First Search for Quirks at the LHC with FASER

Roshan Mammen Abraham,[1] Xiaocong Ai,[2] Saul Alonso Monsalve,[3] John Anders,[4] Emma Kate Anderson,[5] Akitaka Ariga,[6,7] Tomoko Ariga,[8] Jeremy Atkinson,[6] Florian U. Bernlochner,[9] Jianming Bian,[1] Tobias Boeckh,[9] Eliot Bornand,[10] Jamie Boyd,[5] Lydia Brenner,[11] Angela Burger,[12] Franck Cadoux,[10] Haichuan Cao,[13] Roberto Cardella,[10] David W. Casper,[1] Charlotte Cavanagh,[3] Shiyang Chen,[13] Xin Chen,[13] Xing Cheng,[13] Dhruv Chouhan,[9] Andrea Coccaro,[14] Fabio Cufino,[3] Stephane Débieux,[10] Ansh Desai,[15] Sergey Dmitrievsky,[16] Radu Dobre,[17] Monica D'Onofrio,[4] Sinead Eley,[4] Yannick Favre,[10] Jonathan L. Feng,[1] Carlo Alberto Fenoglio,[10] Didier Ferrere,[10] Max Fieg,[18] Wissal Filali,[9] Elena Firu,[17] Haruhi Fujimori,[7] Edward Galantay,[10,5] Stephen Gibson,[19] Sergio Gonzalez-Sevilla,[10] Yuri Gornushkin,[16] Yotam Granov,[20] Jinjing Gu,[13] Carl Gwilliam,[4] Elie Hammou,[11] Daiki Hayakawa,[7] Michael Holzbock,[5] Shih-Chieh Hsu,[21] Zhen Hu,[13] Giuseppe Iacobucci,[10] Tomohiro Inada,[8] Luca Iodice,[10] Sune Jakobsen,[5] Cesar Jesus-Valls,[5] Arash Jofrehei,[10] Hans Joos,[5] Enrique Kajomovitz,[20] Alex Keyken,[19] Felix Kling,[9] Daniela Köck,[15] Pantelis Kontaxakis,[10] Jelle Koorn,[11] Umut Kose,[3] Peter Krack,[11] Susanne Kuehn,[5] Thanushan Kugathasan,[10] Sebastian Laudage,[9] Lorne Levinson,[22] Botao Li,[3] Jinmian Li,[23] Jiaxi Liu,[1] Jinfeng Liu,[13] Yi Liu,[2] Margaret S. Lutz,[5] Joern Mahlstedt,[9] Toni Mäkelä,[1] Yasuhiro Maruya,[8] Anna Mascellani,[3] Lawson McCoy,[1] Josh McFayden,[24] Andrea Pizarro Medina,[10] Hiroaki Menjo,[25] Théo Moretti,[10] Toshiyuki Nakano,[25] Laurie Nevay,[5] Ken Ohashi,[7] Hidetoshi Otono,[8] Lorenzo Paolozzi,[10,5] Annabelle Parry,[4,5] Pawan Pawan,[4] Junle Pei,[26] Brian Petersen,[5] Titi Preda,[17] Markus Prim,[9] Junkai Qin,[13] Michaela Queitsch-Maitland,[27] Juan Rojo,[11] Hiroki Rokujo,[8] André Rubbia,[3] Osamu Sato,[25] Paola Scampoli,[28,6] Kristof Schmieden,[9] Matthias Schott,[9] Cristiano Sebastiani,[5] Anna Sfyrla,[10] Davide Sgalaberna,[3] Mansoora Shamim,[5] Yosuke Takubo,[29] Kakeru Tanaka,[8] Simon Thor,[3] Eric Torrence,[15] Serhan Tufanli,[6] Oscar Ivan Valdes Martinez,[27] Svetlana Vasina,[16] Emanuele Villa,[3] Benedikt Vormwald,[5] Chi Wang,[13] Yuxiao Wang,[13] Eli Welch,[1] Aaron White,[5] Monika Wielers,[30] Benjamin James Wilson,[27] Zhongyi Wu,[1] Yue Xu,[21] Heng Yang,[13] Lekai Yao,[13] Daichi Yoshikawa,[8] Stefano Zambito,[10] Shunliang Zhang,[13] Yuxuan Zhang,[13] Xingyu Zhao,[3] and Zijian Zhao[13]

(FASER Collaboration)

[1]*Department of Physics and Astronomy, University of California, Irvine, California 92697-4575, USA*
[2]*School of Physics, Zhengzhou University, Zhengzhou 450001, China*
[3]*Institute for Particle Physics, ETH Zürich, Zürich 8093, Switzerland*
[4]*University of Liverpool, Liverpool L69 3BX, United Kingdom*
[5]*CERN, CH-1211 Geneva 23, Switzerland*
[6]*Albert Einstein Center for Fundamental Physics, Laboratory for High Energy Physics, University of Bern, Sidlerstrasse 5, CH-3012 Bern, Switzerland*
[7]*Department of Physics, Chiba University, 1-33 Yayoi-cho Inage-ku, 263-8522 Chiba, Japan*
[8]*Kyushu University, 744 Motooka, Nishi-ku, 819-0395 Fukuoka, Japan*
[9]*Universität Bonn, Regina-Pacis-Weg 3, D-53113 Bonn, Germany*
[10]*Département de Physique Nucléaire et Corpusculaire, University of Geneva, CH-1211 Geneva 4, Switzerland*
[11]*Nikhef National Institute for Subatomic Physics, Science Park 105, 1098 XG Amsterdam, Netherlands*
[12]*L2IT, Université de Toulouse, CNRS/IN2P3, UPS, Toulouse, France*
[13]*Department of Physics, Tsinghua University, Beijing, China*
[14]*INFN Sezione di Genova, Via Dodecaneso, 33–16146, Genova, Italy*
[15]*University of Oregon, Eugene, Oregon 97403, USA*
[16]*Affiliated with an international laboratory covered by a cooperation agreement with CERN*
[17]*Institute of Space Science—INFLPR Subsidiary, Bucharest, Romania*
[18]*Theoretical Physics Division, Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*
[19]*Royal Holloway, University of London, Egham, TW20 0EX, United Kingdom*
[20]*Department of Physics and Astronomy, Technion—Israel Institute of Technology, Haifa 32000, Israel*
[21]*Department of Physics, University of Washington, PO Box 351560, Seattle, Washington 98195-1460, USA*
[22]*Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 76100, Israel*
[23]*College of Physics, Sichuan University, Chengdu 610065, China*
[24]*Department of Physics and Astronomy, University of Sussex, Sussex House, Falmer, Brighton, BN1 9RH, United Kingdom*
[25]*Nagoya University, Furo-cho, Chikusa-ku, Nagoya 464-8602, Japan*

[26]Institute of Physics, Henan Academy of Sciences, Zhengzhou 450046, China
[27]University of Manchester, School of Physics and Astronomy, Schuster Building, Oxford Rd, Manchester M13 9PL, United Kingdom
[28]Dipartimento di Fisica "Ettore Pancini," Università di Napoli Federico II,
Complesso Universitario di Monte S. Angelo, I-80126 Napoli, Italy
[29]National Institute of Technology (KOSEN), Niihama College, 7-1, Yakumo-cho Niihama, 792-0805 Ehime, Japan
[30]Particle Physics Department, STFC Rutherford Appleton Laboratory, Harwell Campus, Didcot, OX11 0QX, United Kingdom

(Received 28 July 2026; accepted 23 September 2026; published 30 October 2026)

We present the first LHC search for quirks, particles with standard model and QCD-like infracolor (IC) interactions. The analysis uses data collected in 2022–2024 by FASER, with a 186 fb^{-1} integrated luminosity of proton-proton collisions at a center-of-mass energy $\sqrt{s} = 13.6$ TeV. Scintillator charge and timing information are used to search for slow pairs of quirks with no QCD color and electric charge ± 1. No events are seen in this nearly background-free search, resulting in the first exclusion of quirks with masses above the weak scale for IC confinement scales in the broad range of $300\ \text{eV} < \Lambda < 100\ \text{keV}$.

DOI: 10.1103/z156-xk5y

The standard model (SM) of particle physics is the current framework to describe the fundamental building blocks of the universe and their interactions. However, it faces many theoretical and experimental challenges, including its failure to resolve the gauge hierarchy problem or provide a viable particle candidate for cold dark matter. A generic extension of the SM can be made by adding a new gauge group. While the Abelian gauge groups have been thoroughly investigated, non-Abelian groups are much less studied, despite being more generic. Such hidden non-Abelian sectors are further motivated by neutral naturalness models [1–3], where noncolored top partners, charged under a hidden gauge group, can ameliorate the gauge hierarchy problem.

Quirks [4] are hypothetical matter particles charged under infracolor (IC), a new, QCD-like gauge group. IC models and quirks are motivated by the facts that IC glueballs are viable dark matter candidates [5,6] and the framework arises in attempts to resolve the gauge hierarchy problem [2,7]. Despite this motivation, there have been few direct searches for quirks at colliders. The D0 Collaboration searched for quirks at the Tevatron [8], but there has not yet been any dedicated search for quirks at the Large Hadron Collider (LHC), although there have been a number of proposed search strategies for ATLAS and CMS [9–15] and LHCb [16].

For this Letter, following the benchmark model studied in Ref. [17], we consider the IC gauge group $\text{SU}(N_{\text{IC}})$ with a single quirk flavor in the $(1, 1, -1, N_{\text{IC}})$ representation of $\text{SU}_{\text{C}}(3) \times \text{SU}_{\text{L}}(2) \times \text{U}_{\text{Y}}(1) \times \text{SU}(N_{\text{IC}})$, corresponding to the same SM gauge charges as right-handed charged leptons, with no SM color and electric charge -1. We take the number of infracolors to be $N_{\text{IC}} = 2$. This quirk model is, then, completely specified by the quirk mass $m_{\mathcal{Q}}$ and the IC confinement scale Λ. Quirks with SM color are not considered here, given the complexity of including QCD hadron emission, but they can be probed with a similar analysis.

Quirk-antiquirk pairs are dominantly produced at the LHC through Drell-Yan processes mediated by a photon or Z boson. At leading order, the quirk pair is produced with no transverse momentum (p_{T}) and thus close to the beam collision axis. For $\Lambda \ll m_{\mathcal{Q}}$, quirk pairs remain bound to each other via the IC force and oscillate around their common center of mass, while their high mass does not allow for the creation of new quirk pairs out of the vacuum. The oscillation amplitude is approximately $1\ \text{cm} \times (m_{\mathcal{Q}}/100\ \text{GeV})(1\ \text{keV}/\Lambda)^2$. The pair of quirks can travel long distances leaving exotic electromagnetic traces. Experimentally, quirk pairs can be identified as forward, long-lived signatures with trajectories different from those of SM particles. In particular, their weak-scale masses allow them to be produced with speeds significantly below the speed of light c. The IC force redirects part of each quirk's velocity into transverse oscillations, further reducing its longitudinal velocity. The resulting slow speed provides a powerful discriminant in the forward direction, as discussed in Ref. [17].

The FASER detector [18] at the LHC is well suited to search for quirk signatures [17,19]. It is located 480 m downstream of the ATLAS experiment's interaction point, aligned with the beam collision axis. The LHC dipole magnets and the material between ATLAS and FASER, including more than 100 m of rock, shield FASER from most SM particles. Muons and neutrinos are the only SM particles reaching FASER, making it a suitable environment to search for particles beyond the SM.

As shown in Fig. 1, the FASER detector consists of a 1.1-ton FASERν emulsion detector, a decay volume, three

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.

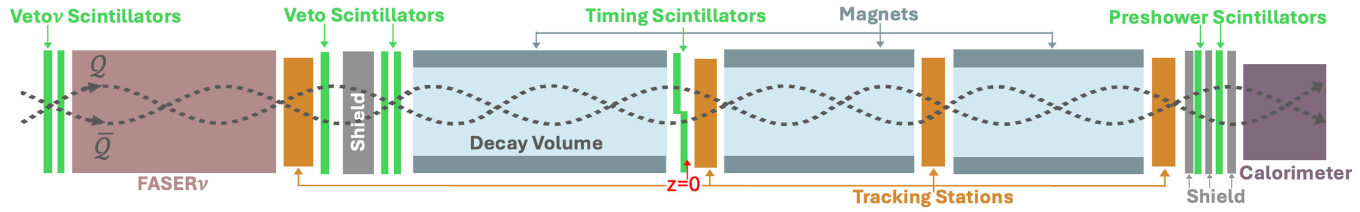

FIG. 1. Schematic of a quirk pair propagating through various components of the FASER detector. The detector geometry is not shown to scale.

permanent dipole magnets, a sampling electromagnetic calorimeter with lead absorbers, and four stations of a tracking spectrometer, each consisting of three 240×240 mm^2 double-sided layers of silicon strip detectors which reconstruct track segments per station. The tracker readout window is limited to 75 ns. In addition, eight scintillators, placed across 5.78 m in the z direction and each with a transverse size of at least 300×300 mm^2 and thickness of 10 or 20 mm, provide information about both the charge and arrival time of charged particles, crucial in the detection of slow quirk pairs with charge deposition twice that of background muons. The scintillators and calorimeters are read out with photomultiplier tubes, and the corresponding waveforms are recorded for each event in a time window of at least 560 ns, from which the arrival time and integrated charge can be reconstructed. The z coordinate is defined along the ATLAS beam collision axis, with $z = 0$ at the timing scintillator station. Transverse x and y coordinates are zero at the center of the detector.

This analysis is based on a data sample of proton-proton collisions at $\sqrt{s} = 13.6$ TeV collected by FASER from 2022 to 2024, corresponding to an integrated luminosity of 186 fb^{-1}. Events are triggered by measured charge above the noise level in any of the scintillators or the calorimeter. For each run, the scintillator time is calibrated to the average arrival time of muons, propagating at velocity c. The calibrated time of each scintillator is zero for a particle produced in a colliding bunch in ATLAS and traveling at speed c. A position-dependent correction, based on the tracker-inferred hit position, accounts for photon propagation within the scintillator. As detailed in Appendix B, this process leads to position- and charge-dependent timing resolutions per event and scintillator, with averages in the range of 0.25–0.6 ns.

We simulate the Drell-Yan production of quirk pairs at ATLAS at leading order using MadGraph5_aMC@NLO [20] interfaced with PYTHIA 8 [21] for parton showering. Events within the fiducial phase space of $\theta < 0.5$ mrad, where θ is the angle relative to the z axis, are retained for quirk propagation. We use the LHCb tune [22] of PYTHIA 8 and apply a normalization correction derived from LHCb measurements of forward Z boson production [23] to improve the modeling of initial-state radiation, to which the FASER signal acceptance is particularly sensitive.

Quirk propagation through material and LHC magnetic fields is simulated [24] following Ref. [25]. For $\Lambda > 3$ keV, deflections from magnetic fields and ionization in material are negligible because the IC force dominates and rapid oscillations average out residual effects. The quirks are therefore propagated analytically without stepwise simulation, while ionization energy loss is still accounted for. Energy loss from IC glueball radiation and quirk-pair deflection outside the detector acceptance are included using survival probabilities parameterized by m_Q, Λ, and the quirk-pair energy. The IC glueball-emission probability per oscillation is set to $\epsilon = 0.1$, following Refs. [12,17].

The interaction of quirks with the FASER detector is simulated in Geant4 [26] using an adaptation of the quirks physics extension within the Athena software framework [27], which accounts for the IC force in the propagation of individual quirks. The digitized detector response is reconstructed using the same pipeline as the collision data. Scintillator times are smeared according to the data-driven resolutions obtained after the timing calibration. Details of the quirk simulation in Geant4 and the timing calibration are discussed in Appendixes A and B, respectively.

The background model includes two sources. The first process, referred to as "through-going muons," consists of muons coming from near the ATLAS interaction point and depositing charge in all FASER scintillators. The second, referred to as the "kinked" background, arises from muon-induced secondary showers in the detector material. Kinked events exhibit timing consistent with propagation at velocity c upstream of a fitted position $z = z_{\text{kink}}$, followed by delayed signals in downstream scintillators due to the additional path length of the shower particles. Contributions from other sources including cosmic muons and beam-induced backgrounds from the two LHC beams were found to be negligible. Simulated muon samples with kinematics predicted by FLUKA [28,29] are used for validations of the scintillator charge response, the modeling of timing, and the source of the kinked background.

Events are selected based on tracker activity, calorimeter energy, timing, and measured charge in the scintillators. Events with one track segment in each of the two upstream tracker stations and more than one segment in either of the two downstream stations are rejected to suppress the kinked background. The quirk oscillation pattern in the x-y plane

ranges from approximately linear to circular as the pair angular momentum increases. An ellipse is therefore fitted to the tracker-hit distribution. The reconstructed tracker-hit ellipse is required to have at least 95% of its area contained within a tracker fiducial region defined by $|x| < 115$ mm and $|y| < 115$ mm, inscribing a circle corresponding to polar angle $\theta \simeq 0.24$ mrad. Events for which the semiminor axis of this ellipse is greater than 3 cm are rejected to suppress showers.

To avoid calorimeter showers that could backsplash to the preshower scintillators, a loose upper cut of 15 (3) GeV on calorimeter activity for 2022–2023 (2024) events is applied, according to the noise level in this period with an integrated luminosity of 67(120) fb^{-1}.

The scintillator charge signal from a muon follows a Landau distribution convoluted with detector noise. A quirk pair typically produces approximately twice the measured charge of a single muon, allowing a minimum per-scintillator charge requirement to reject most through-going muons. Secondary showers, however, generally produce larger signals. A quirk-like charge (QLC) window is therefore optimized for each scintillator to maximize the expected signal significance. A preselection of data includes all signal selections mentioned so far, as well as requiring QLC in at least one Vetoν and one preshower scintillator. The single-muon QLC efficiency varies by scintillator and is weakly correlated across planes: 22.24% of muons have QLC in the first Vetoν layer, 15.70% in both Vetoν layers, and 0.01% in all eight scintillators. Signal events must have QLC in all eight scintillators. This requirement is typically 99% efficient for simulated quirks entering the tracker fiducial region.

Figure 2 shows the calibrated time as a function of scintillator position for a selected event in data compared to a simulated slow signal event. A linear regression extracts t_0, the arrival time at $z = 0$ modulo the 25 ns LHC bunch

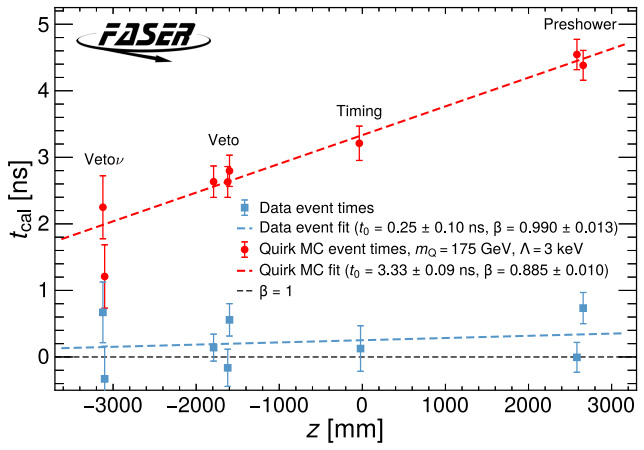

FIG. 2. Scintillator time vs position for a selected event in data (blue) compared to a selected slow simulated quirk event (red). Scintillator times are calibrated to the arrival time of particles traveling at the speed of light (black).

spacing, as well as $\beta = v/c$, with v being the extracted speed from the slope of the fitted line. Events with $\chi^2/N_{\text{d.o.f.}} > 2$ are rejected to ensure a precise extraction of β and t_0.

Figure 3(a) shows the distribution of arrival times t_0 and velocities β for preselected events in data, compared to a few representative quirk models, each with 300 simulated events. Because the per-run timing calibration is referenced to the mean arrival time, the muon-dominated data cluster near $t_0 = 0$ and $\beta = 1$. Delayed radio-frequency (rf) satellite-bucket events [30] arise when the 400 MHz rf cavities push a small fraction of protons from the main colliding bucket into the neighboring rf buckets, creating satellite bunches whose collisions occur at times offset by integer multiples of the 2.5 ns rf period relative to the nominal collisions. Massive quirks, however, are more likely to be slow and arrive out of time, possibly later than the 25 ns LHC bunch spacing.

Through-going muons are rejected using a combined arrival-time and speed selection, with ellipses in the t_0-β plane defining background regions excluded from the final selection. Using speed alone removes substantial signal near the nominal bunch, while arrival time alone is susceptible to satellite-bunch backgrounds. The preselected data distribution is fitted with a two-dimensional Gaussian, whose 1σ ellipse is scaled by a factor n chosen to maximize the signal significance against the background model. This gives $n = 10.0$ for the main timing ellipse shown in Fig. 3(a). Satellite events are rejected using 2.5 ns-spaced copies of the main ellipse, with their scaling factor optimized similarly, giving $n = 4.1$ for 2022–2023 and $n = 4.4$ for 2024. The optimization is performed separately for the two periods because differences in the LHC optics and collision crossing angle lead to differences in background-muon kinematics. Together, the ellipses define the excluded in-time region, with their lower halves forming its lower boundary.

Events are additionally required to have $0.6 < \beta < 1$. The lower bound allows for sufficient propagation time through all tracking stations within their readout time window, while the upper bound removes a region predominantly populated by background events.

The final allowed, out-of-time region is shown unhatched in Fig. 3(a). Depending on m_Q, 0.03%–0.04% of signal events produced at the ATLAS interaction point reach FASER within the tracker fiducial region. The efficiency of reconstructing and selecting the remaining events is approximately 1% (10%) for $m_Q = 50(175)$ GeV.

Statistical fluctuations in timing information and scintillator charge can lead to inefficiencies in rejection of through-going muons. This background is estimated with a data-driven ABCD method. Regions A–D span $0.6 < \beta < 1.0$: A and C lie outside the timing ellipses, while B and D lie inside; A and B require QLC in all

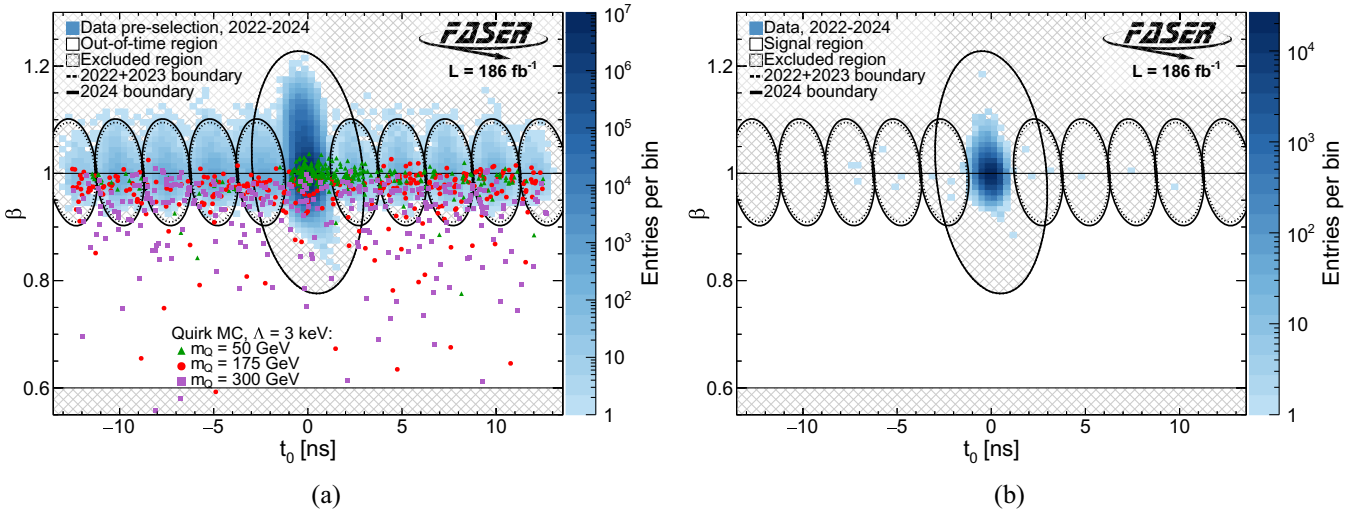

FIG. 3. (a) Preselected data (blue) and three representative signal models, each containing 300 events (green, red, purple) in the t_0-β plane, where t_0 is the arrival time modulo 25 ns and $\beta = v/c$ is the reconstructed speed in units of c. (b) The same distribution for data after requiring QLC in all eight scintillators. The ellipses define the in-time regions for the two data-taking periods, with events above their lower halves classified as in-time and those below as out-of-time. The unhatched region in (b) is the signal region of this analysis.

scintillators, whereas C and D require QLC in 2–5 scintillators. The estimated background rate in the signal region A is fixed to $N_A = N_B N_C / N_D$, giving 0.046 ± 0.010 in 2022–2023 and 0.082 ± 0.014 in 2024. The uncertainties are determined by the Poisson uncertainties on the number of events in each control region. This method relies on negligible correlation between charge and timing variables, which was confirmed by subdividing various control regions and repeating a similar method. Moreover, signal simulation studies confirm that there is negligible contamination of the signal in regions B–D.

The kinked-background rate is estimated using preselected data events with $n > 5.5$ that populate the non-Gaussian tail within the main ellipse of Fig. 3. Their rate density falls approximately as a power law in n and is extrapolated beyond the ellipse into the signal region. Kinked events have showers produced in both upstream and downstream detector material, but the downstream contribution is found to be negligible. The extrapolated rate is scaled by the probability that the upstream muon deposits QLC in the Vetoν scintillators, as described in Appendix C. The kinked background rate in the signal region is found to be 0.033 (0.051) events, with a 68% confidence level (CL) interval of [0.0015, 0.080] ([0.0036, 0.113]) in the 2022–2023 (2024) data-taking period.

The systematic uncertainties on the signal normalization depend on mass and confinement scale. A 2%–4% uncertainty is assigned to the inclusive production cross section, from QCD scales and parton distribution function uncertainties. Initial- and final-state radiation during the production affect the p_T of the quirk pair and consequently the fiducial acceptance. Variations of shower scales in PYTHIA 8 by factors of 0.5 and 2 result in approximately 25% uncertainty on the signal normalization. The uncertainty from IC glueball radiation is 1%–20%, relevant only for $\Lambda \gtrsim 30$ keV and derived by modifying the radiation probability ϵ by $\pm 30\%$. The uncertainty on the efficiency of requiring a good linear regression in the extraction of t_0 and β is 1%–5%. The finite-simulation-sample systematic uncertainty is approximately 6% for most relevant signal models but up to 28% for small confinement scales where the large oscillation amplitude translates to low signal acceptance. Finally, the uncertainty in the luminosity measurement is 2.2% [31,32].

Figure 3(b) shows the remaining events in data after applying all analysis selections except timing requirements. The unhatched region represents the final signal region of this analysis containing zero observed events. We proceed to set limits on the effective fiducial cross section, $\sigma_{\text{fid}}^{\text{eff}}$, and signal strength, $\mu(\Lambda, m_Q) = \sigma_{\text{fid}}^{\text{eff}} / \sigma_{\text{fid}}^{\text{pred}}$, using a fully frequentist (CL_s) method [33] within the CMS Combine framework [34]. The effective fiducial cross section, $\sigma_{\text{fid}}^{\text{eff}}$, is defined as the cross section within the fiducial phase space of $\theta < 0.5$ mrad weighted by the theoretically calculated probability of survival from IC glueball radiation in a 480-meter propagation of the quirk pair toward FASER. It is theoretically predicted to be $\sigma_{\text{fid}}^{\text{pred}}$.

Figure 4 shows the observed and expected 95% CL upper limits on $\sigma_{\text{fid}}^{\text{eff}}$ as a function of the quirk mass for a confinement scale of 10 keV, compared to the theoretical prediction. It also shows the 95% CL upper limit on $\mu(\Lambda, m_Q)$, with the expected uncertainty band including all experimental and theoretical uncertainties. Masses up to 179 GeV are excluded for $\Lambda = 10$ keV.

Figure 5 shows the observed and expected excluded regions in the Λ, m_Q signal parameter space compared to previously excluded regions [8,10,35]. For the first time,

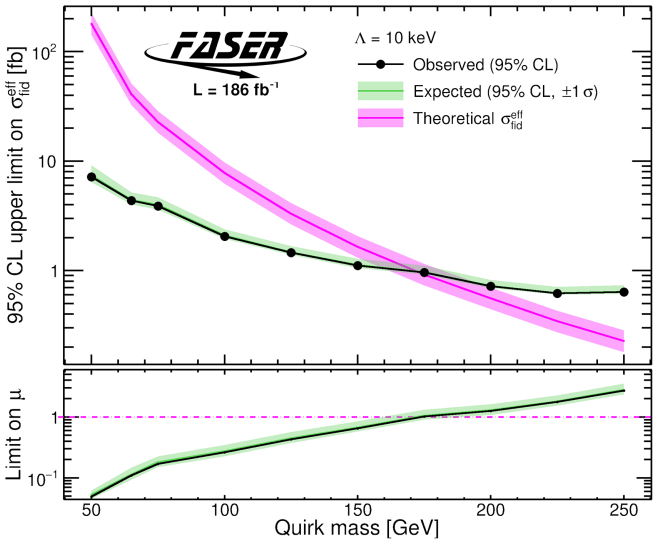

FIG. 4. Top: observed (black) and expected (green) 95% CL upper limit on the effective fiducial cross section, $\sigma_{\mathrm{fid}}^{\mathrm{eff}}$, compared to the theoretical prediction $\sigma_{\mathrm{fid}}^{\mathrm{pred}}$ (magenta). Bottom: 95% CL upper limit on $\mu(\Lambda, m_Q) = \sigma_{\mathrm{fid}}^{\mathrm{eff}}/\sigma_{\mathrm{fid}}^{\mathrm{pred}}$, accounting for all theoretical and experimental uncertainties.

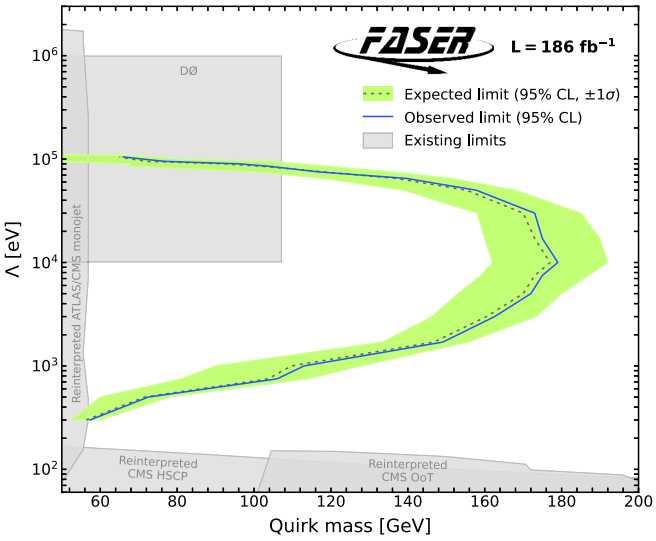

FIG. 5. Observed (solid blue) and expected (dashed blue) 95% CL exclusion contours in the m_Q-Λ plane, along with previously excluded regions (gray) [8,10,35]. The green band reflects the $\pm 1\sigma$ uncertainty on the expected limit.

quirk models with masses beyond the electroweak scale are excluded for a wide range of confinement scales. Tabulated results, including limits on $\sigma_{\mathrm{fid}}^{\mathrm{eff}}$, are provided in the HEPData record for this analysis [36].

In summary, a search for quirk particles at the LHC was performed using the data recorded by FASER in 2022–2024, approximately half of the recorded data in Run 3 of the LHC. As no quirk candidates were observed, lower limits were set on the quirk mass for a wide range of confinement scales covering three orders of magnitude. As the first direct search for quirks at the LHC, this analysis sets a precedent and offers a baseline methodology for similar searches in other collider experiments.

We acknowledge the ATLAS Collaboration software, Athena, on which FASER's offline software system is based. We also thank James Black and Timothy Nelson for their work in developing the comprehensive Geant4 quirks extension within the Athena software framework.

Acknowledgments—We thank CERN for the excellent performance of the LHC and the technical and administrative staff members at all FASER institutions for their contributions to the success of the FASER experiment. We thank the ATLAS Collaboration for providing us with accurate luminosity estimates for the Run 3 LHC proton-proton (pp) collision data. We thank the CERN STI group for providing detailed FLUKA simulations of the muon fluence at FASER. This work was supported in part by Heising-Simons Foundation Grants No. 2019-1179 and No. 2020-1840; Simons Foundation Grant No. 623683; JSPS KAKENHI Grants No. 19H01909, No. 22H01233, No. 20K23373, No. 23H00103, No. 20H01919, and No. 21H00082; the joint research program of the Institute of Materials and Systems for Sustainability, ERC Consolidator Grant No. 101002690; DFG Grant No. SCHO 1527/13-1; Royal Society Grant No. URF\R1 \201519; UK Science and Technology Funding Councils Grant No. ST/T505870/1; the National Natural Science Foundation of China; the Tsinghua University Initiative Scientific Research Program; and the Swiss National Science Foundation.

Data availability—The exclusion limits and effective fiducial cross sections are tabulated in HEPData [36]. The dataset used to obtain these results and the analysis framework are not publicly available. The PreFASERQuirkSim software used for quirk propagation is publicly available [24]. The authors may be contacted for further information on data or software availability.

[1] Z. Chacko, H.-S. Goh, and R. Harnik, The twin Higgs: Natural electroweak breaking from mirror symmetry, Phys. Rev. Lett. **96**, 231802 (2006).
[2] G. Burdman, Z. Chacko, H.-S. Goh, and R. Harnik, Folded supersymmetry and the LEP paradox, J. High Energy Phys. 02 (2007) 009.
[3] B. Batell, M. Low, E. T. Neil, and C. B. Verhaaren, Review of neutral naturalness, Phys. Rep. **1165**, 1 (2026).
[4] J. Kang and M. A. Luty, Macroscopic strings and "quirks" at colliders, J. High Energy Phys. 11 (2009) 065.
[5] J. E. Juknevich, D. Melnikov, and M. J. Strassler, A pure-glue hidden valley I. States and decays, J. High Energy Phys. 07 (2009) 055.
[6] P. Asadi, G. D. Kribs, and M. A. Luty, Quirks live in cool universes, J. High Energy Phys. 05 (2025) 092.

[7] H. Cai, H.-C. Cheng, and J. Terning, A quirky little Higgs model, J. High Energy Phys. 05 (2009) 045.
[8] D0 Collaboration, Search for new fermions ("quirks") at the Fermilab Tevatron collider, Phys. Rev. Lett. **105**, 211803 (2010).
[9] G. Burdman, Z. Chacko, H.-S. Goh, R. Harnik, and C. A. Krenke, The quirky collider signals of folded supersymmetry, Phys. Rev. D **78**, 075028 (2008).
[10] M. Farina and M. Low, Constraining quirky tracks with conventional searches, Phys. Rev. Lett. **119**, 111801 (2017).
[11] S. Knapen, H. K. Lou, M. Papucci, and J. Setford, Tracking down quirks at the Large Hadron Collider, Phys. Rev. D **96**, 115015 (2017).
[12] J. A. Evans and M. A. Luty, Stopping quirks at the LHC, J. High Energy Phys. 06 (2019) 090.
[13] Q. Sha, D. Murnane, M. Fieg, S. Tong, M. Zakharyan, Y. Fang, and D. Whiteson, Learning to reconstruct quirky tracks, Phys. Rev. D **114**, 015005 (2026).
[14] J. Forsyth, M. Low, C. Tenney, and C. B. Verhaaren, Visible collider signals of natural quirks, Phys. Rev. D **111**, 115021 (2025); **113**, 079901(E) (2026).
[15] D. Curtin, S. Dreyer, M. F. Costa, S. Heim, G. Kasieczka, L. Moureaux, D. Rousso, D. Shih, and M. Sommerhalder, Soft-unclustered-energy patterns from quirks, Phys. Rev. D **113**, 015010 (2026).
[16] X. Cid Vidal *et al.*, Searching for quirks at LHCb, J. High Energy Phys. 26 (2026) 037.
[17] J. L. Feng, J. Li, X. Liao, J. Ni, and J. Pei, Discovering quirks through timing at FASER and future forward experiments at the LHC, J. High Energy Phys. 06 (2024) 197.
[18] FASER Collaboration, The FASER detector, J. Instrum. **19**, P05066 (2024).
[19] J. Li, J. Pei, L. Ran, and W. Zhang, The quirk signal at FASER and FASER 2, J. High Energy Phys. 12 (2021) 109.
[20] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, J. High Energy Phys. 07 (2014) 079.
[21] T. Sjostrand, S. Mrenna, and P. Z. Skands, A brief introduction to PYTHIA 8.1, Comput. Phys. Commun. **178**, 852 (2008).
[22] LHCb Collaboration, Supplemental Material for "Measurement of charged hadron production in Z-tagged jets in proton-proton collisions at $\sqrt{s} = 8$ TeV", Phys. Rev. Lett. **123**, 232001 (2019).
[23] LHCb Collaboration, Precision measurement of forward Z boson production in proton-proton collisions at $\sqrt{s} = 13$ TeV, J. High Energy Phys. 07 (2022) 026.
[24] E. Welch, PreFASERQuirkSim, (2026), software.
[25] J. Li, T. Li, J. Pei, and W. Zhang, Uncovering quirk signal via energy loss inside tracker, Phys. Rev. D **102**, 056006 (2020).
[26] S. Agostinelli *et al.* (GEANT4 Collaboration), Geant4—A simulation toolkit, Nucl. Instrum. Methods Phys. Res., Sect. A **506**, 250 (2003).
[27] ATLAS Collaboration, Athena (version 21.0.127) [computer software], Zenodo, (2021).
[28] A. Ferrari, P. R. Sala, A. Fassò, and J. Ranft, FLUKA: A multi-particle transport code (program version 2005), Technical Report No. CERN-2005-010, CERN, 2005.
[29] G. Battistoni *et al.*, Overview of the FLUKA code, Ann. Nucl. Energy **82**, 10 (2015).
[30] A. Jeff *et al.*, Measurement of satellite bunches at the LHC, in *Proceedings of the 3rd International Particle Accelerator Conference (IPAC'12)*, New Orleans, LA, paper MOEPPB010 (2012), pp. 97–99.
[31] ATLAS Collaboration, Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC, Eur. Phys. J. C **83**, 982 (2023).
[32] ATLAS Collaboration, Preliminary luminosity calibration of the ATLAS 13.6 TeV data recorded in 2024 and combination with the 2022 and 2023 measurements, Technical Report No. ATL-DAPR-PUB-2025-001, CERN, Geneva, 2025.
[33] A. L. Read, Modified frequentist analysis of search results (the CL_s method), CERN Report No. CERN-OPEN-2000-205, 2000, 10.5170/CERN-2000-005.81.
[34] CMS Collaboration, The CMS statistical analysis and combination tool: Combine, Comput. Software Big Sci. **8**, 19 (2024).
[35] CMS Collaboration, Search for decays of stopped exotic long-lived particles produced in proton-proton collisions at $\sqrt{s} = 13$ TeV, J. High Energy Phys. 05 (2018) 127.
[36] FASER Collaboration, HEPData for first search for quirks at the LHC with FASER, https://www.hepdata.net/record/181308 (2026).

End Matter

Appendix A: Simulation of quirks in Geant4—Quirk propagation through the detector is simulated in Geant4 using an adaptation of the `Simulation/G4Extensions/Quirks` physics extension from the Athena software framework [27]. The quirk pair is treated as a coupled two-body system connected by an IC string. A custom `G4UserStackingAction` class, which controls the order in which tracks are processed during the Geant4 event, assigns the quirk and antiquirk to separate waiting categories, enforcing alternating propagation. After each step, the string configuration is updated before the partner track is advanced, so the pair is transported as a coupled system. The string force is applied by a dedicated transportation process, while standard Geant4 electromagnetic processes describe material interactions.

The implementation is validated with three checks. First, track continuity is verified by comparing trajectories after upstream propagation with those at the detector simulation interface. Second, the scintillator charge response is compared with a straight-line reference using the same kinematics and relevant electromagnetic processes, confirming

the expected response: similar to that of two minimum ionizing particles (MIPs). Third, the scintillator timing response is compared with the same reference. The charge and timing distributions are consistent within differences expected from the oscillatory quirk motion, which changes scintillator intersections and path lengths. These checks validate the detector-response quantities used in the analysis.

Appendix B: Timing calibration—Reconstructed scintillator times are measured relative to the LHC clock and are calibrated using high-energy, single-track muon events. Run-dependent reference offsets are first subtracted to remove relative timing shifts between channels and data-taking periods. The resulting times are mapped to a common 25 ns bunch period and the remaining bunch-index ambiguity is resolved using the absolute digitizer times. Position-dependent photon propagation times are measured in 2.5×2.5 cm^2 bins across each scintillator, and the calibrated hit time is defined as $t_{\mathrm{cal}} = t - \bar{t}(x, y)$.

For each scintillator plane, a *leave-one-out* (LOO) estimate of the event arrival time at $z = 0$, t_0^{LOO} is reconstructed from the measured times in all other valid planes, excluding the plane under study. The LOO procedure is iterated during the calibration, reducing the width of the residual distribution until it converges to the approximate intrinsic timing resolution. The timing uncertainty is then parametrized as a function of hit position and measured charge by fitting its charge dependence in each spatial bin to a power law. The data-derived timing resolutions are used to smear the simulated hit times before applying the timing calibration. The tracker hits are used to define a containment ellipse for the expected hit region in each scintillator plane, over which the calibration values are averaged, with their spatial variation included as an additional uncertainty. This region is small and approximately circular for most muons and high-Λ quirk events, while for low-Λ quirks it ranges from a narrow, elongated ellipse to a large circle, depending on the angular momentum of the quirk system.

The resolutions of the calibrated time, t_{cal}, and the uncalibrated raw time, t_{raw}, are reported using muon data. For each scintillator plane, an LOO estimate of the arrival time at $z = 0$, $t_{0,\mathrm{cal}}^{\mathrm{LOO}}$ is reconstructed from the calibrated times and uncertainties of all other valid planes, excluding the plane under study, assuming that $v = c$. $t_{0,\mathrm{cal}}^{\mathrm{LOO}}$ is subtracted event by event from both t_{cal} and t_{raw} in the excluded plane to correct for the arrival-time due to the intrinsic collision-time spread of approximately 180 ps arising from the finite bunch length. Gaussian fits are performed to the cores of the resulting time distributions, and the mean variance of the corresponding LOO arrival-time estimate, $\langle \sigma^2_{t_0,\mathrm{LOO}} \rangle$, is subtracted in quadrature from the fitted widths. Using the calibrated event estimate for both measurements places them on the same reference, although it slightly improves the apparent raw time resolution relative to a fully independent raw time estimate.

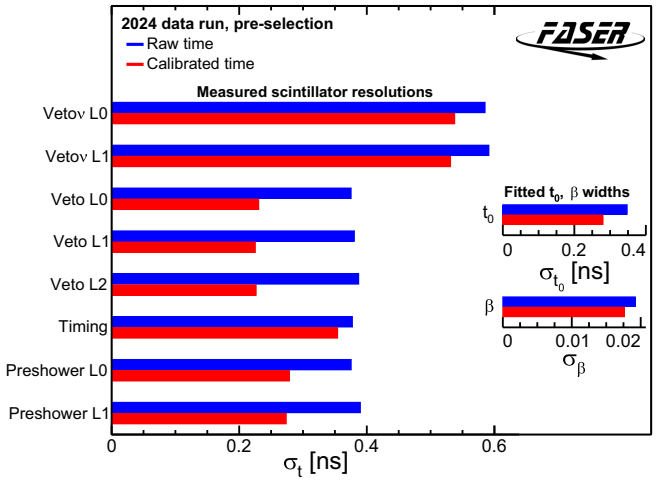

FIG. 6. Measured timing resolutions for a 2024 data run. The right panels show the widths of the fitted t_0 and β distributions for the raw and calibrated timing inputs. The t_0 distributions include collision-time spread.

The resulting typical scintillator timing resolutions are shown in Fig. 6. The timing resolution is slightly worse for Vetoν scintillators as the light is extracted using a wavelength-shifting rod which widens the timing spread. The widths of the reconstructed t_0 and β distributions are shown on the right side of Fig. 6, with the t_0 distribution including the intrinsic collision-time spread. For quirk signals, the scintillator charge is typically about two MIPs, corresponding to smaller charge-parameterized timing uncertainties than for the muon sample.

Appendix C: Kinked background—The kinked background arises when a muon traversing FASER at $v = c$ interacts in the detector and produces secondary shower particles. The additional path length accumulated by the shower particles delays the downstream scintillator signals, changing the timing profile from the $v = c$ behavior of the incoming muon to an apparently slower downstream propagation. Events with this characteristic kink can mimic the timing signature of a slow particle and appear at large n relative to the nominal through-going-muon distribution in the t_0-β plane.

Muon simulation identifies two main sources of such events: showers initiated in the FASERν tungsten and hadronic backsplash from the calorimeter lead. To distinguish these sources, the scintillator times are fit with a piecewise timing model with three free parameters: the arrival time, the kink z position z_{kink}, and the apparent downstream speed, with $v = c$ being the fixed upstream of z_{kink}. An example fit is shown in Fig. 7.

The fitted kink z position z_{kink} is used to classify the likely shower origin, with $z_{\mathrm{kink}} < 0$ assigned to the

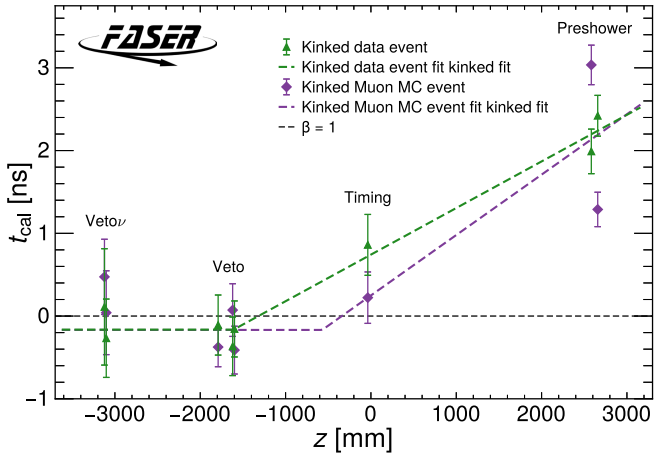

FIG. 7. Calibrated scintillator times for a candidate kinked-background event in data, compared to a simulated event in the same category. The kinked timing fit is used to extract the effective shower z position z_{kink}.

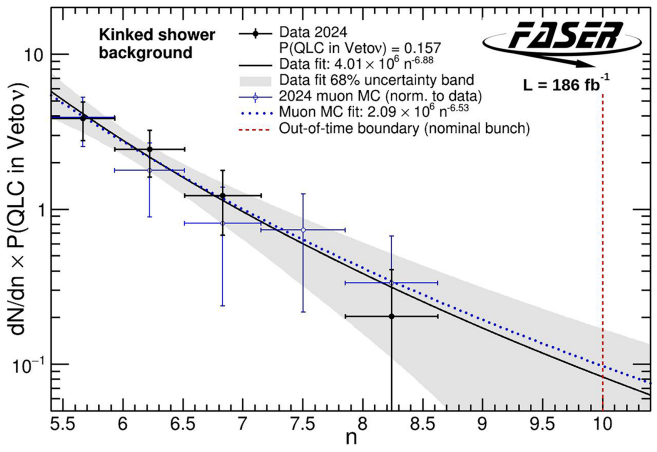

FIG. 8. The 2024 upstream kinked-background estimate. The black points show the data event density multiplied by the Vetoν QLC probability, $P(\mathrm{QLC\ in\ Veto}\nu) = 0.157$, as a function of the ellipse-scaling variable n. The solid black curve shows the power-law fit used to extrapolate the kinked background into the signal region, and the gray band shows the 68% uncertainty from the fit. The blue points show the 2024 muon simulation normalized to the data for a shape comparison, with the dotted blue curve showing the corresponding power-law fit. The dashed red line indicates the nominal out-of-time signal-region boundary.

upstream FASERν category and $z_{\mathrm{kink}} > 0$ assigned to the downstream calorimeter category.

To transfer the kinked background rate to the signal region, scintillators downstream of z_{kink} are conservatively assumed to satisfy the QLC requirement, while scintillators upstream of z_{kink} are assigned a probability of producing QLC signatures measured from data, $P(\mathrm{QLC})$. This makes the upstream category the more pessimistic component estimate since here $P(\mathrm{QLC}) = P(\mathrm{QLC\ in\ Veto}\nu) = 0.157$ and all downstream scintillators are assumed to satisfy the QLC requirement.

The kinked-background rate in the out-of-time region is estimated from data before applying the final timing and charge requirements. In both data and muon simulation, the kinked-event density falls approximately as a power law in n, motivating a power-law model for the extrapolation. The density in the preselected sample is fit for $n > 5.5$, where contamination from through-going muons is suppressed; including lower-n events would bias the fitted slope and underestimate the background at large n. No events in this fit region are classified as downstream, largely due to the calorimeter energy requirement, and the corresponding transfer factor would further suppress this contribution; only the upstream category is therefore considered in the background estimate.

The 2024 upstream kinked-background distribution and corresponding fit are shown in Fig. 8. The fitted density is integrated over the nominal out-of-time signal-region range in n. The same fitted density is integrated over the corresponding satellite out-of-time regions, with the satellite integral scaled by the measured ratio of events in the nominal bunch period to those in the satellite bunch periods. The sum of the nominal and scaled satellite integrals gives the upstream kinked yield before the final QLC requirement. This yield is then multiplied by the probability for QLC in the Vetoν scintillators.