B}) &\sim \mathcal{N}(\mu_{\bar{B}}, \sigma_{\bar{B}}), \\ P(L) &\sim \mathcal{N}(0, \sigma_L/K), \\ P(F_0) &\sim \mathcal{N}(0, \sigma_F), \end{aligned} \quad (18)$$

where $\mu_\alpha = 2$, $\sigma_\alpha = 3$, $\mu_{\bar{B}} = 50$ nG, $\sigma_{\bar{B}} = 50$ nG, $\sigma_L = 10^{43}$ yr⁻¹, and $\sigma_F = 0.01$ km⁻²yr⁻¹. The hyperparameters have additional lower bounds: α at 1, L and F_0 at 0, and $\bar{B}$ between 0 and 100 nG. The upper limit on $\bar{B}$ reflects that we do not expect such large extra-Galactic magnetic fields that would result in even nearby sources having almost isotropic arrival direction distributions. The parameters are never found on their boundaries in any of the model fits presented in this work, and a larger upper bound on $\bar{B}$ has no effect on the results.

For the analyses performed in Sections 4 and 5, a sensitivity analysis was performed to ensure that the main conclusions were robust to the choice of hyperprior distributions, under a range of plausible alternatives.

3.4 Inference

We perform inference on the model parameters for a given data set within the framework of Bayesian probability theory. By combining the likelihood (equations 15 and 16) and the priors (equation 18), we obtain an expression for the joint posterior distribution of the model parameters given the data

$$P(L, \alpha, \bar{B} | \hat{E}, \hat{\omega}) \propto P(\hat{E}, \hat{\omega} | L, \alpha, \bar{B}) P(L, \alpha, \bar{B}) \quad (19)$$

where $P(\hat{E}, \hat{\omega} | L, \alpha, \bar{B})$ is the likelihood, given in equation (15) and $P(L, \alpha, \bar{B})$ is the joint prior distribution of the priors described in Section 3.3.

We compute the posterior distribution numerically by generating samples from it using STAN² (Carpenter et al. 2017). STAN is a software based on an implementation of adaptive Hamiltonian Monte Carlo, as outlined in Betancourt (2017). We use a set of diagnostics to assess the convergence of STAN to the target posterior distribution. These include the number of effective (or uncorrelated)

samples, n_{eff} , and the Gelman–Rubin statistic, $\hat{R}$, to monitor the mixing of separate chains (Gelman & Rubin 1992). We require $n_{\text{eff}} > 1000$ and $\hat{R} < 1.1$ with no divergent transitions for all model parameters in the analyses presented below.

We also perform model checking to evaluate the fit of the model to data using posterior predictive checks (PPCs, Gelman, Meng & Stern 1996; Sinharay & Stern 2003; Lynch & Western 2016). This involves generating data under the assumptions of the fitted model by drawing samples from the joint posterior predictive distribution. The generated data are then compared with the observed data to verify that the inferences of the model are sensible. The method and checks are detailed in Appendix C.

3.5 Arrival directions only

In order to test the significance of including the UHECR energy data into the model and to facilitate comparison with previous work, we also develop a model for only the UHECR arrival directions. This follows the model described in S13, but with some key improvements. We use the above parametrization, but as there is no energy model in this case we simply infer the width of the vMF distribution representing the UHECR deflections, κ , as a parameter instead of $\bar{E}_i$, α , and $\bar{B}$ (see equation A4). Thanks to the implementation in STAN, we are also able to infer κ directly instead of conditioning upon it as in S13. As we do not use a Gibbs sampling approach, the choice of priors for L , F_0 , and κ are not constrained by the need to form closed-form conditional distributions for the Gibbs steps which can be directly sampled from. This makes it much simpler to implement the model and perform sensitivity analyses. Finally, the discrete latent label parameters, λ_i , are marginalized over. This allows for better exploration of the tails of the posterior distribution and more efficient sampling, as detailed in Stan Development Team (2018). The likelihood of this reduced model can be written (cf. S13, equation 19)

$$P(\hat{\omega} | L, \kappa) = \exp \left[- \sum_{k=0}^K F_k \epsilon_k(\kappa) \right] \prod_{i=0}^I \sum_{k=0}^K F_k P(\hat{\omega}_i | \varpi_k, \kappa). \quad (20)$$

4 SIMULATIONS

We test our model on simulated data sets in order to verify its ability to recover the input parameters. To do this, we use the catalogue of 23 SBGs within 250 Mpc chosen in A18 to define the source locations (see Section 5.2). A large number of UHECRs are simulated according to the assumptions of the model with initial energies from 20 EeV, chosen as UHECRs with arrival energy below 20 EeV have negligible probability of being detected above 52 EeV, given $f_E = 0.12$. A threshold is then applied at $E_{\text{th}} = 52$ EeV to mimic the publicly available data, yielding 249 UHECRs as shown in Fig. 6. The source spectrum is defined by $\alpha = 3$, the root mean square magnetic field strength is taken to be $\bar{B} = 20$ nG and the associated fraction of the simulated sample is set to $f = 0.5$. The observatory exposure and the detection uncertainties are modelled as described in Section 2.5.

We fit two models to the simulated sample: a model for the UHECR arrival directions only (as described in Section 3.5); and the full joint model for both UHECR energies and arrival directions. The resulting marginal posterior distributions for f are shown in Fig. 7. We see that by ignoring the information in the UHECR energy data, we are not able to properly recover the input $f = 0.5$. The arrival direction model underestimates f and is more constrained around this false value, with a narrower posterior distribution. In contrast,

²<http://mc-stan.org>




Figure 6. The simulated data set from the generative model, shown in Galactic coordinates using the Hammer–Aitoff projection. The coloured tissots represent UHECR energies, with larger tissots corresponding to higher energy, and the grey scale shows the time-averaged exposure of the PAO. The locations of sources in the SBG catalogue are also shown by solid black tissots.




Figure 7. The marginal posterior distribution of f for the case of the arrival direction only model and the full joint model. The input value of $f = 0.5$ is also shown by the black line for comparison.

the joint model correctly recovers $f = 0.5$ and the wider posterior reflects the true uncertainty in this result due to the inclusion of the UHECR energy model.

It should also be noted that by not providing a physical model for the UHECR energies, we are not able to include, and thus fit for, physically interesting parameters such as α and $\bar{B}$. In this way, including the UHECR energies into the model is both more accurate and more informative. The joint model is able to correctly infer all input hyperparameters, as demonstrated in Fig. 8. This result was verified for a broad range of relevant input hyperparameters.

4.1 CRPROPA comparison

As we recognize that our model is an approximation to the underlying physical processes involved, we also test it with the results of a Monte Carlo simulation using the publicly available CRPROPA 3 code (Batista et al. 2016). We use this tool to simulate the energy losses and deflections of protons more realistically, but otherwise set-up an identical simulation to that described above. We include all available energy loss mechanisms due to interactions of protons

with both the CMB and the EBL, with photomeson production simulated using the SOPHIA code (Mücke et al. 2000). The magnetic field is implemented as a turbulent field with a Kolmogorov power spectrum on length-scales of 1–3.135 Mpc, giving $l_c = 1$ Mpc to match the assumptions described in Section 2.4.2. We define $\bar{B} = 20$ nG on a 3D grid with 100 kpc spacing and 500 points in each dimension, resulting in a cube of side 50 Mpc which is repeated in space to cover the required volume. UHECRs are simulated from and tracked down to a minimum energy of 20 EeV. Again, 20 EeV is chosen as with an energy detection uncertainty of $f_E = 0.12$, UHECRs with arrival energy below 20 EeV have negligible probability of being detected above 52 EeV and thus affecting the final sample. We set-up an observer sphere at the Earth with a radius of 1 Mpc (in its current implementation, CRPROPA 3 requires a large observer radius in order to recover large numbers of events with reasonable computing resources). In order to offset the effect of shorter trajectories on the magnetic deflections of protons in the simulation, we also shift all sources to be 1 Mpc further away, although the impact of this on the end result is minimal.

Despite these differences, using our model to fit the CRPROPA 3 data set we find that the marginal posterior distributions for α and B are consistent with the input values, as shown in Fig. 8. This demonstrates that despite its simplicity, the model used captures the essence of the underlying physics and as such is a justifiable approximation.

5 APPLICATION TO UHECR DATA

We apply the model to the publicly available data of the 231 highest energy cosmic rays detected at the PAO, as documented in Aab et al. (2015b). The observatory covers an area of 3000 km² and consists of 1660 surface detectors, with a duty cycle of ~ 100 per cent, and 24 fluorescence detectors which are limited by the need to operate in clear night conditions, resulting in a duty cycle of ~ 13 per cent.

As in Section 4, we compare the case of the arrival direction only model presented in Section 3.5 and the full joint model for UHECR energies and arrival directions as presented in Section 3.




Figure 8. The joint marginal posterior distributions for the model hyperparameters from fits to the simulated data sets. The left-hand plot shows that all input hyperparameters are inferred correctly for the case of the fit to data simulated from the generative model. The right-hand plot shows the key physical parameters inferred from a fit to the CRPROPA 3 simulated data set. In both cases, the shaded contours show the 30, 60, 90, and 99 per cent highest posterior density credible regions and the solid lines show the input values for the simulation.

5.1 UHECR data

The available data were collected between 2004 January 1 and 2014 March 1, covering a ~ 10 yr period. The sample contains 231 UHECR above a threshold energy of $E_{\text{th}} = 52$ EeV. At these energies, the trigger and selection efficiency is assumed to be 100 per cent, so the exposure of the observatory is purely a function of the detector geometry and the Earth's rotation and can be modelled as described in Section 2.5. We also take into account the relative exposures of different observation periods, due to how the observatory has grown in size over the years, as described in Abreu et al. (2010). Additionally, the different exposure of PAO to ‘vertical’ ($\theta_i < 60^\circ$) and ‘inclined’ ($60^\circ \leq \theta_i < 80^\circ$) samples is also accounted for. These exposure effects come into the likelihood in the form of different A_i and θ_i , as shown in equations (A5) and (15).

5.2 Source catalogues

The model presented is partly defined in terms of a catalogue of potential UHECR sources, thus the choice of sources is central to the results of the fit and the eventual conclusions of such an analysis. In this work, we follow the source catalogues used in A18, and also by several other authors, in order to allow comparison between the different statistical approaches used. These catalogues were chosen to represent nearby AGNs and SBG populations, motivated by the observation that the rate of energy production of UHECRs is $\sim 10^{45} \text{ erg Mpc}^3 \text{ yr}^{-1}$, which is similar to that observed in the γ -ray band by *Fermi* for AGNs and SBGs (Dermer & Razzaque 2010).

For completeness, we summarize the chosen catalogues here. We consider the *Fermi*-LAT 2FHL catalogue of hard γ -ray sources (Ackermann et al. 2016), and select only radio-loud sources within a 250 Mpc radius, resulting in 17 AGNs (blazars and radio galaxies). For SBGs, we take the 23 SBG which have radio fluxes > 0.3 Jy and are within 250 Mpc from the list of SBGs considered in a *Fermi*-LAT search, reported in Ackermann et al. (2012). We also consider a flux-limited sample, the *Swift*-BAT hard X-ray sources with fluxes above $13.4 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$ (Oh et al. 2018). These catalogues are shown along with the UHECR data in Fig. 9. Unlike A18, we do not consider the 2MRS catalogue of $\sim 40\,000$ galaxies within 250 Mpc due to the large computational resources required. Under the assumptions of our model, we would expect a higher association fraction due to the sheer number of sources present. In this way, a meaningful comparison with the other catalogues in our analysis would be difficult.

Several issues have been pointed out with the specification of these potential sources. Matthews et al. (2018) highlight that the nearby radio galaxy Fornax A does not appear in the 2FHL catalogue as it is an extended source, and could reasonably provide an explanation for one of the hotspots mentioned in the A18 analysis. Abbasi et al. (2018a) note that the exclusion of Local Group Objects (SMC, LMC, M33, and M31) is contradictory to the logic of weighting the SBG sources according to their radio fluxes, as in this case these sources would dominate the flux maps generated in A18 due to their close proximity. Considerations of statistical completeness and the selection effects of candidate source catalogues also present challenges. The specification of a physically motivated catalogue of potential UHECR sources is a non-trivial exercise and should be approached with caution in future work.




Figure 9. The source catalogues and UHECR data used in the analysis are shown in Galactic coordinates using the Hammer–Aitoff projection. The coloured tissots represent UHECR energies, with larger tissots corresponding to higher energy, and the grey solid line shows the exposure limit of the PAO. The source catalogues are also shown by solid black tissots.

which attempts similar analyses. Despite the issues discussed, we proceed with the above source catalogues in the interest of a more direct comparison with previous work.

A18 also weight the sources according to their flux in γ -rays, radio and X-rays for the 2FHL, SBG, and *Swift*-BAT catalogues, respectively. The details of this assumed correlation are highly model dependent, with the exact relationship between γ -ray and UHECR flux dependent on the relevant spectra of particles and radiation in the source environment, or indeed during propagation, the UHECR composition and extra-Galactic magnetic fields (Murase et al. 2012; Böttcher et al. 2013; Fang & Murase 2018). Given that the details of this are still an open question, we instead chose the simple case of equal luminosities for the sources in our analysis.

6 RESULTS

The results of the comparison between the arrival direction and joint models are shown in Fig. 10 in the form of the marginalized posterior distributions for f , and the corresponding 68 per cent regions of highest posterior density are given in Table 1. The fits of the arrival direction model find lower associated fractions compared to those of the joint model, and that $P(f|\hat{E}, \hat{\omega})$ is broader for the joint model. This is similar to what was seen for the comparison of the two models using the simulated data set, as discussed in Section 4 and shown in Fig. 7. It seems by including the energy information in the context of this model, we are able to uncover stronger associations between the candidate source catalogues and the UHECRs. We interpret the broader marginal posterior distributions as a more realistic representation of the uncertainty in the results, due to the inclusion of a model for the UHECR energy losses and reconstruction. The largest associated fractions of ~ 23 per cent are found for the case of the SBG and *Swift*-BAT catalogues. The fact that we find roughly the same association fraction for both the SBG and *Swift*-BAT catalogues demonstrates that very different catalogues can produce similar results under the assumptions of our model. The SBG catalogue is made up of 23 predominantly nearby sources, whilst the *Swift*-BAT catalogue contains 213 sources, most of which are more distant. By allowing all model parameters to be free in the fit, we see that different values of the magnetic field strength allow for the density of sources in each case to explain a reasonable fraction of the UHECR observations. Both catalogues trace the local matter distribution in different ways.

Our results are broadly consistent with the most recent results of the PAO for single catalogues against isotropy (Section 4.2 in A18), which find associated fractions of 7 ± 4 , 10 ± 4 , and 7^{+4}_{-3} per cent for the 2FHL, SBG, and *Swift*-BAT catalogues, respectively, with additional systematic errors of ~ 0.3 per cent. We see that the results are generally closer to those of the arrival direction model, with our joint model favouring higher associated fractions in each case. Additionally, we tend to have larger uncertainties on our results which are naturally represented by the width of the marginal posterior distribution. It is important to note that larger uncertainties do not necessarily mean that our method is less constraining, but rather reflects that it consistently takes into account the uncertainties at each level of the model hierarchy. By reducing the number of assumptions we make with regards to the model in order to make it more physical, we naturally introduce uncertainties into the problem. For example, by fixing α or $\bar{B}$, we obtain narrower marginal posterior distributions for f , but the apparent constraining power is simply a result of making stronger assumptions. As we are

able to infer physical parameters, there is no need to fix or scan over any of the parameters in the model.

A direct comparison between our results and the work of A18 is not possible for a number of reasons. First, the PAO collaboration have access to the full, more recent UHECR data set with a total exposure of $\alpha_T = 89, 720 \text{ km}^2 \text{ sr yr}$, and also to lower energy events, with their analysis considering 5514 events down to 20 EeV. Secondly, we consider a pure proton model, whereas they compute the attenuation weights for a mixed composition based on the results of a joint analysis of the UHECR spectrum and composition (Aab et al. 2018b), which are then included in the generation of the presented flux maps. Finally, they also assume a range of fixed source spectral indices ($\alpha = \{-1.5, 1, 2\}$) in the context of different scenarios which are chosen to represent a range of compositions and maximum rigidities at the sources. In this way, there are many factors that are different between the two analyses in addition to the statistical method employed and so the results should be interpreted accordingly.

In order to illustrate the fits and further analyse the results for f , we compute the association probabilities of each source-UHECR pair as described in Appendix B and summarize the results in Fig. 11. As expected for equal source luminosities, the associations are dominated by the nearest sources. For the 2FHL catalogue, Centaurus A dominates with a small contribution from M 87. For the SBGs, NGC 4945, M 83, and NGC 253 have the strongest associations with small contributions from almost all of the catalogue reflecting the higher value of f . In the case of the *Swift*-BAT catalogue, Centaurus A and the Circinus galaxy dominate. Here, the large number of sources mean that the contribution to the higher f value mostly comes from many low-probability associations. Fig. 11 also allows us to visualize the implication of the value of $\bar{B}$ resulting from the fit. We can see that for the 2FHL catalogue, the large tail in $P(\bar{B}|\hat{E}, \hat{\omega})$ allows for low-probability associations with low-energy UHECRs that are far from Centaurus A and M 83. The larger $\bar{B}$ found for the SBG case is represented by the stronger association probabilities for more deflected UHECRs, whereas the smaller $\bar{B}$ found for the *Swift*-BAT case is shown by the preference for smaller deflections.

For the joint model fits, the marginal posterior distributions of key parameters are summarized in Fig. 12. For each catalogue, we see that the fit favours a very steep energy spectrum of $\alpha \sim 6.5$. This can be understood by the fact that with the majority of UHECRs being associated with the background component of the model, the α resulting from the fit really represents the arrival energy spectrum and not that of the sources, and is therefore the same for each catalogue. The large α is needed to fit the arrival spectrum due to the lack of an exponential cut-off in our model for both the injection spectrum and the arrival spectrum. All fits generally favour $\bar{B} \sim 10 - 30 \text{ nG}$, with a longer tail out to $\sim 50 \text{ nG}$ for the 2FHL catalogue. These values correspond to deflections on the scale of $1.15^\circ (\bar{B}/1 \text{ nG})$ for a UHECR of 100 EeV travelling a distance of 10 Mpc (see equation 4). The deflections are not small for lower energy UHECRs and the approximation used is not strictly valid, but reflects the underlying relationship of the relevant parameters with the data.

Given the simplicity of our model, we do not attempt to draw direct astrophysical conclusions from the fit values of α and $\bar{B}$, but it is important to demonstrate that the described framework is able to include such parameters and these results can inform on future extensions to the propagation model which would allow for a more physical interpretation. An obvious advantage of the framework is that it is possible to simultaneously perform inference on a magnetic




Figure 10. Marginal posterior distributions for f for each of the source catalogues. In each case, we consider two different models: the arrival directions only and the joint model for UHECR energies.

Table 1. The results for f for both the arrival direction and joint models. The results are given as the mean value with the lower and upper bounds representing the 68 per cent region of highest posterior density.

	f (per cent)	
	Arrival direction	Joint
2FHL	$4.4^{+1.7}_{-2.6}$	$9.5^{+2.4}_{-5.9}$
SBG	$11.2^{+4.6}_{-7.6}$	$22.7^{+6.6}_{-12.4}$
Swift-BAT	$7.4^{+2.6}_{-5.4}$	$22.8^{+6.6}_{-8.0}$

field model and a UHECR model avoiding the need to fully specify one in order to interpret the other. This is one of the main goals of the IMAGINE consortium,³ which aims to bring together data from the fields of Galactic science, large-scale structure formation cosmic rays and to build a consistent picture of the phenomenology using Bayesian methods (Boulanger et al. 2018).

The model for the background component assumes no UHECR energy losses and thus has the flexibility required to provide a good fit to the data. A more physical model in which the background UHECRs are modelled as coming from beyond 250 Mpc was also considered and tested against the data. This model is much more constraining, requiring UHECRs which cannot be explained by the catalogued sources to be of lower energies, due to their more distant origin, and was unable to provide a good fit to the data for any of the source catalogues considered. This can be understood by considering the location of several of the highest energy UHECRs, shown in Fig. 9, which are either far from potential sources (2FHL and SBG), or the sources which seem ‘nearby’ on the celestial sphere are in fact distant (> 50 Mpc for *Swift-BAT*). We see in Fig. 11 that we do not have any source-UHECR associations for three of the highest energy UHECRs. Within the context of our model, UHECRs of these energies (> 100 EeV) are not strongly deflected in magnetic fields and cannot be from distant sources (see Fig. 5). In this way, the catalogues proposed cannot explain the UHECR data observed and thus the fit favours the more flexible isotropic background component, even if large deflections are considered.

We use PPCs to assess the ability of the joint model to fit the data and the results are shown in Appendix C. We find reasonable fits for all cases, with possible indications of extensions to the spectral and luminosity function modelling being able to better represent the data.

7 CONCLUSIONS

We present a hierarchical model for the UHECR energies and arrival directions, enabling a joint fit of a phenomenological model to the UHECR observations.

Through calculations and simulations within the framework of the model, we demonstrate the impact of including UHECR energy data and detection effects on the physical interpretation. Modelling reconstruction uncertainties of 12 percent for the UHECR energies results in an effective extension of the GZK horizon due to the misidentification of lower energy events in a UHECR sample. The size of this effect depends on the shape of the source spectrum, the model for the UHECR energy losses and the distribution of sources, but is non-negligible in many cases with ~ 30 per cent of a selected UHECR sample coming from beyond 250 Mpc for the example presented in Fig. 3. We also include the arrival direction reconstruction uncertainty of 1° and show the influence of the detection effects on the particle-specific effective GZK horizon. Again, the distribution of possible source locations of the UHECRs is considerably extended, with the energy uncertainty being mostly responsible.

By comparing fits of simulated data with the joint and simpler model for only the UHECR arrival directions, we show that including the UHECR energy information is essential in order to recover the input parameters. We interpret these results as a strong motivation for explicitly modelling the individual UHECR energies in future comparisons of candidate source populations and UHECR data. The joint model also enables to fit for the source spectral index and the rms magnetic field strength and can reasonably recover the input parameters of simulations run using CRPROP 3, showing that the model is a justifiable approximation capable of inferring physical parameters.

The model is applied to the publicly available data of 231 UHECR with $\hat{E} > 52$ EeV detected by the PAO and the 2FHL, SBG, and *Swift-BAT* catalogues of potential sources. We find the largest association fractions of $22.7^{+6.6}_{-12.4}$ and $22.8^{+6.6}_{-8.0}$ per cent for the SBG and *Swift-BAT* catalogues, respectively. The fact that very different source distributions can give similar results motivates extension of the model in order to test more constraining scenarios against the data. Whilst our model is able to fit the data well, several of the highest energy cosmic rays in the sample are left without strong source associations. This implies that the chosen catalogues cannot fully explain the observations under the assumptions of the model, and can inform on the choice of future candidate source catalogues.

Both the physical model and the formalism presented here can be extended to include more complex acceleration and propagation models. Additionally, we plan to include the available TA data (Abbasi et al. 2014) and detector exposure into this framework

³<https://www.astro.ru.nl/imagen>




Figure 11. The association probabilities of source-UHECR pairs resulting from the fit of the joint model are plotted together with the UHECR data for the different source catalogues. The lines are shaded according to the association probabilities (normalized to the maximum) and are plotted for the case of $P(\lambda_i = k|\hat{E}, \hat{\omega}) > 0.001$. The dominant sources are described in the text. As in Fig. 9, the sky map is shown in Galactic coordinates using the Hammer–Aitoff projection, the coloured tissots represent UHECR energies, with larger tissots corresponding to higher energy, and the source catalogues are also shown by solid black tissots.




Figure 12. Joint marginal posterior distributions of α , $\bar{B}$, and f for each of the candidate source catalogues. The shaded contours enclose the 30, 60, 90, and 99 per cent highest posterior density credible regions.

in future work and to account for the ~ 14 per cent systematic uncertainties on the UHECR energy scale.

Results from the PAO favour a heavier composition at the highest energies over pure protons (Aab et al. 2016). An obvious extension of our model would be to introduce the possibility of a mixed composition. This could be achieved by modelling the composition as an additional latent mark on the production, resulting in a marked Poisson process with three marks: $\{\lambda_i, \tilde{E}_i, C_i\}$. The simplest case would be to approximate the mark distribution for the C_i as a mixture model over five representative elements at the source, e.g. over H, He, N, Si, and Fe, as is typically done when modelling the composition. The weights of the mixture model could be left as free parameters to directly infer the source properties based on the data. A propagation model then needs to be specified for the composition-dependent energy losses and deflections. For the deflections, a straightforward option would be to consider equation (4) multiplied by the initial charge of the UHECRs, Z , and the energy losses could also be approximated to be purely a function of the initial composition of the UHECRs. The reconstruction of the composition from air shower data can be treated using Gumbel distributions (Gumbel 1935) which can be used to model the distribution of extreme values ($X_{\max}$ in this case) together with the parametrization described in De Domenico et al. (2013).

However, a realistic physical model needs to account for the changing composition during the UHECR propagation due to photodisintegration and the effect of this on the energy losses and magnetic deflections. For this reason, we model the UHECRs as protons here in order to demonstrate the use of the hierarchical model and only highlight the capability of the model to include the UHECR composition as a possibility for future work. An interesting avenue to explore in this case would be to interface a Monte Carlo simulation code for the UHECR propagation, such as CRPROPA 3, with the hierarchical model and statistical framework described here.

ACKNOWLEDGEMENTS

We thank J. M. Burgess and C. Fuglesang for useful discussions. This work made use of NUMPY (Oliphant 2006), SCIPY (Jones et al. 2001), ASTROPY (Price-Whelan et al. 2018), MATPLOTLIB (Hunter 2007), SEABORN (Waskom et al. 2017), H5PY (Collette 2013), and the High Performance Computing facilities at the PDC center of the Swedish National Infrastructure for Computing. FC was supported by the Alexandra and Bertil Gylling Foundation.

REFERENCES

- Aab A. et al., 2015a, *Nucl. Instrum. Methods A*, 798, 172
 Aab A. et al., 2015b, *ApJ*, 804, 15
 Aab A. et al., 2016, *Phys. Lett. B*, 762, 288
 Aab A. et al., 2018a, *ApJ*, 853 (A18)
 Aab A. et al., 2018b, *J. Cosmol. Astropart. Phys.*, 1704, 038
 Abbasi R. U. et al., 2018a, *ApJ*, 867, L27
 Abbasi R. U. et al., 2018b, *ApJ*, 76, 858
 Abbasi R. U. et al., 2014, *ApJ*, 790, L21
 Abraham J. et al., 2007, *Science*, 318, 938
 Abreu P. et al., 2010, *Astropart. Phys.*, 34, 314
 Abu-Zayyad T. et al., 2012, *Nucl. Instrum. Methods A*, 689, 87
 Achterberg A., Gallant Y. A., Norman C. A., Melrose D. B., 1999, preprint (arXiv:astro-ph/9907060)
 Ackermann M. et al., 2012, *ApJ*, 755, 164
 Ackermann M. et al., 2016, *ApJ*, 222, 19
 Aharonian F. A., Kelner S. R., Prosekin A. Y., 2010, *Phys. Rev. D*, 82, 043002
 Aloisio R., 2017, *Prog. Theor. Exp. Phys.*, 2017, 12A102
 Alvarez-Muñiz J., Engel R., Stanev T., 2002, *ApJ*, 572, 185
 Anchordoqui L. A., 2018, preprint (arXiv:1807.09645)
 Anchordoqui L. A., Dova M. T., Epele L. N., Swain J. D., 1997, *Phys. Rev.*, D55, 7356
 Arnold B. C., 2015, *Pareto Distributions Second Edition*. CRC press, Boca Raton, Florida
 Baring M. G., 2001, *AIP Conf. Proc. Vol. 516, Cosmic Ray Origin, Acceleration, and Propagation*. American Institute of Physics (AIP), Melville, NY, p. 153
 Batista R. A. et al., 2016, *J. Cosmol. Astropart. Phys.*, 2016, 038
 Bell A. R., 1978, *MNRAS*, 182, 147
 Berezhinsky V. S., Astrophysics S. G. A., 1988, *Astron. Astrophys.*, 199, 1
 Betancourt M., 2017, preprint (arXiv:1701.02434)
 Biermann P. L., Strittmatter P. A., 1987, *ApJ*, 322, 643
 Blasi P., Epstein R. I., Olinto A. V., 2000, *ApJ*, 533, L123
 Bonifazi C., 2009, *Nucl. Phys. B – Proc. Suppl.*, 190, 20
 Böttcher M., Reimer A., Sweeney K., Prakash A., 2013, *ApJ*, 768, 54
 Boulanger F. et al., 2018, *J. Cosmol. Astropart. Phys.*, 8, 049
 Carpenter B. et al., 2017, *J. Stat. Softw.*, 76, 1
 Chodorowski M. J., Zdziarski A. A., Sikora M., 1992, *ApJ*, 400, 181
 Collette A., 2013, *Python and HDF5*. O'Reilly, Boston, MA
 Dermer C. D., Razzaque S., 2010, *ApJ*, 724, 1366
 Dermer C. D., Razzaque S., Finke J. D., Atoyan A., 2009, *New J. Phys.*, 11, 065016
 De Domenico M., Insolia A., 2012, *J. Phys. G: Nucl. Part. Phys.*, 40, 015201
 De Domenico M., Settimo M., Riggi S., Bertin E., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 050
 Durrer R., Neronov A., 2013, *A&AR*, 21, 62
 Fang K., Murase K., 2018, *Nat. Phys.*, 14, 396
 Fang K., Kotera K., Olinto A. V., 2012, *ApJ*, 750, 118
 Fisher R., 1953, *Proc. R. Soc. A: Math. Phys. Eng. Sci.*, 217, 295
 Gelman A., 2006, *Bayesian Anal.*, 1, 515
 Gelman A., Rubin D. B., 1992, *Stat. Sci.*, 7, 457
 Gelman A., Meng X. L., Stern H., 1996, *Stat. Sinica*, 6, 733
 Gelman A., Simpson D., Betancourt M., 2017, *Entropy*, 19, 555
 Goulding A. D., Alexander D. M., Lehmer B. D., Mullaney J. R., 2010, *MNRAS*, 406, 597

- Greisen K., 1966, *Phys. Rev. Lett.*, 16, 748
- Gumbel E. J., 1935, *Ann. l'Institut Henri Poincaré*, 5, 115
- Hanlon W., 2018, preprint (arXiv:1801.00001)
- Harari D., Mollerach S., Roulet E., Sánchez F., 2002, *J. High Energy Phys.*, JHEP03(2002), 045
- Harari D., Mollerach S., Roulet E., 2006, *J. Cosmol. Astropart. Phys.*, 0611, 012
- Hill C. T., Schramm D. N., 1985, *Phys. Rev.*, D31, 564
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Jansson R., Farrar G. R., 2012, *ApJ*, 757, 14
- Jones E. et al., 2001, SciPy: Open Source Scientific Tools for Python. Available at: <http://www.scipy.org/>
- Kamary K., Mengersen K., Robert C. P., Rousseau J., 2014, preprint(arXiv:1412.0444)
- Khanin A., Mortlock D. J., 2016, *MNRAS*, 460, 2765
- Kim H. B., Kim J., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 006
- Kotera K., Lemoine M., 2008, *Phys. Rev. D*, 77, 123003
- Kotera K., Olinto A. V., 2011, *ARA&A*, 49, 119
- Kulsrud R. M., Zweibel E. G., 2008, *Rep. Prog. Phys.*, 71, 046901
- Loredo T. J., 2004, AIP Conf. Proc. Vol. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering. American Institute of Physics (AIP), Boston, MA, p. 195
- Lynch S. M., Western B., 2016, *Sociol. Methods Res.*, 32, 301
- Matthews J. H., Bell A. R., Blundell K. M., Araudo A. T., 2018, *MNRAS*, 479, L76
- Mücke A., Engel R., Rachen J. P., Protheroe R. J., Stanev T., 2000, *Comput. Phys. Commun.*, 124, 290
- Murase K., Dermer C. D., Takami H., Migliori G., 2012, *ApJ*, 749, 63
- Norman C. A., Melrose D. B. A. A., 1995, *ApJ*, 454, 60
- Oh K. et al., 2018, *ApJS*, 235, 14
- Oikonomou F., Connolly A., Abdalla F. B., Lahav O., Thomas S. A., Waters D., Waxman E., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 015
- Oliphant T., 2006, A Guide to NumPy. Tregol Publishing, USA
- Price-Whelan A. M. et al., 2018, *AJ*, 156, 123
- Ryu D., Kang H., Hallman E., Jones T. W., 2003, *ApJ*, 593, 599
- Simpson D., Rue H., Riebler A., Martins T. G., Sørbye S. H., 2017, *Stat. Sci.*, 32, 1
- Sinharay S., Stern H. S., 2003, *J. Stat. Plan. Inference*, 111, 209
- Soiaporn K. et al., 2013, *Ann. Appl. Stat.*, 7, 1249 (S13)
- Sommers P., 2001, *Astropart. Phys.*, 14, 271
- Stan Development Team, 2018, Stan Modeling Language Users Guide and Reference Manual. Available at: <http://mc-stan.org>
- Stanev T., Biermann P. L., Evans J. L., Rachen J. P., Watson A. A., 1995, *Phys. Rev. Lett.*, 75, L3056
- Streit R. L., 2014, Poisson Point Processes. Springer-Verlag, New York
- Tinyakov P. G., Tkachev I. I., 2001, *J. Exp. Theor. Phys. Lett.*, 74, 1
- Tokuno H. et al., 2012, *Nucl. Instrum. Methods A*, 676, 54
- Vallée J. P., 2004, *New Astron. Rev.*, 48, 763
- Vietri M., 1995, *ApJ*, 453, 883
- Waskom M., et al., 2017, mwaskom/seaborn: v0.8.1
- Watson L. J. et al., 2011, *MNRAS*, 418, 206
- Waxman E., 1995, *Phys. Rev. Lett.*, 75, 386
- Waxman E., Miralda-Escude J., 1996, *ApJ*, 472, L89
- Zatsepin G. T., Kuzmin V. A., 1966, *JETP Lett.*, 4, 78

APPENDIX A: DERIVATION OF THE LIKELIHOOD

Following the ideas introduced in Section 3.1, we proceed by starting at the top of the model hierarchy with the source population. The total population can be modelled as a superposition of homogeneous Poisson point processes from each component. UHECRs are produced in a marked Poisson point process, which means that they are produced with a set of associated properties, referred to as ‘marks’: $\{\lambda_i, \tilde{E}_i\}$. The λ_i are integer valued labels associating UHECRs i to a source components k . Their categorical

mark distribution is given by

$$P(\lambda = k | F_k) = \frac{F_k}{\sum_{j=0}^K F_j}. \quad (\text{A1})$$

These labels implicitly define the initial directions of the UHECRs, since they are produced at the source position. The $\tilde{E}_i$ are distributed according to the source power-law spectrum, characterized by α and normalized as described in Section 2.2. The resulting mark probability distribution for the latent energy marks, $\tilde{E}_i$, is thus a Pareto distribution (Arnold 2015)

$$P(\tilde{E}_i | \alpha) = \frac{\alpha - 1}{E_{\min}} \left(\frac{\tilde{E}_i}{E_{\min}} \right)^{-\alpha}. \quad (\text{A2})$$

UHECR propagation involves both energy losses and magnetic deflections, as described in Section 2.4. The arrival energies E_i are calculated directly from the $\tilde{E}_i$ by solving equation (2) for a given source distance, D_k , and UHECR redshift, z . In this way, $P(E_i | \tilde{E}_i, D_k)$ is a delta function

$$P(E_i | \tilde{E}_i, D_k) = \delta[E_i - E(\tilde{E}_i, D_k)]. \quad (\text{A3})$$

For the background component with $k = 0$, there are no energy losses and $E_i = \tilde{E}_i$, which is the equivalent of the above statement with $D_0 = 0$ Mpc. Magnetic deflections are modelled using the vMF distribution (defined in Section 2.5) centred on the source locations

$$P(\omega_i | \varpi_k, D_k, \tilde{E}_i, \bar{B}) = \frac{\kappa_{k,i}}{4\pi \sinh(\kappa_{k,i})} e^{\kappa_{k,i} \omega_i \cdot \varpi_k}, \quad (\text{A4})$$

where $\kappa_{k,i} = \kappa_{k,i}(D_k, \tilde{E}_i, \bar{B})$, as given by θ_{rms} in equation (4), which is then converted to the dimensionless vMF parameter using equation (9).

We now proceed to UHECR detection and reconstruction. The form of this component of the likelihood is given in equation (14). We first focus on the rate of the process. It is possible to marginalize over the latent UHECR arrival directions and find an analytic solution to the integral such that (S13)

$$\begin{aligned} P(\hat{\omega}_i | \varpi_k, D_k, \tilde{E}_i, \bar{B}) &= \int d\omega_i P(\hat{\omega}_i | \omega_i) P(\omega_i | \varpi_k, \kappa_{k,i}) A_{\perp}(\omega_i, t_i) \\ &\approx A_i \cos(\theta_i) \int d\omega_i P(\hat{\omega}_i | \omega_i) P(\omega_i | \varpi_k, \kappa_{k,i}) \\ &= A_i \cos(\theta_i) \begin{cases} \frac{\kappa_{k,i} \kappa_d}{4\pi \sinh(\kappa_{k,i}) \sinh(\kappa_d)} \frac{\sinh(|\kappa_d \omega_i + \kappa_{k,i} \varpi_k|)}{|\kappa_d \omega_i + \kappa_{k,i} \varpi_k|} & \text{if } k \geq 1 \\ \frac{1}{4\pi} & \text{if } k = 0 \end{cases}, \quad (\text{A5}) \end{aligned}$$

where $A_{\perp}$ is the effective area of the detector at time t , θ_i is the detected UHECR zenith angle and we have assumed that the angular reconstruction uncertainty, σ_{ω} , is small compared to the scale over which $A_{\perp}$ varies in the approximation. The dependence of $\kappa_{k,i} = \kappa_{k,i}(D_k, \tilde{E}_i, \bar{B})$ has been omitted for clarity. A_i denotes the area of the detector at UHECR arrival time t_i . This changes slowly as the size of a detector grows during its construction, as for the different periods described in Abreu et al. (2010). The resulting expression for the Poisson rate is given in equation (15).

The expected number of events, $\bar{N}$, is given by the sum over the contributions from all possible sources, as described in Section 3.1. The effective exposure, ε_k , is indirectly dependent on the source UHECR energy distribution, the source distance, and the magnetic field strength through the dependence of the vMF distribution on $\kappa_{k,i} = \kappa_{k,i}(D_k, \tilde{E}_i, \bar{B})$. As $\kappa_{k,i}$ is different for each UHECR, and we do not know the source-UHECR associations, $\tilde{E}_i$ or $\bar{B}$ a priori, we approximate the spread of the vMF by considering an expected κ for each source, $\bar{\kappa}_k$. We have

$$\varepsilon_k(\bar{\kappa}_k) = \begin{cases} \int d\omega P(\omega | \varpi_k, \bar{\kappa}_k) \varepsilon(\omega) & \text{if } k \geq 1 \\ \varepsilon_{\text{AT}}/(4\pi) & \text{if } k = 0 \end{cases}, \quad (\text{A6})$$

which is weakly dependent on $\bar{\kappa}_k$, particularly in the limit of $\bar{\kappa}_k \gg 1$.

In order to summarize the distribution of energies at the source, we consider the median of the source power-law distribution: $\tilde{E}_m = 2^{1/(\alpha-1)} \tilde{E}_{th}$. Consequently, $\bar{\kappa}_k = \bar{\kappa}_k(D_k, \alpha, \tilde{E}_{th}, \tilde{B})$ and is independent of the UHECR index, i . Bringing this result together with equation (10) and summing over all source components, we arrive at equation (16).

APPENDIX B: SOURCE-UHECR ASSOCIATION

As described in Section 3.1, UHECRs are modelled as having integer labels, λ_i , which denote their sources. Having marginalized out the λ_i in the likelihood (equation 15), meaning that the λ_i are not explicitly sampled during inference, we now show how to recover the marginal posterior distribution for λ_i , as described for a similar problem in Stan Development Team (2018). The likelihood function can be expressed as

$$P(\hat{E}, \hat{\omega}|L, \alpha, \tilde{E}, \tilde{B}) \propto \prod_{i=0}^N \sum_{k=0}^K P(\lambda_i = k, \hat{E}_i, \hat{\omega}_i|L, \alpha, \tilde{E}, \tilde{B}) \quad (B1)$$

$$\propto \prod_{i=0}^N \sum_{k=0}^K P_{k,i},$$

where the inner term, $P_{k,i}$, is the joint probability for $\lambda_i = k$ and the observed data, given the model parameters. During sampling, assuming convergence, the values for the model parameters are drawn from the joint posterior distribution, $P(L, \alpha, \tilde{E}, \tilde{B}|\hat{E}, \hat{\omega})$. At each iteration, the associated $P_{k,i}$ is calculated. Averaging this term over the iterations gives an estimate of the unnormalized marginal posterior for $\lambda_i = k$

$$P(\lambda_i = k|\hat{E}, \hat{\omega}) \propto q(\lambda_i = k|\hat{E}, \hat{\omega}) = \frac{1}{S} \sum_{s=0}^S P_{k,i}^s, \quad (B2)$$

where S is the number of iterations, or posterior samples. We can then normalize this by dividing by the sum over all possible associations of the i th UHECR

$$P(\lambda_i = k|\hat{E}, \hat{\omega}) = \frac{q(\lambda_i = k|\hat{E}, \hat{\omega})}{\sum_{l=0}^K q(\lambda_i = l|\hat{E}, \hat{\omega})}, \quad (B3)$$

thus recovering the normalized marginal posterior distribution for the individual UHECR labels λ_i .

APPENDIX C: POSTERIOR PREDICTIVE CHECKS

In order to assess the ability of our model to fit the data, we perform PPCs as introduced in Section 3.4. The posterior predictive distribution for this data set is defined as

$$P(\hat{E}_{rep}, \hat{\omega}_{rep}|\hat{E}, \hat{\omega}) = \int d\Theta P(\hat{E}_{rep}, \hat{\omega}_{rep}|\Theta) P(\Theta|\hat{E}, \hat{\omega}), \quad (C1)$$

where Θ is the set of all model parameters and $\hat{E}_{rep}$ and $\hat{\omega}_{rep}$ are the replicated data that we wish to generate. We draw $\hat{E}_{rep}$ and $\hat{\omega}_{rep}$ from this distribution and compare with the observed $\hat{E}, \hat{\omega}$ in order to evaluate the how well the model represents the data. Systematic differences between the replicated and observed data are an indicator of model misfit. In this way, in addition to assessing the quality of a fit, PPCs can be used as a tool to indicate in which ways to extend a model in order to better represent the data.

We compute 100 posterior predictive draws for both the fit of the simulated data discussed in Section 4 and the fits of the PAO data set with different source catalogues presented in Section 5. The results are shown in Fig. C1. The simulated data show a good fit, as is expected given that the model is the correct data generating process in this case. There are no systematic deviations between $\hat{E}$ and $\hat{E}_{rep}$, and the clustering in the distribution of $\hat{\omega}$ is reflected in $\hat{\omega}_{rep}$. For the fits to data, the $\hat{E}_{rep}$ capture the overall shape of the detected energy spectrum but not the detail, and the slight shoulder in lower energies of $\hat{E}$ is missed. This could be due to modelling the UHECR energy spectrum with a simple power law with no cut-off and using the same spectral index for the source and background component, both of which are approximations used in this more simplistic model. The $\hat{\omega}_{rep}$ distributions are more isotropic than that of the simulation, due to smaller association fractions being found. However, there is some slight clustering around the nearest source locations, as would be expected from the equal luminosity assumption.

These results are consistent with an associated fraction $f < 1$ due to some UHECRs coming from an unresolved background component due to distant sources, and others from sources which are not present in the chosen catalogues.




Figure C1. 100 draws from the posterior predictive distribution, with the data, for the case of the simulated data set (upper left) and fits to the PAO data using the 2FHL (upper right), SBG (lower left), and *Swift*-BAT (lower right) catalogues. In each case, the energy spectrum and arrival direction distributions are shown separately with the data in solid colour and the replicated data overplotted in transparent grey.

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.