

Tuning PYTHIA for forward physics experiments

Max Fieg¹, Felix Kling², Holger Schulz³, and Torbjörn Sjöstrand⁴

¹*Department of Physics and Astronomy, University of California, Irvine, California 92697-4575, USA*

²*Deutsches Elektronen-Synchrotron DESY, Notkestrasse 85, 22607 Hamburg, Germany*

³*Tharsus Limited, Coniston Road, Blyth NE24 4RF, United Kingdom*

⁴*Department of Physics, Lund University, Box 118, 221 00 Lund, Sweden*



(Received 2 October 2023; accepted 28 November 2023; published 9 January 2024)

Event generators like PYTHIA play an important role in physics studies at the Large Hadron Collider (LHC). While they make accurate predictions in the central region, i.e., at pseudorapidities $\eta < 5$, a disagreement between PYTHIA and measurements in the forward region $\eta > 7$ has been observed. We introduce a dedicated forward physics tune for the PYTHIA event generator to be used for forward physics studies at the LHC, which uses a more flexible modeling of beam remnant hadronization and is tuned to available particle spectra measured by LHCf. Furthermore, we provide an uncertainty estimate on the new tune in a data-driven way which can be used as a means of flux uncertainty for future forward physics studies. We demonstrate an application of our tune by showing the updated neutrino and dark photon spectra at the FASER experiment.

DOI: [10.1103/PhysRevD.109.016010](https://doi.org/10.1103/PhysRevD.109.016010)

I. INTRODUCTION

The Large Hadron Collider (LHC) has been instrumental in constraining physics both within and beyond the Standard Model (BSM). Its main experiments, ATLAS, CMS, LHCb, and ALICE, have discovered and measured properties of the Higgs, constrained dark sectors, probed new physics in the flavor sector, and, more generally, furthered our understanding of fundamental particle physics. These experiments benefit greatly from Monte Carlo event generators, which can make accurate predictions of particle distributions in the central region with pseudorapidities $\eta \lesssim 5$. Much work has been put into improving, validating, and tuning these generators for the experiments at the LHC, and often excellent agreement has been reached.

Recently, there has been new interest in particle production in the forward direction at the LHC, corresponding to $\eta \gtrsim 7$, where much less data has been collected as compared to the central experiments. The implementation of the FASER experiment has already set leading bounds in certain BSM scenarios [1] and lead to the first direct observation of neutrinos produced at a collider [2,3]. Additionally, the Forward Physics Facility (FPF) has been proposed to house a suite of experiments to further study

particles produced in the forward direction during the high-luminosity LHC era [4,5]. The success of these experiments will be greatly enhanced if similar event generators can be used to make accurate predictions.

However, in the context of the LHC, the popular event generator PYTHIA [6,7] has been tuned only in the central region, and, thus, one should not expect reliable predictions in the forward direction. Indeed, the LHCf experiment, which can measure distributions of neutral particles with $\eta \gtrsim 9$, shows a distinct disagreement with PYTHIA's predictions obtained using the popular tune relying on data from central experiments—the so-called *Monash* tune [8]. Notably, PYTHIA predicts an excess of mesons but a deficit of baryons when compared to LHCf data [9–12].

In this paper, we provide a forward physics tune for the PYTHIA event generator by fitting hadronization parameters to LHCf measurements of neutral pion, photon, and neutron production. In particular, we will primarily fit parameters that have little impact on central physics so as to not spoil the success of PYTHIA in this region.

In addition to our forward tune, we will also provide an uncertainty estimate on these parameters. Currently, existing generators typically provide only one central prediction but no measure of uncertainty. One approach often used in astroparticle physics is to define an uncertainty based on the spread of event generators' predictions. While this definition captures a spread of underlying physics modeling, it is not data-driven and it is not clear if it has any statistical meaning. Here, for the first time, we follow a different approach and provide the uncertainty on a single generator in a data-driven way.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

This paper is organized as follows. In Sec. II, we discuss how hadronization is done in PYTHIA in the forward direction. In Sec. III, we discuss our tuning procedure to the LHCf measurements and provide our tune on these kinematic parameters. In Sec. IV, we show how our tune impacts the predictions for forward neutrino and dark photon production at the FASER experiment. In Sec. V, we summarize and conclude.

II. MODELING OF FORWARD PARTICLE PRODUCTION IN PYTHIA

There are few theory constraints in the modeling of forward physics. While at least some aspects of central physics are governed by perturbation theory, such as jet production, the forward region is entirely of nonperturbative origin.

An early assumption was so-called Feynman scaling [13], i.e., that the $x_E dn/dx_F$ distribution should be collision-energy independent. Here, $x_F = 2p_z/E_{CM}$ and $x_E = 2E/E_{CM}$ in the rest frame of the event, and n is the number of produced particles per event. Perfect Feynman scaling would correspond to a collision-energy-independent central rapidity plateau dn/dy , while data instead show this distribution to be rising with energy, suggesting that an increasing fraction of the total energy is taken from the forward region to produce more particles in the central one.

Central particle production in PYTHIA is generated by multiparton interactions (MPIs). That is, since hadrons are composite objects, several parton-parton subcollisions can occur inside a single pp event. The high- p_\perp tail of these subcollisions corresponds to regular jet production. The bulk of them occur at a few GeV, however, where they are not distinguished individually but are visible only by their collective effect. The rise of the rapidity plateau is mainly driven by an increasing average number of MPIs.

The beam remnants stem from those partons that are *not* kicked out of the incoming protons by MPIs. The remnants and the MPIs are related to each other by flavor and color. An MPI can take away both valence and sea quarks from the original uud valence content of a proton, giving rise to varied remnant topologies, e.g., ud or $uud\bar{s}$. Each kicked-out (anti)quark also carries away some (anti)color, and each gluon both a color and an anticolor, that have to be compensated in the remnant so as to preserve the overall color state. In the Lund string model [14], each separated color-anticolor pair gives rise to a linear confinement field, a string, that will fragment into hadrons. This would mean that the momentum of a remnant often had to be shared between many string systems, making it difficult to obtain a leading baryon that carries a significant fraction of the incoming proton energy. Also note that the number of MPIs goes up with increasing collision energy, implying softening baryon spectra.

Indeed, the problem in PYTHIA is to produce a spectrum with a fair amount of high-momentum baryons, and some

corrections have to be introduced to the baseline picture, as will be outlined in this section. We here do not consider the class of elastic scattering, which obviously is quite separate and not of interest here. We also leave diffraction aside for now but return to it later.

Early on [15], it was realized that a picture of fully independent MPIs does not reproduce collider phenomenology, e.g., the rise of the average transverse momentum of charged particles with increasing multiplicity. Hence the need for color reconnection (CR), the assumption that nature has a tendency to rearrange colors such that the total string length to be drawn out is reduced. Many possible scenarios have been proposed over the years, and a few of them are implemented in PYTHIA. We will study two of them here.

In the default CR scenario, it is assumed that the partons pulled out from the colliding protons are strongly correlated in color, in a way that the color of one such parton may be the same as the anticolor of another such. In a picture where only gluons are pulled out, the resulting remnant would then be in a color octet state, which conveniently can be subdivided into a triplet single quark and an antitriplet diquark. If, in addition, one valence quark is kicked out, only a diquark remains. These are the two most common outcomes, but others are possible and modeled. One is that all three valence quarks are kicked out. Then a single gluon is assigned to carry the remaining energy and momentum. Another is that the removal of sea quarks leaves their antipartners behind. Then the remnant is simplified by splitting off a hadron, e.g., $uud\bar{s} \rightarrow ud + u\bar{s} \rightarrow ud + K^+$.

The other scenario is the QCDCR one [16]. In it, explicit colors are assigned to both quarks and gluons, and reconnections can occur between identical colors if they reduce the total string length. Such a detailed tracing of color is not done in the default scenario. Another distinguishing feature of QCDCR is so-called junction reconnections. In it, two triplet strings can combine into an antitriplet one, according to the familiar color algebra $\mathbf{3} \otimes \mathbf{3} = \bar{\mathbf{3}} \oplus \mathbf{6}$. This leads to Y-shaped color topologies that carry nonvanishing baryon numbers. Notably, the QCDCR model correctly predicts an increased fraction of charm baryons in pp relative to e^+e^- events, which the default does not [17,18].

Zooming in on the remnant region, the QCDCR starting point is again to assign explicit colors to each parton pulled out of the incoming protons, with opposite colors in the remnant. This allows a bigger color charge to accumulate in the remnant than assumed in the default scenario, and this requires additional remnant gluons. In a first instance, the remnant is only simplified when, e.g., the color of one gluon equals the anticolor of another gluon. But again, high remnant color charges are deemed less likely, so an exponential suppression in the size of the remnant multiplet is introduced, whereby more remnant color lines are forced to cancel.

In the following, we will introduce a new forward physics tune that uses the QCDCR scenario with its suggested parameter values [16] as a starting point. On top of that, some old or new parameters are varied, with a special eye toward consequences in the forward region. An alternative tune that uses the default CR scenario and the Monash tune [8] as a starting point is presented in the Appendix.

Whenever the remnant consists of more than one parton, the remnant energy and (longitudinal) momentum have to be shared between them. To this end, there are assumed shapes for valence and sea quark momentum fractions x , as well as for gluons. With each x first picked at random according to these shapes and then rescaled to unit sum, each parton energy is now assigned a fraction x_{rescaled} of the full remnant energy. A diquark receives the sum of the constituent quark x values but is, in addition, allowed a further enhancement factor, by default 2. A remnant hadron receives the sum of its constituent momenta. The bottom line is that, in the two most common cases, either a diquark carries the full remnant momentum or it carries an average of 80% of it.

It is this diquark that primarily can fragment to produce the leading baryon, e.g., the neutron measured by LHCf. In spite of the steps already taken to make the diquark hard, it still turns out that the default fragmentation results in too soft neutrons. We have, therefore, sought ways to further harden the leading baryon spectrum. This requires modifications to the fragmentation of a leading diquark, relative to the normal string scenario.

To give some background, consider the normal string fragmentation, as probed most directly in e^+e^- annihilation events, $e^+e^- \rightarrow \gamma^*/Z^0 \rightarrow q_0\bar{q}_0$. There, the string between the q_0 and \bar{q}_0 breaks by the production of new $q_i\bar{q}_i$ pairs, to give a sequence $q_0\bar{q}_1 - q_1\bar{q}_2 - q_2\bar{q}_3 - \dots - q_{n-1}\bar{q}_0$ of n mesons. Here, $q_0\bar{q}_1$ is called the first-rank hadron of the q_0 jet, $q_1\bar{q}_2$ the second-rank one, and so on. The simplest extension to baryon production is to allow also antidiquark-diquark breaks, where the color antitriplet diquark takes the role of an antiquark, and vice versa. Thereby the baryon and

antibaryon are nearest neighbors in rank, giving rise to both flavor and momentum correlations. Specifically, since two flavor pairs are shared, you could not produce a $\Xi - \bar{p}$ combination this way. Studies mainly at LEP have shown that baryon-antibaryon pairs are more decorrelated than this picture allows for.

This is where the popcorn mechanism enters. In it, diquarks are not bound objects, but quarks can drift apart along the string, in such a way that a meson can be produced between the baryon and antibaryon, whereby the latter two share only one $q_i\bar{q}_i$ pair. Tunes to LEP data suggest that half the time the baryon and antibaryon are nearest neighbors, and half the time they are separated by a meson in between. Translated to the fragmentation of a leading diquark, this means that the production of a baryon and of a meson as the first-rank particle are about equally likely. But we do not have quite as nice a test bed for diquark fragmentation as e^+e^- offers for the quark one and also have not spent a corresponding effort at tuning, so this assumption is untested. On the contrary, it is plausible that an initial diquark from an incoming proton sticks together better than assumed for new string breaks. Therefore, we introduce a new parameter d_{pop} (see Table I for the full name in the code) uniquely for diquarks at the end of strings. If zero, then such a diquark will never break up, while if unity, such a split is as likely as inside a string. A second-rank baryon takes less average momentum than a first-rank one does, so a reduced admixture of the former gives a harder baryon spectrum.

For an initial parton in a string aligned along the z axis, the first-rank hadron takes a fraction z_1 of the total light-cone momentum $E + p_z$, the second-rank a fraction z_2 of what is left after the first, i.e., a fraction $z_2(1 - z_1)$ of the original amount, and so on. In each step, we assume the z value to be picked at random according to the Lund symmetric fragmentation function (LSFF). In its full generality, the LSFF allows for one separate parameter for each quark or diquark flavor species and quark or diquark mass correction factors for the first-rank hadron.

TABLE I. The main PYTHIA parameters studied in this article, their default parameters in the QCDCR tune (according to the *mode 2* configuration in Ref. [16]), and their values in the forward physics tune obtained in this study. The last column shows the uncertainty range for $\sigma_{\text{soft}} = \sigma_{\text{remn}}$ as discussed in Sec. III D.

Full name	Shorthand	Baseline (QCDCR)	Forward tune	Uncertainty
BeamRemnants:dampPopcorn	d_{pop}	1	0	
BeamRemnants:hardRemnantBaryon	f_{remn}	Off	On	
BeamRemnants:aRemnantBaryon	a_{remn}	...	0.36	
BeamRemnants:bRemnantBaryon	b_{remn}	...	1.69	
BeamRemnants:primordialKTsoft	σ_{soft}	0.9	0.58	0.26...1.27
BeamRemnants:primordialKTthard	σ_{hard}	1.8	1.8	
BeamRemnants:halfScaleForKT	Q_{half}	1.5	10	
BeamRemnants:halfMassForKT	m_{half}	1	1	
BeamRemnants:primordialKTremnant	σ_{remn}	0.4	0.58	0.26...1.27

In practice, this full generality is seldom used, and then the LSFF simplifies to

$$f(z) \propto \frac{1}{z} (1-z)^a \exp\left(-\frac{bm_{\perp}^2}{z}\right). \quad (1)$$

Here, $m_{\perp}^2 = m^2 + p_{\perp}^2$ is the squared transverse mass of the produced hadron, and a and b are free parameters to be tuned. A relevant aspect is that hadrons with a larger mass also take a larger average z value. Nevertheless, it appears that the forward baryon spectrum needs to be harder than is default. For the purposes of this tune, we have, therefore, allowed a and b to be set separately when a diquark jet produces a first-rank baryon, hence the new parameters a_{remn} and b_{remn} which can be turned on by setting $f_{\text{remn}} = \text{on}$. In a future with more data and understanding at hand, alternative modifications could be considered.

Note that the Lund string model is soft and collinear safe, in that emissions of such gluons only mildly affect the hadronic final state [19]. This is relevant for the role of parton showers in the forward direction. The PYTHIA initial-state shower is described by standard Dokshitzer-Gribov-Lipatov-Altarelli-Parisi evolution, with a smooth lower turnoff at around a scale of 2 GeV. Thus, the theoretically less-well-understood low- Q^2 and low- x region is only mildly probed. Presumably, a detailed low- p_{\perp} modeling would result in a softer diquark momentum than currently, which then would have to be absorbed into a somewhat harder leading diquark fragmentation than the here-presented tunes will give.

In addition to the flavor and longitudinal structure of particle production, also the transverse fragmentation must be considered. Here, the discussion can be split into the partonic setup and the string fragmentation.

In the first stage, each parton taken out of the incoming proton to become part of an MPI is assumed to have some transverse motion, “primordial k_{\perp} .” This is expected to be of the order of the quark constituent mass, say, a third of a GeV. For hard processes, notably Z -boson production, empirically a higher scale of the order of 2 GeV is required. This could be owing to an imperfect modeling of the low- p_{\perp} behavior of initial-state parton showers, but whatever the reason an interpolation is introduced wherein soft systems receive a lower primordial k_{\perp} and hard systems a higher one. The full expression for the Gaussian width σ is

$$\sigma = \frac{\sigma_{\text{soft}} Q_{\text{half}} + \sigma_{\text{hard}} Q}{Q_{\text{half}} + Q} \frac{m}{m + m_{\text{half}} \sqrt{\frac{E}{m}}}. \quad (2)$$

Here, the Q , m , and E are the hard scale, mass, and energy, respectively, of the MPI subsystem, while σ_{soft} , σ_{hard} , Q_{half} , and m_{half} are free parameters. The second factor is intended to reduce σ for low-mass systems, especially if these are strongly boosted in the forward direction ($E \gg m$).

Also, the left-behind constituents of the beam remnants, mainly quarks and diquarks, are each assigned a primordial k_{\perp} with a Gaussian width σ_{remn} . Taken together, the MPI initiators and the remnant constituents add to give a net p_{\perp} . An opposite recoil is shared evenly by them all, except that the damping factor for low-mass systems in Eq. (2) is used also here, such that transverse momentum overall is conserved.

With the kinematics of partons fully defined, string fragmentation can be applied. Again, consider a string piece aligned along the z axis. Then, in each string break, the new q_i and \bar{q}_i are assumed to receive opposite and compensating p_{\perp} kicks, which add vectorially to give the total p_{\perp} of each $q_i \bar{q}_{i+1}$ hadron. Again, a Gaussian distribution is used, with width σ . The full p_{\perp} of a hadron is then obtained after the rotation and boosted back to the full pp collision frame, which notably depends on the primordial k_{\perp} assigned to the remnant in the previous step.

A final note.—So far, we have considered nondiffractive events. Diffraction in PYTHIA is based on the Ingelman-Schlein picture [20], wherein a diffractive system can be modeled as a proton-glueball collision, where the glueball “hadron” is viewed as a representation of a Pomeron. Notably, the proton end of this system, which is in the forward direction, is next to identical with the one of a nondiffractive system. The glueball end usually is at more central rapidities and has negligible impact on the forward region. The picture is slightly modified for low-mass diffraction but is there assumed dominated by the production of a string with one leading diquark. Therefore, the modifications already introduced for nondiffractive events can be reused, without the introduction of any further ones.

In summary, the two main new modifications of the PYTHIA code are to allow a reduced probability for a remnant diquark to break up and to allow a harder fragmentation function for it. In addition, some existing parameters are also modified within the tuning effort.

III. TUNING KINEMATICS

As described in the previous section, the modeling of forward particle production introduces a number of phenomenological parameters. Their role is to parametrize the inability to make first-principle predictions in the absence of perturbative methods. For the simulation to have predictive power, it is imperative that these parameters are set to values (“tuned”) in such a way that the simulation reproduces a wide range of measured datasets, in this case from LHCf. In this section, we first discuss the datasets, parameters, and methodology before presenting the results in the form of a forward physics tune that is based on the QCDCR scenario. The tuning parameters and their values for both the baseline tune and the forward physics tune are shown in Table I. Results for an alternative tune that is based on the default CR scenario and the Monash tune are presented for comparison in the Appendix. We take the

default values for the remaining PYTHIA parameters that we do not tune, including the choice of parton distribution function (PDF) which we take to be NNPDF2.3 QCD + QED LO $\alpha_s(M_Z) = 0.130$ [21].

A. Datasets

We exclusively use data measured by the LHCf experiment for tuning purposes in this study, as it is by far the most relevant source of information on forward particle production. LHCf measured neutral hadron and photon fluxes at forward rapidities $\eta \gtrsim 8.8$ [22]. It is worth noting that forward photon production is dominated by $\pi^0 \rightarrow \gamma\gamma$ decay. We reasonably assume that the same mechanisms govern hadronization mechanisms at $\sqrt{s} = 7$ and 13 TeV collision energies. We therefore use LHCf data from both energies. The following list is a summary of the LHCf datasets we use to tune our phenomenological parameters with:

- (i) neutron energy spectra at 7 TeV [9],
- (ii) neutron energy spectra at 13 TeV [10],
- (iii) π^0 energy spectra at 7 TeV [11], and
- (iv) photon p_z spectra at 13 TeV [12].

The data are publicly available in the form of histograms of cross sections that are differential in either η or p_\perp .

We note that we use a very recently published LHCf measurement on η mesons [23] for validation of our methodology. We further validate our result by confronting the tuned simulation with more central measurements from CMS and TOTEM in Sec. III E.

B. Tuning parameters

Our mission is to identify and tune the value of phenomenological parameters relevant to forward physics while at the same time keeping the excellent predictive power of PYTHIA for central physics intact. In this context, working with parameters related to the modeling of the beam remnants (Table I) is a natural choice. They predominantly influence forward particle production while, as we will show, their influence on central particle production is limited. In the following, we discuss the effects these parameters have on the predictions of forward particle spectra and how the parameters are tuned to data, and, finally, we present a robust uncertainty estimate for the most relevant parameters.

Compared to the experimental data, the default PYTHIA configuration predicts too many hard pions in the LHCf phase space. Disabling the popcorn mechanism for meson production from beam remnants (i.e., setting $d_{\text{pop}} = 0$) leads to the desired reduction of hard pions. We note that we studied the effect of varying d_{pop} but found only little sensitivity for small $d_{\text{pop}} > 0$ and, hence, set this parameter to 0. A side effect of disabling the popcorn mechanism in beam remnants is an increase in the production of hard neutrons, simply because remnant diquarks can no longer hadronize into mesons. This turns out to be fortuitous, as

PYTHIA's default predicts too few hard neutrons in the most forward direction $\eta > 10.76$.

By adjusting other parameters associated with the beam remnant, we can tune the overall normalization of the forward hadronic flux. In particular, we can modify the initial k_\perp of the partons in the incoming protons: Partons with a relatively larger k_\perp will generally pull hadrons toward distributions of smaller η . The phenomenology of this effect is governed by the width of the primordial k_\perp distribution for the MPI initiators. The corresponding tuning parameters are σ_{soft} , σ_{hard} , and Q_{half} , and, for the beam remnant, σ_{remn} . The net effect is a nonzero p_\perp imparted on hadrons, the manifestation of which can be seen in the forward neutron and pion spectrum.

The overall effects of σ_{soft} , σ_{hard} , and σ_{remn} on PYTHIA's predictions for LHCf measurements are qualitatively similar, while their sensitivities are not (see our discussion in Sec. II). An increase in any of these parameters makes it more likely that forward hadrons inherit larger transverse momenta and, therefore, populate more central phase-space regions (i.e., bins with smaller η in the LHCf data). We exploit this freedom the model gives us and take a pragmatic approach. To keep σ_{hard} at its default value of 1.8 GeV, we reduce its sensitivity by increasing the (poorly constrained) Q_{half} to 10 GeV. As can be seen in Eq. (2), this makes the k_\perp distribution more dependent on σ_{soft} . To remove the remaining degeneracy between σ_{soft} and σ_{remn} , which have default values of 0.9 and 0.4 GeV, respectively, we define a parameter σ that relates the two: $\sigma = \sigma_{\text{soft}} = f\sigma_{\text{remn}}$, where f is a number that fixes the ratio. We studied the effect of tuning σ when choosing different values of f in the vicinity of $f = 1$. Since we found only marginal improvement, we choose to fix f at a value of $f = 1$ and keep only σ as a tuning parameter.

Two parameters, a_{remn} and b_{remn} , that govern the baryon fragmentation function complete our set of tuning parameters. They allow us to have an almost exclusive handle on the neutron spectrum, without much impact on the pion spectrum. In our setup, lowering (raising) a_{remn} while raising (lowering) b_{remn} results in slightly harder (softer) forward neutron spectra. Initially, we studied the effect of treating a_{remn} and b_{remn} as independent tuning parameters. However, we found that equally good quality of PYTHIA predictions can be achieved by fixing a_{remn} to the base tune's value for the LSFF of 0.36 and tuning only b_{remn} .

C. Tuning methods

The observations detailed in the previous paragraph lead to a reduction of the dimensionality of the tuning problem. We are left with two free parameters σ and b_{remn} , which we will explore in the ranges of [0–1.5] and [0.5–5], respectively. We use PYTHIA to simulate seven million pp collisions at $\sqrt{s} = 7$ TeV and five million collisions at 13 TeV for each point we initially explore in the $(\sigma, b_{\text{remn}})$ space. We analyze the simulated events with Rivet, enabling

the analysis routines that correspond to the experimental data listed in Sec. III A. The result of the analyses is a set of histograms obtained from simulation that can immediately be compared with the corresponding experimentally measured histograms. It should be noted that we obtain a set of histograms for each point in the so-explored parameter space.

Equipped with experimentally measured histograms and a method to obtain simulated histograms for any point in the parameter space, we could define a goodness-of-fit measure and numerically find a best-fit point that minimizes the measure. However, the computational cost to do so is prohibitively expensive. Instead, we construct an analytic surrogate model of the simulation response to shifts in the parameter space. The model allows us to predict the simulation outcome at any point in the parameter space at a fraction of the cost of computing the actual simulation. Not only is the model cheap to evaluate, but, due to its analytic nature, it is also straightforward to compute first- and second-order derivatives. These qualities make it an ideal fit for numerical minimization. We use the *Apprentice* toolkit for event generator tuning [24] to facilitate the construction of the surrogate, the definition of a goodness-of-fit measure, and the minimization thereof. We explored different options for the surrogates and found no benefit in going beyond quadratic polynomials. As input to the surrogate, we use the full simulation results at 64 uniformly distributed points in the specified range for σ and b_{remn} .

The *Apprentice* toolkit allows to bias the goodness-of-fit measure using a weighting mechanism for individual histograms and even bins. In general, one might wish to better reproduce either the neutron spectra, photon spectra, pion spectra, or a subset of certain η bins. We, however, wish to be agnostic and place the neutron, photon, and pion spectra measured at LHCf on equal footing. Since the datasets under consideration have quite different numbers of bins, we decided on a democratic weighting such that each of the four analyses is normalized according to the number of data points in that analysis. For a given particle spectrum and collision energy from Sec. III A, the weighting can be expressed as $w = (N_{\text{bins}})^{-1}$, where N_{bins} is the number of data points across η (or p_{\perp}) bins in that set.

Apprentice is then used to minimize the weighted goodness-of-fit measure. The outputs are a best-fit point $\sigma_0, b_{\text{remn},0}$, and predicted spectra at that point, computed from the surrogate model. These spectra are compared against the actual output of the simulation when run with the parameters of the best-fit point in a necessary effort to validate the method. The best-fit values for σ and b_{remn} for our forward physics tune can be found in Table I.

D. Tuning uncertainty

In addition to the central tuning prediction, we wish to provide a measure of uncertainty on our best fit.

An approach to estimate the uncertainty sometimes used in astroparticle physics is taken to be the spread in different event generators' predictions. While this does capture differences in underlying physics modeling, this definition is not data-driven and the error band lacks statistical meaning.

Naively, one might follow the usual method of looking for $\Delta\chi^2 = 1$ to obtain a 68% confidence interval. However, due to unknown correlations in experimental data and imperfections in the physics modeling, the goodness-of-fit measure does not follow a χ^2 distribution. If one were to nonetheless follow that approach with our model, the observed χ^2_{min} results in an unusable underestimate of uncertainties.

In light of this, we take a more practical approach. Our goal is to provide a well-defined range for our tuning parameters that can return a spread of particle fluxes for future studies at the FPF. This range can be obtained by varying the prediction in the vicinity of the best fit and testing how much the predictions change. The question remains: How much should one vary the tuning parameters to find the corresponding upper and lower bound? A practical parameter uncertainty range is one that covers distances of PYTHIA's prediction at the best-fit point from the experimentally measured data and data uncertainties.

We find that our fitting parameters σ and b_{remn} are not strongly correlated and that deviations about the best-fit point are most sensitive to σ . We therefore choose to vary and provide an uncertainty on σ . To obtain this uncertainty, we define a prediction band specified by two points $(f \times \sigma_0, \sigma_0/f)$, where f is a number that is increased until the band contains 68% of points (for $f = 1$, the band obviously contains zero points). Now, even for extremal values of σ in our range, there are a small number of data points which PYTHIA has difficulty describing; the central value of these points lies just outside the prediction range specified by $\sigma \in [0-1.5]$ and are typically found in the highest or lowest bins of the energy spectrum. Since we do not want those points to drive our estimation of uncertainty, we exclude them when counting the fraction of points inside the band specified by f . Across the four analyses, there are 20 of these out of 306 total data points.

The method yields two parameter points σ_- and σ_+ which define a robust uncertainty band containing 68% of points: $0.26 < \sigma < 1.27$.

E. Discussion of results

Turning to the tuned LHCf particle spectra, we show our results in Fig. 1. Here, we show the baseline QCDCR prediction (dashed line), our obtained forward physics tune result (solid line), and its error band (shaded band) against LHCf measurements.

The pion and photon spectra show similar behavior as most of the photons come from pion decay, so we discuss them together. The pion (photon) spectra can be found in

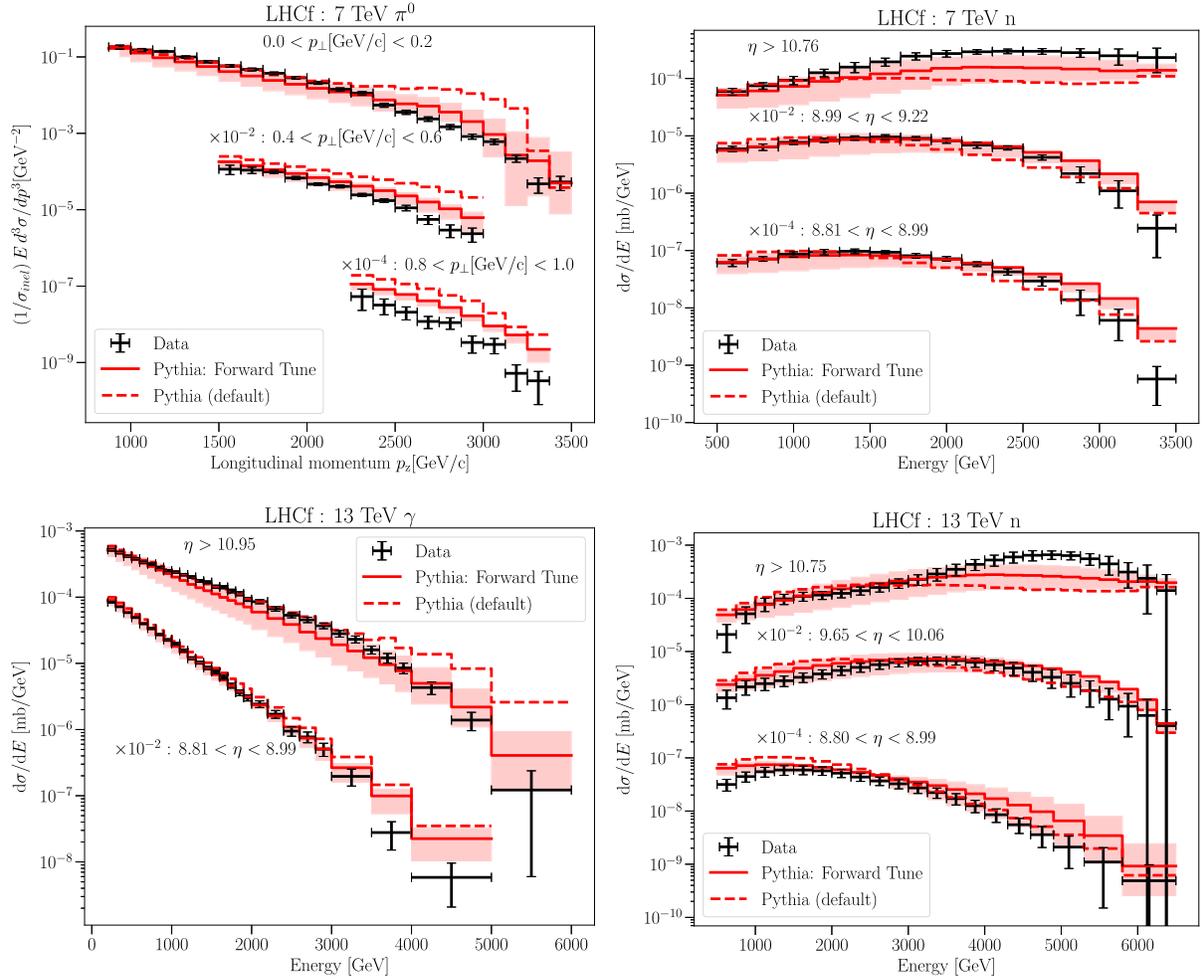


FIG. 1. LHCf measurements of pions (upper left), photons (lower left), and neutrons at $\sqrt{s} = 7$ TeV (upper right) and $\sqrt{s} = 13$ TeV (lower right) as compared to our tune and the default PYTHIA prediction. The solid curve is the central prediction of our forward tune, and the shaded region defines our uncertainty band. The dashed curve is the default PYTHIA prediction, and the black error bars are the measured data points. The text near the curves indicates the η (or p_{\perp}) of the curve, as well as a multiplicative shift that we use for display purposes.

the upper (lower) left panel in Fig. 1. For the pion spectra, two p_{\perp} bins are excluded for display purposes, but this discussion also applies to them. We see that the default configuration predicts too many particles, with a pronounced excess for the most forward bins at high p_z, E . Our tune greatly reduces this excess at $E_{\pi^0, \gamma} \approx 3$ TeV energies, which can in large part be attributed to the removal of the popcorn mechanism on the beam remnant. At smaller momenta $p_z \sim$ TeV, the default curves do better for the largest η (smallest p_{\perp}) pion (photon) bins, but this is a small improvement compared to the excess that is reduced in other bins. For most curves, our uncertainty band envelopes most of the data points with the exception of some curves which are still in tension (e.g., pions with $0.8 < p_{\perp} [\text{GeV}/c] < 1.0$).

The predicted and measured neutron spectra are shown in the upper (lower) right panels in Fig. 1 for LHC center-of-mass energies of 7 TeV (13 TeV). For the $\sqrt{s} = 13$ TeV

neutrons, we show three of six representative η bins. The most clear deficiency of the default PYTHIA prediction is an underproduction of neutrons with $\eta > 10.76$, resulting in a spectrum that peaks at lower energies relative to the measured peak. As with the pions and photons, by disabling the popcorn mechanism on the beam remnant, our tune can address this deficiency at both LHC energies by producing more hard neutrons.

We show the impact of our tune on the forward η meson distribution as measured by LHCf in the upper left panel in Fig. 2. The default PYTHIA configuration overpredicts the number of η mesons by almost 2 orders of magnitude for some bins. While we did not tune to this dataset at all, we see that our tune improves on this by producing less η 's.

In the remaining panels in Fig. 2, we show our tune as compared to the rapidity distribution of charged particles at CMS and TOTEM's T2 telescope [25–27], measurements

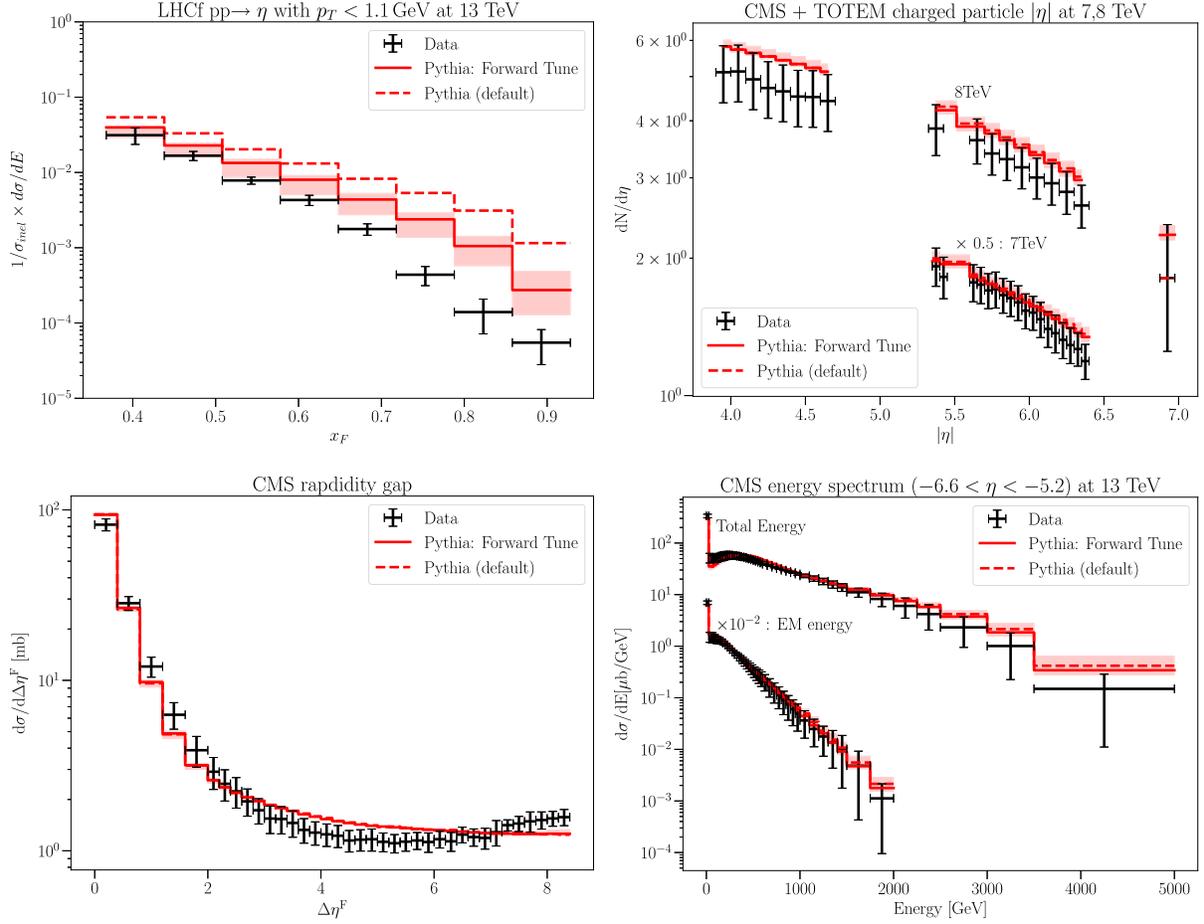


FIG. 2. In the upper left panel, we show the η meson distribution as measured by LHCf [23]. Our tune (solid red line) improves on this distribution, as compared to the default configuration (dashed red line) [16]. In the remaining panels, we compare our tune to more central measurements. In particular, we show CMS and TOTEM charged particle pseudorapidity distribution [25–27] (upper right), CMS rapidity gap measurement [25] (lower left), and CMS energy spectrum from $-6.6 < \eta < -5.2$ [28] (lower right). These measurements are expected to be the most sensitive to our tuning parameters, and we see a small deviation from the default prediction.

of the rapidity gap distribution at CMS [25], and the energy spectrum measured by CMS’ CASTOR calorimeter at $-6.6 < \eta < -5.2$ [28]. There is also a similar rapidity gap analysis from ATLAS [29] that we checked but do not show, which, in addition to the CMS rapidity gap, was used to tune the parameters in PYTHIA associated with the modeling of the elastic, diffractive, and total cross section [30]. Besides LHCf, these measurements are the most sensitive to the beam remnant, with TOTEM and CASTOR covering $\eta \sim 5 \dots 7$, respectively. If our tune had an impact on central physics, we would expect to see an effect on the predicted spectra at these experiments, with a subleading impact on predictions of the rapidity gap at CMS and ATLAS. In all cases, we find a negligible difference between our forward physics tune and the default PYTHIA prediction, while our uncertainty band produces at most a 5% variation (seen in the CMS and TOTEM measurements of charged particle pseudorapidity distribution).

IV. APPLICATION AT FASER

In this section, we discuss how our tune can be applied to current and future forward physics experiments. As our tune modifies forward hadron production rates, the decay products of these hadrons will also be affected. Forward hadrons may decay into neutrinos and, as a result, produce a highly collimated intense neutrino beam along the collision axis of the LHC. Similarly, these hadrons might also decay into yet undiscovered light and weakly interacting particles. As the LHC ring curves away, these neutrinos and BSM particles will travel unimpeded down the collision axis. A new class of experiments has recently begun operating to exploit this particle flux.

One of these experiments is FASER [31], which is located along the collision axis, 480 m downstream of the ATLAS interaction point (IP), and covers $\eta \gtrsim 9$. Located at the front of the experiment is the FASER ν neutrino detector which is a $25 \text{ cm} \times 25 \text{ cm} \times 1 \text{ m}$ tungsten emulsion

detector [32,33]. The FASER detector also consists of a long-lived particle detector which searches for the decay products of BSM particles via a series of trackers and a calorimeter. The SND@LHC experiment is also currently taking data and is located 480 m from IP on the opposite side of the ATLAS as FASER [34]. SND@LHC collects off-axis neutrinos from the pp collision and covers $7.2 < \eta < 8.7$.

To fully utilize the HL-LHC era, upgrades to these experiments have been envisioned, as well as the implementation of further forward physics experiments. These proposed experiments would be located in the FPF [4,5], which is a dedicated cavern for forward physics, located 620 m from the ATLAS IP with space to host a suite of experiments. This includes three detectors aimed at studying neutrinos as well as FASER2 for long-lived particle decays and the FORMOSA experiment for millicharged particle detection.

In the following, we apply our tune to make predictions for neutrino fluxes and the dark photon search sensitivity at FASER. These predictions can, of course, also be applied for other experiments at the FPF.

A. Neutrinos

The LHC produces an intense flux of high-energy neutrinos. This was first realized in the 1980s [35], but no detection has been made until recently. The first candidates were first detected using the FASER ν pilot detector in 2021 [2] and further observed by the FASER

detector in 2023 [3]. These neutrinos are expected to originate from pion, kaon, and charm meson decays.

The first estimate of the neutrino flux was provided in Ref. [36], which takes into account both the prompt flux from charm meson decay occurring at the IP and the displaced decays of long-lived hadrons. This estimate uses a variety of Monte Carlo (MC) event generators from cosmic-ray physics (EPOS-LHC [37], Sibyll 2.3d [38], QgsJet 2.04 [39], and DpmJet 3.2019.1 [40]) as well as PYTHIA to model the hadron production at the LHC. The average and spread of these generators have then been used to define a first rough neutrino flux estimate and its uncertainty.

Using our improved forward physics tune, we make predictions for the event rate at FASER ν . For this, we use the dedicated fast simulation as introduced in Ref. [36] to model the propagation and decay of long-lived hadrons when passing through the LHC beam pipe and magnetic fields. We have updated the magnet field configuration to those used at the beginning of run 3 and use the same beam crossing angle of $160 \mu\text{rad}$ downward. We then convolute the neutrino flux obtained using PYTHIA with the interaction cross sections obtained from GENIE [41] to calculate the number of expected events in FASER ν .

Our results are shown in Fig. 3 for an integrated luminosity of 150 fb^{-1} . The left and right panels are the electron and muon neutrino spectrum, respectively. The red line is our central prediction for our forward tune, and the dashed black line is the spectrum with the default configuration of PYTHIA. The red shaded region is our uncertainty band as determined in Sec. III D. For comparison, we also

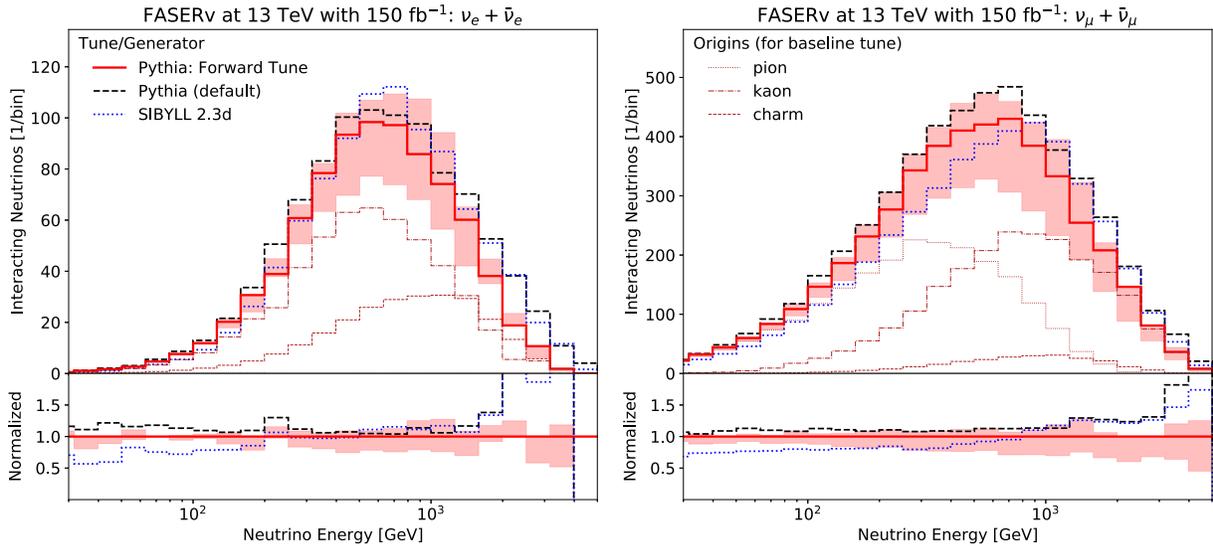


FIG. 3. Predicted neutrino energy spectrum at FASER ν for $\nu_e + \bar{\nu}_e$ (left) and $\nu_\mu + \bar{\nu}_\mu$ (right). The solid red curve is the spectrum computed using the neutrino flux from our tune, and the shaded region is our uncertainty band. The dotted, dash-dotted, and dashed red curves show the composition of the neutron flux in terms of the parent meson. For comparison, we show the interaction spectrum predicted by the default PYTHIA configuration (dashed black line) as well as the Sibyll event generator (dotted blue line). In the bottom panel in each figure, we show the ratio of the curves to our tune—our uncertainty analysis gives about a 20%–30% uncertainty on the interacting neutrino spectrum.

show the predictions from the SIBYLL event generator. In the bottom panel, we show the ratios of the curves to our tuned curve—we see that our uncertainty gives roughly a 20% uncertainty in the neutrino interaction rate.

Also indicated in Fig. 3 is composition of the neutrinos in terms of their parent mesons, shown in dotted, dash-dotted, and dashed curves for the pion, kaon, and charm meson, respectively. Clearly, the majority of electron neutrinos come from kaon decay, with a significant charm component at higher energies. Muon neutrinos, on the other hand, are dominantly produced by pion decay at lower energies and kaon decay at high energies. While PYTHIA models charm production, we note that there are ongoing efforts to provide refined predictions of forward charm production mode using perturbative QCD [42–45], some of which predict significantly enhanced charm production rates. In the regime where light hadron decays dominate the neutrino composition, the obtained flux uncertainty with our tune roughly agrees with that in Ref. [36].

We note that, currently, we include only uncertainties associated with the kinematic distribution. There could be additional sources of uncertainties associated with the flavor composition, especially the kaon to pion production fraction. Indeed, observation from astroparticle physics suggest that forward kaon production might be different than predicted by existing hadronic interaction models. Over more than two decades, cosmic-ray experiments have reported significant discrepancies between the number of observed muons in high-energy cosmic-ray air showers and

model predictions [46–48]. This observation is commonly referred to as the muon puzzle. Extensive studies have suggested that an enhanced rate of strangeness production in the forward direction could explain the discrepancy [49–51]. While forward strange measurements could shed light on this discrepancy, no attempt was made to include this in our tune due to the lack of data.

B. Dark photons

The other main purpose of FASER is the search for light long-lived particles with MeV–GeV masses [52–54]. These are, for example, motivated by dark matter and, more generally, dark sectors. One of the primary examples discussed in the literature is dark photons. The dark photon is a gauge field of a broken U(1) in the dark sector. Through kinetic mixing with the SM photon, the dark photon A' can interact with SM fields. This interaction is suppressed by the kinetic mixing parameter ϵ with an interaction Lagrangian $\mathcal{L} \supset \epsilon/2 F'^{\mu\nu} F_{\mu\nu}$, where F (F') is the field strength of the (dark) photon. For massive dark photons with $2m_e < m_{A'} < 2m_\mu$, the dark photon will primarily decay into e^+e^- . With sufficiently small ϵ , the dark photon will travel several hundred meters before decaying and could decay inside FASER, which has been designed to detect this signal.

Recently, FASER reported the first results for the search for dark photons [1]. In the probed regime, dark photons mainly come from neutral pion decay with small

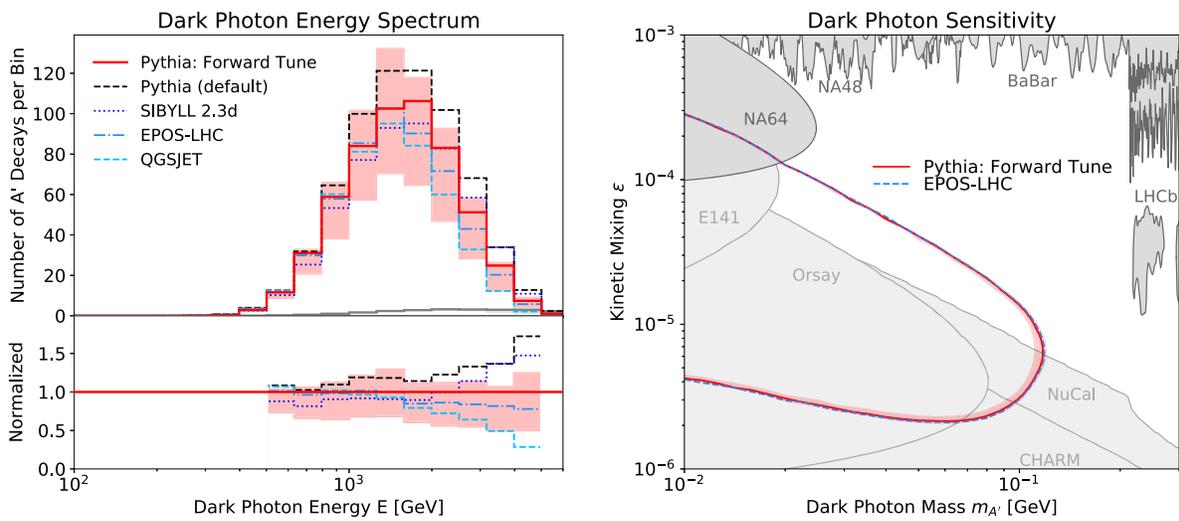


FIG. 4. Dark photon spectrum at FASER for $m_{A'}, \epsilon = 25 \text{ MeV}, 3 \times 10^{-5}$ (left) and the discovery reach for FASER using the spectrum predicted by our tune. In the left panel, we show the spectra predicted by our tune (solid red line) as well as the associated uncertainty that we calculate (shaded red). For comparison, we show the spectra predicted by the default PYTHIA configuration (dashed black line), SIBYLL (dotted blue line), EPOS-LHC (dash-dotted blue line), and QGSJET (dashed blue line). In the bottom section in the left panel, we show the ratio of the curves to our tune and see that our uncertainty imparts an $\approx 50\%$ uncertainty on the number of dark photon decays in FASER. In the right panel, we show FASER’s sensitivity in dark photon parameter space for our tune (solid red line), our associated uncertainty (shaded red), and the sensitivity predicted by EPOS-LHC for comparison (dashed blue line). While the uncertainty we calculate has a small impact on FASER’s sensitivity, see that the uncertainty is most important when FASER is limited by exposure (i.e., at small ϵ and large $m_{A'}$).

contributions from eta-meson decay and dark bremsstrahlung. The FASER Collaboration has estimated the dark photon flux using EPOS-LHC. The signal uncertainty was estimated by comparing with Sibyll and QgsJet.

We use our PYTHIA forward physics tune to model forward particle production and FORESEE [55] to then obtain the expected dark photon event rate at FASER. The left panel in Fig. 4 shows the energy spectrum of dark photons decaying in FASER during run 3 with 150 fb^{-1} integrated luminosity for $m_{A'} = 25 \text{ MeV}$ and $\epsilon = 3 \times 10^{-5}$. This point lies at the edge of the previously excluded region. The red curve is our main prediction, and the shaded band is the error band. The bottom panel shows the ratio of the curves to our central prediction and shows that our uncertainty is roughly 30%. For comparison, we also show the dark photon spectrum from the default PYTHIA configuration (dashed black line) and the prediction from Sibyll, EPOS-LHC, and QgsJet in dotted, dash-dotted, and dashed blue curves, respectively. We can see that the predictions from these other generators are consistent with our prediction. We note that our uncertainty is slightly larger than the uncertainty obtained by comparing generators at low energy and similar at higher energy.

The right panel shows the FASER sensitivity for run 3 with 150 fb^{-1} in the dark photon parameter space spanned by ϵ and $m_{A'}$. The gray shaded areas are excluded by existing experiments (from above by prompt resonance searches and from below by long-lived particle searches in beam dumps) as obtained from DarkCast [56,57]. The constraints shown in light gray are obtained by recasting experimental analyses, while dark gray bounds were reported directly by the experiments. Using our tune, we draw our expected sensitivity contour in red with our uncertainty as the shaded contour and compare with the sensitivity contour as calculated with EPOS-LHC in dashed blue. We find that the sensitivity calculated with each configuration is comparable. We also note that the overall effect of the flux uncertainty on the sensitivity reach is small. This is due to an exponential (ϵ^4) suppression of the event rate at large (small) ϵ .

V. CONCLUSION

In recent years, a new set of experiments has begun their operation in the forward direction of the LHC, with the purpose of observing and studying collider neutrinos as well as searching for light long-lived particles. This emerging forward neutrino and particle search program requires precise predictions of the anticipated particle fluxes. Currently, forward particle production is often simulated using specialized MC event generators developed for cosmic-ray physics, such as EPOS-LHC, QgsJet, and Sibyll. Additionally, multipurpose event generators like PYTHIA can also be utilized. However, it has been noticed that the corresponding predicted spectra exhibit some

discrepancies when compared to the measured flux obtained from the LHCf experiment.

This paper addresses this issue by introducing a new dedicated forward tune for PYTHIA, specifically designed for forward physics studies at the LHC. This newly proposed tune is based on the QCDCR scenario introduced in Ref. [16], offers a more adaptable approach for modeling beam remnant hadronization, and is tuned to match the available forward particle spectra measured by LHCf. A comprehensive list of the relevant parameters and their corresponding values can be found in Table I. We also explored an alternative tune based on the well-established Monash configuration utilizing the default CR scenario. However, we found that this alternative tune exhibits a poorer agreement with LHCf data compared to the QCDCR-based approach, as discussed in the Appendix.

When fine-tuning event generators, the process currently lacks a well-established method for quantifying and incorporating measures of uncertainty. In addition to our fit, we also provide an uncertainty in a data-driven way for the first time. What has sometimes been done is to take the spread in event generators' predictions to define an uncertainty band on the true particle distribution. In this paper, we vary the relevant tuning parameter σ around the best fit such that 68% of the data points are captured. This band can then be used for further applications to study the impact of flux uncertainties.

To demonstrate an application of our tune, we also show its impact on the predicted neutrino and dark photon sensitivity. A precise understanding of the neutrino flux that better agrees with forward physics data is important to study TeV neutrino interactions, and an improved understanding of the dark photon flux will increase experiments' search sensitivity. Our tune also provides a means of understanding the flux uncertainty in each case. For both cases, we find that our tune is consistent with the Sibyll, EPOS-LHC, and QgsJet generators, and that our uncertainty band is a bit wider than the spread of these generators' predictions.

In conclusion, our forward tune of PYTHIA enables enhanced precision in the exploration of forward physics phenomena. Our approach presents a data-guided mechanism for honing the neutrino flux and its associated uncertainty. By gaining better control over the uncertainty in neutrino flux, it opens the gateway to improved investigations, including a refined modeling of neutrino production through hadron decay [58], exploration of sterile neutrino production, and a deeper understanding of neutrino interactions within experiments designed to unveil proton structure [59], and potential avenues toward uncovering new signatures of physics.

ACKNOWLEDGMENTS

We thank Aki Ariga, Tomoko Ariga, Eugenio Berti, Andy Buckley, Anatoli Fedynitch, Jonathan Feng, Hiroaki Menjo, Kenneth Österberg, Stefan Hoeche, Tanguy Pierog,

and Christophe Royon for many fruitful discussions. We are grateful to the authors and maintainers of many open-source software packages, including SCIKIT-HEP [60]. This work utilized the infrastructure for high-performance and high-throughput computing, research data storage and analysis, and scientific software tool integration built, operated, and updated by the Research Cyberinfrastructure Center (RCIC) at the University of California, Irvine (UCI). The RCIC provides cluster-based systems, application software, and scalable storage to directly support the UCI research community. The work of M. F. was supported by National Science Foundation (NSF) Grant No. PHY-2210283 and was also supported by NSF Graduate Research Fellowship Grant No. DGE-1839285. F. K. acknowledges support by the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy—EXC 2121 Quantum Universe—390833306. T. S. has been supported by the Swedish Research Council, Contract No. 2016-05996.

APPENDIX: ALTERNATE MONASH TUNE

Here, we discuss and show the results of the alternate tune which is based off the well-known Monash tune to

central physics, which we provide for comparison purposes. We show our fitting results in Table II and our fitted spectra against LHCf data in Fig. 5. While we find comparable tuning parameters for this Monash-based tune as our main tune, the QCDCR configuration from Ref. [16] proves to be an important feature for our tuning purposes. We find that the weighted and reduced χ^2 for our primary tune is 2.74 and 3.28 for the Monash-based tune.

While the Monash-based tune has some same advantages of our primary tune, there are some clear deficiencies. In particular, the photon spectra show a significant underproduction of forward photons with $E \lesssim 3$ TeV—a similar effect same can be seen for relatively softer pions, while the Monash-based tune does better at higher energies. We find that the photon and pion contribution to the reduced χ^2 for our primary (Monash-based) tune is 0.43 and 1.69 (1.45 and 0.97), respectively. A further deficiency can be seen in the $\eta > 10.76$ neutron spectra, particularly for $\sqrt{s} = 7$ TeV—the Monash tune does not address the shape of the neutron spectra as well as our primary tune does. We find that the contribution to the χ^2 from the neutron spectra for $\sqrt{s} = 7$ and 13 TeV for our primary (Monash-based) tune is 0.46 and 0.17 (0.53 and 0.33), respectively.

TABLE II. The main PYTHIA parameters studied in this article, their default parameters in the Monash tune, and their values in the Monash-based tune obtained in this study. The last column shows the uncertainty range for $\sigma_{\text{soft}} = \sigma_{\text{remn}}$ as discussed in Sec. III D.

Full name	Shorthand	Baseline (Monash)	Forward tune	Uncertainty
BeamRemnants:dampPopcorn	d_{pop}	1	0	
BeamRemnants:hardRemnantBaryon	f_{remn}	Off	On	
BeamRemnants:aRemnantBaryon	a_{remn}	...	0.68	
BeamRemnants:bRemnantBaryon	b_{remn}	...	1.22	
BeamRemnants:primordialKTsoft	σ_{soft}	0.9	0.56	0.2...1.42
BeamRemnants:primordialKTthard	σ_{hard}	1.8	1.8	
BeamRemnants:halfScaleForKT	Q_{half}	1.5	10	
BeamRemnants:halfMassForKT	m_{half}	1	1	
BeamRemnants:primordialKTremnant	σ_{remn}	0.4	0.56	0.22...1.42

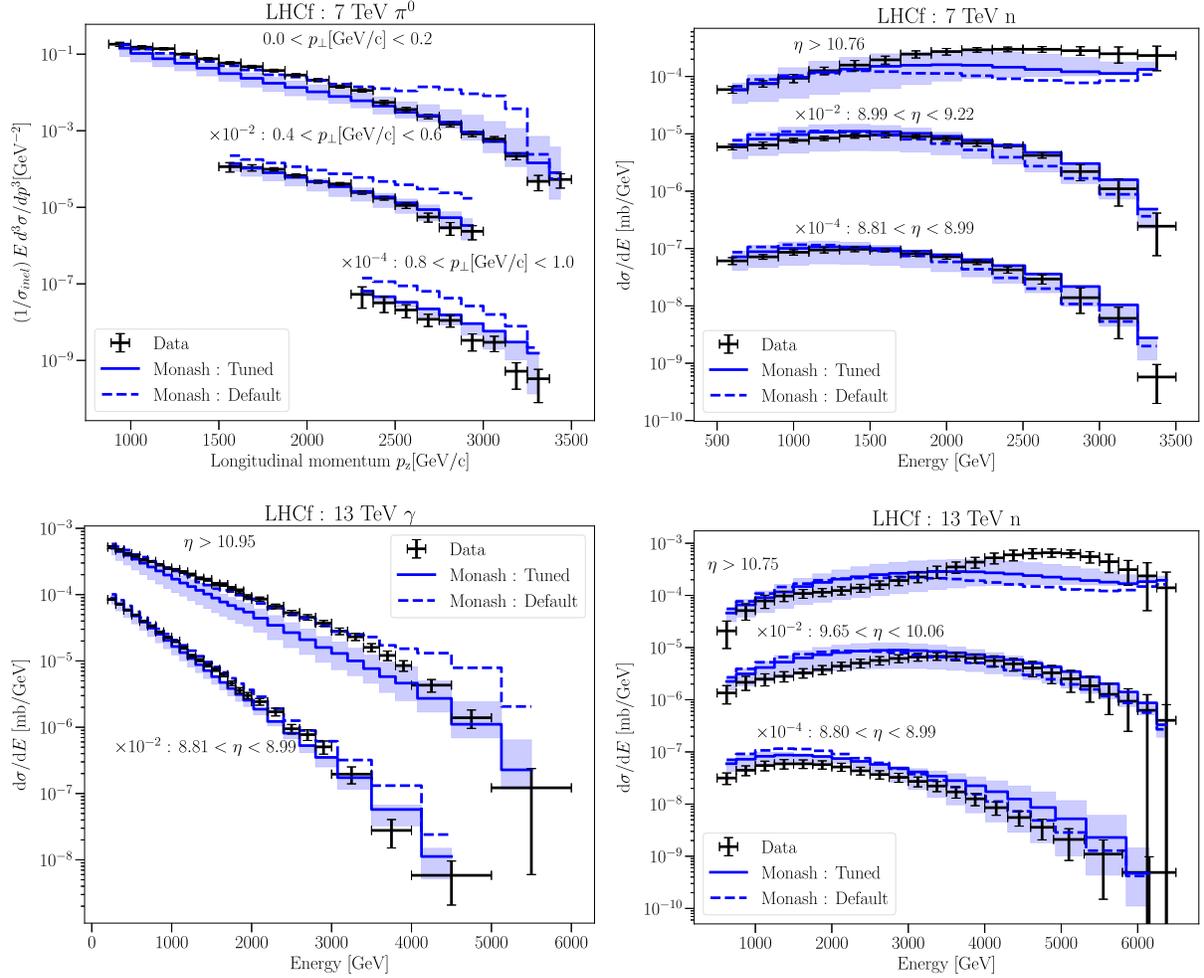


FIG. 5. LHCf spectra as compared to an alternate tune that we explored based on the Monash tune. The LHCf measurements of pions (upper left), photons (lower left), and neutrons at $\sqrt{s} = 7$ TeV (upper right) and $\sqrt{s} = 13$ TeV (lower right) as compared to our alternate, Monash-based tune and the default PYTHIA prediction. The solid curve is the central prediction of our Monash-based tune, and the shaded region defines our uncertainty band. The dashed curve is the default PYTHIA prediction, and the black error bars are the measured data points. The text near the curves indicates the η (or p_{\perp}) of the curve, as well as a multiplicative shift that we use for display purposes.

- [1] H. Abreu *et al.* (FASER Collaboration), Search for dark photons with the FASER detector at the LHC, [arXiv: 2308.05587](#).
- [2] H. Abreu *et al.* (FASER Collaboration), First neutrino interaction candidates at the LHC, *Phys. Rev. D* **104**, L091101 (2021).
- [3] H. Abreu *et al.* (FASER Collaboration), First direct observation of collider neutrinos with FASER at the LHC, *Phys. Rev. Lett.* **131**, 031801 (2023).
- [4] L. A. Anchordoqui *et al.*, The forward physics facility: Sites, experiments, and physics potential, *Phys. Rep.* **968**, 1 (2022).
- [5] J. L. Feng *et al.*, The forward physics facility at the high-luminosity LHC, *J. Phys. G* **50**, 030501 (2023).
- [6] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [7] C. Bierlich *et al.*, A comprehensive guide to the physics and usage of PYTHIA 8.3, *SciPost Phys. Codebases* **8** (2022), 10.21468/SciPostPhysCodeb.8.
- [8] P. Skands, S. Carrazza, and J. Rojo, Tuning PYTHIA 8.1: The Monash 2013 Tune, *Eur. Phys. J. C* **74**, 3024 (2014).
- [9] O. Adriani *et al.* (LHCf Collaboration), Measurement of very forward neutron energy spectra for 7 TeV proton–proton collisions at the Large Hadron Collider, *Phys. Lett. B* **750**, 360 (2015).

- [10] O. Adriani *et al.* (LHCf Collaboration), Measurement of energy flow, cross section and average inelasticity of forward neutrons produced in $\sqrt{s} = 13$ TeV proton-proton collisions with the LHCf Arm2 detector, *J. High Energy Phys.* **07** (2020) 016.
- [11] O. Adriani *et al.* (LHCf Collaboration), Measurements of longitudinal and transverse momentum distributions for neutral pions in the forward-rapidity region with the LHCf detector, *Phys. Rev. D* **94**, 032007 (2016).
- [12] O. Adriani *et al.* (LHCf Collaboration), Measurement of forward photon production cross-section in proton-proton collisions at $\sqrt{s} = 13$ TeV with the LHCf detector, *Phys. Lett. B* **780**, 233 (2018).
- [13] R. P. Feynman, Very high-energy collisions of hadrons, *Phys. Rev. Lett.* **23**, 1415 (1969).
- [14] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand, Parton fragmentation and string dynamics, *Phys. Rep.* **97**, 31 (1983).
- [15] T. Sjöstrand and M. van Zijl, A multiple interaction model for the event structure in hadron collisions, *Phys. Rev. D* **36**, 2019 (1987).
- [16] J. R. Christiansen and P. Z. Skands, String formation beyond leading colour, *J. High Energy Phys.* **08** (2015) 003.
- [17] A. M. Sirunyan *et al.* (CMS Collaboration), Production of Λ_c^+ baryons in proton-proton and lead-lead collisions at $\sqrt{s_{NN}} = 5.02$ TeV, *Phys. Lett. B* **803**, 135328 (2020).
- [18] S. Acharya *et al.* (ALICE Collaboration), Λ_c^+ production in pp and in p -Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, *Phys. Rev. C* **104**, 054905 (2021).
- [19] T. Sjöstrand, The merging of jets, *Phys. Lett.* **142B**, 420 (1984).
- [20] G. Ingelman and P. E. Schlein, Jet structure in high mass diffractive scattering, *Phys. Lett.* **152B**, 256 (1985).
- [21] R. D. Ball, V. Bertone, S. Carrazza, L. Del Debbio, S. Forte, A. Guffanti, N. P. Hartland, and J. Rojo (NNPDF Collaboration), Parton distributions with QED corrections, *Nucl. Phys.* **B877**, 290 (2013).
- [22] O. Adriani *et al.* (LHCf Collaboration), Measurement of photons and neutral pions in the very forward region of LHC, technical design report of the LHCf experiment, <http://cds.cern.ch/record/926196>.
- [23] O. Adriani *et al.*, Measurement of the forward η meson production rate in p - p collisions at $\sqrt{s} = 13$ TeV with the LHCf-Arm2 detector, *J. High Energy Phys.* **10** (2023) 169.
- [24] M. Krishnamoorthy, H. Schulz, X. Ju, W. Wang, S. Leyffer, Z. Marshall, S. Mrenna, J. Müller, and J. B. Kowalkowski, Apprentice for event generator tuning, *EPJ Web Conf.* **251**, 03060 (2021).
- [25] V. Khachatryan *et al.* (CMS Collaboration), Measurement of diffraction dissociation cross sections in pp collisions at $\sqrt{s} = 7$ TeV, *Phys. Rev. D* **92**, 012003 (2015).
- [26] S. Chatrchyan *et al.* (CMS, TOTEM Collaborations), Measurement of pseudorapidity distributions of charged particles in proton-proton collisions at $\sqrt{s} = 8$ TeV by the CMS and TOTEM experiments, *Eur. Phys. J. C* **74**, 3053 (2014).
- [27] G. Antchev *et al.* (TOTEM Collaboration), Measurement of the forward charged particle pseudorapidity density in pp collisions at $\sqrt{s} = 7$ TeV with the TOTEM experiment, *Europhys. Lett.* **98**, 31002 (2012).
- [28] A. M. Sirunyan *et al.* (CMS Collaboration), Measurement of the inclusive energy spectrum in the very forward direction in proton-proton collisions at $\sqrt{s} = 13$ TeV, *J. High Energy Phys.* **08** (2017) 046.
- [29] G. Aad *et al.* (ATLAS Collaboration), Rapidity gap cross sections measured with the ATLAS detector in pp collisions at $\sqrt{s} = 7$ TeV, *Eur. Phys. J. C* **72**, 1926 (2012).
- [30] C. O. Rasmussen and T. Sjöstrand, Models for total, elastic and diffractive cross sections, *Eur. Phys. J. C* **78**, 461 (2018).
- [31] H. Abreu *et al.* (FASER Collaboration), The FASER detector, [arXiv:2207.11427](https://arxiv.org/abs/2207.11427).
- [32] H. Abreu *et al.* (FASER Collaboration), Detecting and studying high-energy collider neutrinos with FASER at the LHC, *Eur. Phys. J. C* **80**, 61 (2020).
- [33] H. Abreu *et al.* (FASER Collaboration), Technical proposal: FASERnu, [arXiv:2001.03073](https://arxiv.org/abs/2001.03073).
- [34] G. Acampora *et al.* (SND@LHC Collaboration), SND@LHC: The scattering and neutrino detector at the LHC, [arXiv:2210.02784](https://arxiv.org/abs/2210.02784).
- [35] A. De Rujula, E. Fernandez, and J. J. Gomez-Cadenas, Neutrino fluxes at future hadron colliders, *Nucl. Phys.* **B405**, 80 (1993).
- [36] F. Kling and L. J. Nevay, Forward neutrino fluxes at the LHC, *Phys. Rev. D* **104**, 113008 (2021).
- [37] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner, EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider, *Phys. Rev. C* **92**, 034906 (2015).
- [38] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, Hadronic interaction model sibyll 2.3d and extensive air showers, *Phys. Rev. D* **102**, 063002 (2020).
- [39] S. Ostapchenko, Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QgsJet-II model, *Phys. Rev. D* **83**, 014018 (2011).
- [40] A. Fedynitch, Cascade equations and hadronic interactions at very high energies. Ph.D. thesis, KIT, Karlsruhe, Department of Physics, 2015.
- [41] C. Andreopoulos *et al.*, The GENIE neutrino Monte Carlo generator, *Nucl. Instrum. Methods Phys. Res., Sect. A* **614**, 87 (2010).
- [42] W. Bai, M. Diwan, M. V. Garzelli, Y. S. Jeong, and M. H. Reno, Far-forward neutrinos at the Large Hadron Collider, *J. High Energy Phys.* **06** (2020) 032.
- [43] R. Maciula and A. Szczurek, Far-forward production of charm mesons and neutrinos at forward physics facilities at the LHC and the intrinsic charm in the proton, *Phys. Rev. D* **107**, 034002 (2023).
- [44] A. Bhattacharya, F. Kling, I. Sarcevic, and A. M. Stasto, Forward neutrinos from charm at Large Hadron Collider, [arXiv:2306.01578](https://arxiv.org/abs/2306.01578).
- [45] L. Buonocore, F. Kling, L. Rottoli, and J. Sominka, Predictions for neutrinos and new physics from forward heavy hadron production at the LHC, [arXiv:2309.12793](https://arxiv.org/abs/2309.12793).
- [46] A. Aab *et al.* (Pierre Auger Collaboration), Muons in air showers at the Pierre Auger Observatory: Mean number in highly inclined events, *Phys. Rev. D* **91**, 032003 (2015); **91**, 059901(E) (2015).
- [47] A. Aab *et al.* (Pierre Auger Collaboration), Testing hadronic interactions at ultrahigh energies with air showers measured

- by the Pierre Auger Observatory, *Phys. Rev. Lett.* **117**, 192001 (2016).
- [48] D. Soldin (EAS-MSU, IceCube, KASCADE-Grande, NEVOD-DECOR, Pierre Auger, SUGAR, Telescope Array, Yakutsk EAS Array Collaborations), Update on the combined analysis of muon measurements from nine air shower experiments, *Proc. Sci. ICRC2021* (2021) 349 [arXiv:2108.08341].
- [49] J. Allen and G. Farrar, Testing models of new physics with UHE air shower observations, in *Proceedings of the 33rd International Cosmic Ray Conference* (2013), p. 1182, arXiv:1307.7131.
- [50] L. A. Anchordoqui, H. Goldberg, and T. J. Weiler, Strange fireball as an explanation of the muon excess in Auger data, *Phys. Rev. D* **95**, 063005 (2017).
- [51] J. Albrecht *et al.*, The muon puzzle in cosmic-ray induced air showers and its connection to the Large Hadron Collider, *Astrophys. Space Sci.* **367**, 27 (2022).
- [52] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, ForwArd Search ExpeRiment at the LHC, *Phys. Rev. D* **97**, 035001 (2018).
- [53] A. Ariga *et al.* (FASER Collaboration), Letter of intent for FASER: ForwArd Search ExpeRiment at the LHC, arXiv:1811.10243.
- [54] A. Ariga *et al.* (FASER Collaboration), FASER's physics reach for long-lived particles, *Phys. Rev. D* **99**, 095011 (2019).
- [55] F. Kling and S. Trojanowski, Forward experiment sensitivity estimator for the LHC and future hadron colliders, *Phys. Rev. D* **104**, 035012 (2021).
- [56] P. Ilten, Y. Soreq, M. Williams, and W. Xue, Serendipity in dark photon searches, *J. High Energy Phys.* **06** (2018) 004.
- [57] C. Baruch, P. Ilten, Y. Soreq, and M. Williams, Axial vectors in DarkCast, *J. High Energy Phys.* **11** (2022) 124.
- [58] F. Kling, T. Mäkelä, and S. Trojanowski, Investigating the fluxes and physics potential of LHC neutrino experiments, *Phys. Rev. D* **108**, 095020 (2023).
- [59] J. M. Cruz-Martinez, M. Fieg, T. Giani, P. Krack, T. Mäkelä, T. Rabemananjara, and J. Rojo, The LHC as a neutrino-ion collider, arXiv:2309.09581.
- [60] E. Rodrigues *et al.*, The SCIKIT HEP project—A overview and prospects, *EPJ Web Conf.* **245**, 06028 (2020).