Editors' Suggestion | Featured in Physics

Validation of the DESI DR2 Lyα BAO analysis using synthetic datasets

L. Casas,[1] H. K. Herrera-Alcantar,[2,3] J. Chaves-Montero,[1] A. Cuceu,[4,*] A. Font-Ribera,[1] M. Lokken,[1] M. Abdul-Karim,[3] C. Ramírez-Pérez,[1] D. Alonso,[5] J. Aguilar,[4] S. Ahlen,[6] U. Andrade,[7,8] E. Armengaud,[3] A. Aviles,[9,10,11] S. Bailey,[4] S. BenZvi,[12] D. Bianchi,[13,14] A. Brodzeller,[4] D. Brooks,[15] R. Canning,[16] A. Carnero Rosell,[17,18] M. Charles,[19] E. Chaussidon,[4] T. Claybaugh,[4] K. S. Dawson,[20] A. de la Macorra,[11] A. de Mattia,[3] Arjun Dey,[21] Biprateep Dey,[22,23] Z. Ding,[24] P. Doel,[15] D. J. Eisenstein,[25] W. Elbers,[26] S. Ferraro,[4,27] J. E. Forero-Romero,[28,29] C. Garcia-Quintero,[25,*] Lehman H. Garrison,[30,31] E. Gaztañaga,[32,16,33] H. Gil-Marín,[34,32,35] S. Gontcho A. Gontcho,[4] A. X. Gonzalez-Morales,[36] C. Gordon,[1] G. Gutierrez,[37] J. Guy,[4] M. Herbold,[19] K. Honscheid,[38,39,19] C. Howlett,[40] D. Huterer,[41,8] M. Ishak,[42] S. Juneau,[21] R. Kehoe,[43] D. Kirkby,[44] T. Kisner,[4] A. Kremin,[4] O. Lahav,[15] M. Landriau,[4] J. M. Le Goff,[3] L. Le Guillou,[45] A. Leauthaud,[46,47] M. E. Levi,[4] Q. Li,[20] M. Manera,[48,1] P. Martini,[38,49,19] A. Meisner,[21] J. Mena-Fernández,[50] R. Miquel,[51,1] J. Moustakas,[52] D. Muñoz Santos,[53] A. D. Myers,[54] S. Nadathur,[16] L. Napolitano,[54] G. Niz,[36,9] H. E. Noriega,[10,11] E. Paillas,[55,56] N. Palanque-Delabrouille,[3,4] W. J. Percival,[57,58,59] Matthew M. Pieri,[53] C. Poppett,[4,60,27] F. Prada,[61] I. Pérez-Ràfols,[62] C. Ravoux,[63] G. Rossi,[64] E. Sanchez,[65] D. Schlegel,[4] M. Schubnell,[41,8] H. Seo,[66] F. Sinigaglia,[17,18] D. Sprayberry,[21] T. Tan,[3] G. Tarlé,[8] P. Taylor,[19] W. Turner,[38,49,19] M. Vargas-Magaña,[11] M. Walther,[67,68] B. A. Weaver,[21] M. Wolfson,[19] C. Yèche,[3] P. Zarrouk,[45] and R. Zhou[4]

(DESI Collaboration)

[1]*Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Edifici Cn, Campus UAB, 08193, Bellaterra (Barcelona), Spain*
[2]*Institut d'Astrophysique de Paris, 98 bis boulevard Arago, 75014 Paris, France*
[3]*IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France*
[4]*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 94720, USA*
[5]*Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom*
[6]*Physics Department, Boston University, 590 Commonwealth Avenue, Boston, Massachusetts 02215, USA*
[7]*Leinweber Center for Theoretical Physics, University of Michigan, 450 Church Street, Ann Arbor, Michigan 48109-1040, USA*
[8]*University of Michigan, 500 South State Street, Ann Arbor, Michigan 48109, USA*
[9]*Instituto Avanzado de Cosmología A. C., San Marcos 11-Atenas 202, Magdalena Contreras, Ciudad de México C. P. 10720, Mexico*
[10]*Instituto de Ciencias Físicas, Universidad Nacional Autónoma de México, Avenida Universidad s/n, Cuernavaca, Morelos, C. P. 62210, Mexico*
[11]*Instituto de Física, Universidad Nacional Autónoma de México, Circuito de la Investigación Científica, Ciudad Universitaria, Cd. de México C. P. 04510, Mexico*
[12]*Department of Physics and Astronomy, University of Rochester, 206 Bausch and Lomb Hall, P.O. Box 270171, Rochester, New York 14627-0171, USA*
[13]*Dipartimento di Fisica "Aldo Pontremoli," Università degli Studi di Milano, Via Celoria 16, I-20133 Milano, Italy*
[14]*INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy*
[15]*Department of Physics and Astronomy, University College London, Gower Street, London, WC1E 6BT, United Kingdom*
[16]*Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth, PO1 3FX, United Kingdom*
[17]*Departamento de Astrofísica, Universidad de La Laguna (ULL), E-38206, La Laguna, Tenerife, Spain*
[18]*Instituto de Astrofísica de Canarias, C/ Vía Láctea, s/n, E-38205 La Laguna, Tenerife, Spain*
[19]*The Ohio State University, Columbus, 43210 Ohio, USA*
[20]*Department of Physics and Astronomy, The University of Utah, 115 South 1400 East, Salt Lake City, Utah 84112, USA*
[21]*NSF NOIRLab, 950 North Cherry Avenue, Tucson, Arizona 85719, USA*
[22]*Department of Astronomy and Astrophysics, University of Toronto, Toronto, Ontario M5S 3H4, Canada*
[23]*Department of Physics and Astronomy and Pittsburgh Particle Physics, Astrophysics, and Cosmology Center (PITT PACC), University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, Pennsylvania 15260, USA*
[24]*University of Chinese Academy of Sciences, Nanjing 211135, People's Republic of China*

[25]*Center for Astrophysics | Harvard and Smithsonian, 60 Garden Street, Cambridge, Massachusetts 02138, USA*
[26]*Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, United Kingdom*
[27]*University of California, Berkeley, 110 Sproul Hall No. 5800 Berkeley, California 94720, USA*
[28]*Departamento de Física, Universidad de los Andes, Carrera 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia*
[29]*Observatorio Astronómico, Universidad de los Andes, Carrera 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia*
[30]*Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, New York 10010, USA*
[31]*Scientific Computing Core, Flatiron Institute, 162 Fifth Avenue, New York, New York 10010, USA*
[32]*Institut d'Estudis Espacials de Catalunya (IEEC), c/ Esteve Terradas 1, Edifici RDIT, Campus PMT-UPC, 08860 Castelldefels, Spain*
[33]*Institute of Space Sciences, ICE-CSIC, Campus UAB, Carrer de Can Magrans s/n, 08913 Bellaterra, Barcelona, Spain*
[34]*Departament de Física Quàntica i Astrofísica, Universitat de Barcelona, Martí i Franquès 1, E08028 Barcelona, Spain*
[35]*Institut de Ciències del Cosmos (ICCUB), Universitat de Barcelona (UB), c. Martí i Franquès, 1, 08028 Barcelona, Spain*
[36]*Departamento de Física, DCI-Campus León, Universidad de Guanajuato, Loma del Bosque 103, León, Guanajuato C. P. 37150, Mexico*
[37]*Fermi National Accelerator Laboratory, PO Box 500, Batavia, Illinois 60510, USA*
[38]*Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, Ohio 43210, USA*
[39]*Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, Ohio 43210, USA*
[40]*School of Mathematics and Physics, University of Queensland, Brisbane, Queensland 4072, Australia*
[41]*Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, Michigan 48109, USA*
[42]*Department of Physics, The University of Texas at Dallas, 800 West Campbell Road, Richardson, Texas 75080, USA*
[43]*Department of Physics, Southern Methodist University, 3215 Daniel Avenue, Dallas, Texas 75275, USA*
[44]*Department of Physics and Astronomy, University of California, Irvine, California 92697, USA*
[45]*Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), FR-75005 Paris, France*
[46]*Department of Astronomy and Astrophysics, UCO/Lick Observatory, University of California, 1156 High Street, Santa Cruz, California 95064, USA*
[47]*Department of Astronomy and Astrophysics, University of California, Santa Cruz, 1156 High Street, Santa Cruz, California 95065, USA*
[48]*Departament de Física, Serra Húnter, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain*
[49]*Department of Astronomy, The Ohio State University, 4055 McPherson Laboratory, 140 West 18th Avenue, Columbus, Ohio 43210, USA*
[50]*Laboratoire de Physique Subatomique et de Cosmologie, 53 Avenue des Martyrs, 38000 Grenoble, France*
[51]*Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain*
[52]*Department of Physics and Astronomy, Siena College, 515 Loudon Road, Loudonville, New York 12211, USA*
[53]*Aix Marseille Université, CNRS, CNES, LAM, Marseille, France*
[54]*Department of Physics and Astronomy, University of Wyoming, 1000 East University, Department 3905, Laramie, Wyoming 82071, USA*
[55]*Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, Arizona 85721, USA*
[56]*Instituto de Estudios Astrofísicos, Facultad de Ingeniería y Ciencias, Universidad Diego Portales, Avenida Ejército Libertador 441, Santiago, Chile*
[57]*Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L 3G1, Canada*
[58]*Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, Ontario N2L 2Y5, Canada*

[59]Waterloo Centre for Astrophysics, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L 3G1, Canada
[60]Space Sciences Laboratory, University of California, Berkeley, 7 Gauss Way, Berkeley, California 94720, USA
[61]Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain
[62]Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 10, 08930 Barcelona, Spain
[63]Université Clermont-Auvergne, CNRS, LPCA, 63000 Clermont-Ferrand, France
[64]Department of Physics and Astronomy, Sejong University, 209 Neungdong-ro, Gwangjin-gu, Seoul 05006, Republic of Korea
[65]CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain
[66]Department of Physics and Astronomy, Ohio University, 139 University Terrace, Athens, Ohio 45701, USA
[67]Excellence Cluster ORIGINS, Boltzmannstrasse 2, D-85748 Garching, Germany
[68]University Observatory, Faculty of Physics, Ludwig-Maximilians-Universität, Scheinerstrasse 1, 81677 München, Germany

(Received 2 April 2025; accepted 10 December 2025; published 16 January 2026)

The second data release (DR2) of the Dark Energy Spectroscopic Instrument (DESI), containing data from the first three years of observations, doubles the number of Lyman-α (Lyα) forest spectra in DR1 and it provides the largest dataset of its kind. To ensure a robust validation of the baryonic acoustic oscillation (BAO) analysis using Lyα forests, we have made significant updates compared to DR1 to both the mocks and the analysis framework used in the validation. In particular, we present `CoLoRe-QL`, a new set of Lyα mocks that use a quasilinear input power spectrum to incorporate the nonlinear broadening of the BAO peak. We have also increased the number of realizations used in the validation to 400, compared to the 150 realizations used in DR1. Finally, we present a detailed study of the impact of quasar redshift errors on the BAO measurement, and we compare different strategies to mask damped Lyman-α absorbers in our spectra. The BAO measurement from the Lyα dataset of DESI DR2 is presented in a companion publication.

DOI: 10.1103/fvgh-kswf

I. INTRODUCTION

One of the central questions in cosmology is the cause behind the accelerated expansion of the Universe. Among the most powerful observables for investigating this phenomenon are baryon acoustic oscillations (BAO), fluctuations in the matter density caused by acoustic density waves in the early Universe. First measured in the distribution of galaxies twenty years ago [1,2], BAO serves as a standard ruler for measuring cosmic expansion.

Measuring BAO at $z > 2$ using galaxies as tracers is challenging, as obtaining a sufficient number of redshifts for distant, faint galaxies is highly time consuming. However, the Lyman-α (Lyα) forest in the spectra of distant quasars (QSO) provides a powerful alternative. This observable consists of absorption lines in the spectra of high-redshift quasars caused by neutral hydrogen in the intergalactic medium. Consequently, the Lyα forest traces the distribution of matter in the Universe, enabling the measurement of BAO at higher redshifts than those accessible by galaxy surveys. The Baryon Oscillation Spectroscopic Survey (BOSS, [3]) was the first to carry out these measurements from the autocorrelation of the Lyα forest [4–6] and its cross-correlation with quasars [7].

The Dark Energy Spectroscopic Instrument (DESI) represents a major step forward in the precise measurement of BAO from the Lyα forest [8,9]. DESI is a robotic, fiber-fed, highly multiplexed spectroscopic surveyor that operates on the Mayall 4-m telescope at Kitt Peak National Observatory [10] that can obtain simultaneous spectra of almost 5000 objects [11–14] using complex planning [15] and reduction pipelines [16]. The DESI instrument and spectroscopic pipelines were tested extensively during a period of survey validation before the start of the main survey [17,18]. The first year of DESI observations (DR1, [19]) yielded the most precise measurement of BAO up to that time from the Lyα forest using over 420 000 Lyα forest spectra and 700 000 quasars [20]. This dataset also led to precise BAO [21] and full-shape [22] measurements from the clustering of galaxies and quasars [23]. In combination with external probes, these measurements from DESI DR1 resulted in one of the tightest constraints on dynamical dark energy and on the sum of the neutrino masses [24,25].

The first three years of DESI observations (DR2) include over 820 000 Lyα forest spectra and 1.2×10^6 quasars at $z > 1.77$, almost doubling the dataset in DR1. The goal of

*NASA Einstein Fellow.

this work is to use synthetic realizations of DESI data (mock catalogs) to validate the Lyα BAO analysis of the second data release of DESI [26]. Mock datasets provide a controlled environment where all possible sources of systematics are known, since these mock datasets provide a controlled environment to study systematic effects. Therefore they are essential for testing the analysis pipeline and assessing the impact of potential systematic errors on cosmological constraints.

Compared to the validation of the DR1 measurement presented in [27], we have introduced several improvements to both the mocks themselves and to their analysis. These include refining the accuracy of quasar small-scale clustering, incorporating nonlinear broadening of the BAO peak, improving the treatment of redshift errors and damped Lyman-α absorbers (DLAs), and increasing the number of mocks from 150 to 400. Since the DESI DR2 Lyα BAO analysis is limited to scales larger than $30h^{-1}$ Mpc, expected to be fairly well described by linear theory [28], we use relatively simple log-normal mocks and ignore the impact of nonlinearities beyond the broadening of the BAO peak.

Our results are part of a comprehensive set of DESI DR2 BAO measurements from the clustering of galaxies, quasars, and the Lyα forest. The companion paper [29] presents the BAO measurements from galaxies and quasars at $z < 2$ and the cosmological interpretation of all BAO measurements.

The outline of the paper is as follows. Section II provides an overview of the mock datasets used in DR1 and DR2, focusing on the updates made for the validation of DR2. Section III describes the analysis methodology applied to the mocks, outlining the steps involved in extracting and interpreting the BAO signal. The main results are presented in Sec. IV, and in Sec. V we discuss the small bias present in BAO measurements when we introduce redshift errors before continuum fitting. Finally, Sec. VI summarizes the findings and presents the conclusions of this work.

II. MOCKS

This section outlines the process used to create the mocks for validating the Lyα BAO analysis in DR2, that can be divided into two stages. In the first one, explained in Sec. II A, we simulate the distribution of quasars and extract the Lyα transmitted flux fraction along the line of sight to each quasar, which we usually refer to as transmission skewers. The second stage transforms these transmission skewers into realistic DESI spectra, as detailed in Sec. II B.

A. Simulations of the Universe at $z > 2$

In DR1, two different types of mocks were used to simulate the distribution of quasars and the fluctuations in the Lyα forest: 100 realizations of `LyaCoLoRe` mocks [30] and 50 realizations of SACLAY mocks [31]. For the analysis validation of DR2 we have used 300 realizations of `CoLoRe-QL` mocks (improved version of `LyaCoLoRe` mocks used in DR1, described in Sec. II A 2) and 100 realizations of the SACLAY mocks. Both sets of mocks are based on local transformations of Gaussian random fields and do not include higher-order correlations arising from the nonlinear evolution of the density field. However, they have different approaches to populate the simulated boxes with quasars and to add redshift-space distortions (RSD) in the Lyα forest.

1. Previous mocks already used in DESI DR1

The first step to generate a `LyaCoLoRe` or `CoLoRe-QL` mock is to produce a three-dimensional 10 Gpc h^{-1} Gaussian random field using an input power spectrum corresponding to the linear power spectrum of density fluctuations at $z = 0$ and extrapolate it back in time to generate an all-sky light cone reaching $z = 3.8$. We then use a "biasing model" to translate this field into fluctuations in the density of quasars that we Poisson sample to obtain our catalog of quasar positions. This is done using the `CoLoRe` package,[1] described in [32], that is also used to compute the Gaussian "skewers" of both densities and line-of-sight velocities from the center of the box (the observer) to each of the quasar positions. We then use the package `LyaCoLoRe`,[2] described in [30], to translate these Gaussian skewers into the redshift-space Lyα optical depth, from which we compute the transmitted flux fraction $F = e^{-\tau}$. This process involves several steps. First, given that the cells of the initial field are $\simeq 2.4h^{-1}$ Mpc, in order to reproduce the one-dimensional fluctuations in the Lyα forest we need to add extra small-scale power to the Gaussian skewers. This is followed by a log-normal transformation in order to obtain positive densities. Next, we apply the fluctuating Gunn-Peterson approximation (FGPA, [33]) to compute the real-space optical depth. We then use the line-of-sight velocities to shift the absorption and obtain the redshift-space optical depth. Finally, we compute the transmitted flux fraction. We refer the reader to [30] for more details on these mocks.

SACLAY mocks, on the other hand, were created using the `SaclayMocks`[3] package. The methodology is similar to that of the `LyaCoLoRe` mocks, but it presents two important differences. Instead of using the same random field to simulate both the quasar positions and the Lyα skewers, here we use two random fields that have the same random seed, but different input power spectra. The input power spectrum for these mocks is chosen so that the resulting quasar power spectrum remains approximately linear, with a scale-independent bias and linear

[1] https://github.com/damonge/CoLoRe.
[2] https://github.com/igmhub/LyaCoLoRe.
[3] https://github.com/igmhub/SaclayMocks.

redshift-space distortions down to small scales. The second difference is the implementation of redshift-space distortions in the Lyα skewers. The velocity gradients along the lines of sight are also computed from the initial Gaussian density field. However, a modified FGPA transformation is then applied to the sum of the density and velocity gradient fields, generating directly the redshift-space optical depth. A multiplicative factor is applied to the velocity gradient to fit the measured amount of redshift-space distortions. This implementation allows for a predictive model of the correlations down to small scales [31].

2. New quasilinear mocks (CoLoRe-QL)

As shown in Fig. 16 of [27], the BAO uncertainties obtained when analyzing the DR1 mocks were systematically smaller than the BAO uncertainties in DESI DR1. The discrepancy could be explained by the nonlinear broadening of the BAO peak, present in data, but absent in our Gaussian mocks. Nonlinear evolution causes an increased broadening of the peak, making it more smeared compared to linear evolution and reducing the precision of its position measurement. In future analyses of DESI Lyα we plan to incorporate this effect by using more complex simulations based on perturbation theory. Meanwhile, in this publication we have used a modified version of the LyaCoLoRe mocks with broadened BAO features. Instead of using the linear power spectrum to generate the Gaussian density field, we use now as input a *quasilinear* power spectrum that is similar to the one described in Sec. III D and used in the BAO fits. In particular we use an isotropic version of Eq. (12), evaluated at $k_{\parallel} = 0$ and with $\Sigma_{\perp} = 3.26h^{-1}$ Mpc. Given that the velocity fields in CoLoRe are derived from the Gaussian density field, and are therefore also affected by this smoothing, the final nonlinear broadening of the BAO peak is somewhat anisotropic. In order to fully capture the anisotropic broadening of the peak one would need to use mocks with more realistic clustering, such as Lagrangian perturbation theory-based mocks, something that we leave for future work. However, as we discuss in Sec. IV, the BAO uncertainties from the CoLoRe-QL mocks are already in better agreement with the ones measured in DESI DR2, while the SACLAY mocks (that have not been updated) still show smaller BAO uncertainties.

As shown in Fig. 6 of [34], the clustering of quasars in the LyaCoLoRe mocks used to validate the analysis of eBOSS DR16 and in DESI DR1 did not reproduce well the clustering seen in the quasar catalogs from the more realistic Abacus simulations [35], which we take as our benchmark for comparison. In order to improve this, the CoLoRe-QL mocks use a different biasing model than the exponential one used in the DR1 mocks [32]. We use a linear biasing model, where the density of quasars is proportional to the (log-normal) density of matter, but modified such that quasars can only populate cells with a (log-normal) matter density threshold following:

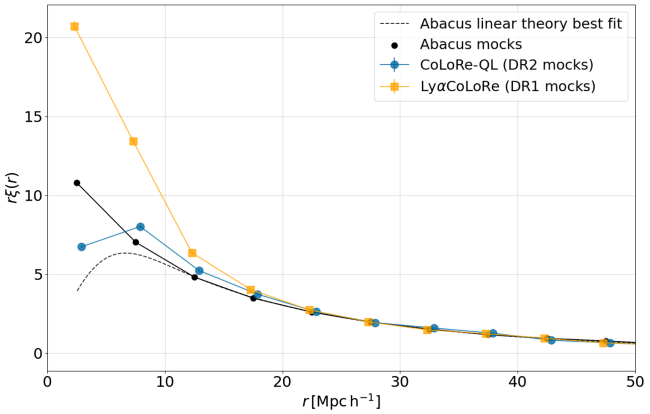

FIG. 1. Comparison of the QSO autocorrelation function measured in the LyaCoLoRe mocks used in DR1 (yellow), the CoLoRe-QL mocks presented here (blue), and the quasar catalogs from the Abacus simulation (black) at $z = 2.5$. The dashed black line is the best-fit linear theory to the Abacus mocks.

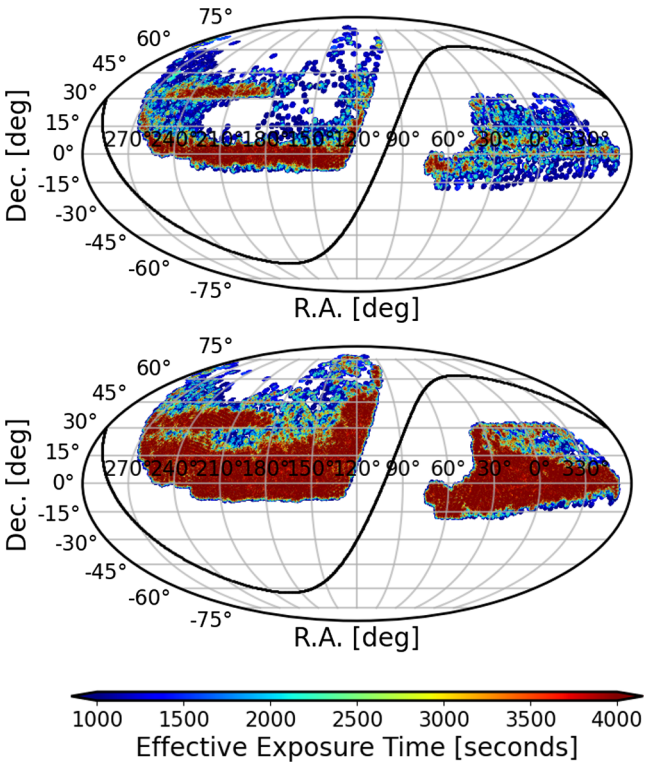

FIG. 2. Comparison of the footprints in DESI DR1 (top) and DR2 (bottom). The colors show the effective exposure time.

$$1 + \delta_Q \propto 1 + b_Q \delta_M \quad \text{for } \delta_M > t$$
$$1 + \delta_Q = 0 \quad \text{for } \delta_M \leq t \tag{1}$$

where t is the threshold and δ_Q and δ_M represent the quasar and matter overdensities, respectively. The values of the linear bias and of the threshold density were tuned to match

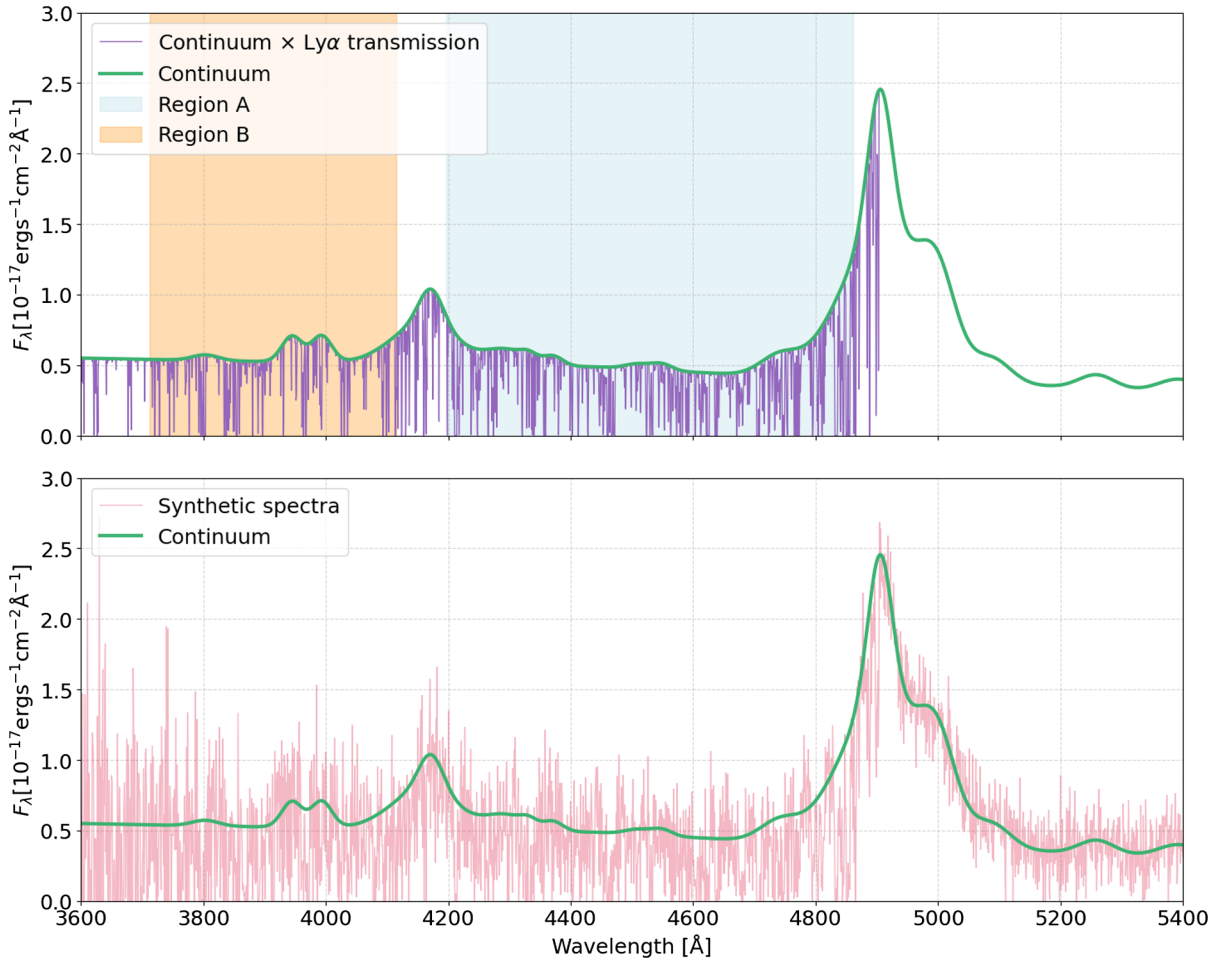

FIG. 3. Simulated spectra from a CoLoRe-QL mock at $z = 3.03$. Top: the green curve represents the quasar continuum as calculated by QUICKQUASARS, while the purple lines show the Lyα absorption, obtained by multiplying the transmission flux fraction with the quasar continuum. The shaded regions, A (blue) and B (orange), highlight the different areas used for calculating the correlations. Bottom: the pink line illustrates the simulated flux after incorporating instrumental noise using SPECSIM.

the amplitude of quasar clustering on linear scales from a preliminary, blinded measurement of quasar clustering in DESI DR1. Both parameters vary as functions of redshift, and their precise configuration details are available in the corresponding configuration file.[4] In Fig. 1, we compare the quasar clustering from LyaCoLoRe, CoLoRe-QL, and Abacus mocks. The results show that the CoLoRe-QL mocks are in better agreement with the Abacus mocks down to smaller scales.

B. Synthetic quasar spectra for DESI DR2

After generating the quasar positions and the skewers of Lyα transmission described in Sec. II A, referred to here as the "raw mocks," we use the DESISIM[5] package to emulate the characteristics of the DESI DR2 Lyα quasar sample as closely as possible.

The first step is to downsample the simulated quasar catalog to match the footprint and redshift distribution of quasars in DR2. Next, we randomly assign an r-band magnitude and number of exposures to each quasar following the observed distributions, as described in Sec. 2.2 of [27]. Fig. 2 shows a comparison of the footprints in DR1 and DR2 of DESI. Since our mocks do not have gravitationally collapsed objects, we mimic the impact of nonlinear peculiar velocities by adding random shifts to the quasar redshifts, following a Gaussian distribution with a dispersion of $\sigma_{\mathrm{FoG}} = 150\ \mathrm{km\,s^{-1}}$, which is the expected velocity dispersion for halos with $\sim 10^{12} h^{-1} M_\odot$ characteristic mass hosting quasars at $z > 2.0$ [36].

The second step is to generate synthetic spectra for each of the simulated quasars. For this we use the script QUICKQUASARS,[6] described in detail in [37]. The Lyα skewers with the transmitted flux fraction a shown in

[4]The configuration files can be accessed at https://github.com/igmhub/LyaCoLoRe/tree/colore-ql/colore-ql.
[5]https://github.com/desihub/desisim.
[6]https://github.com/desihub/desisim/blob/main/py/desisim/scripts/quickquasars.py.

the top panel of Fig. 3 are modified to include various astrophysical contaminants, such as high column density systems (HCDs), broad absorption lines (BALs) and other transition lines (metals). These are then multiplied by random (unabsorbed) quasar spectra (quasar continua), which are generated following the SIMQSO[7] package [38]. Finally, we use the SPECSIM[8] package [39] to add instrumental noise representative of DESI during nominal dark-time conditions. An example of a mock spectrum after adding instrumental noise is shown on the bottom panel of Fig. 3. For further details on the synthetic spectra generation procedure, we refer the reader to Secs. 2 and 3 of [37].

As discussed in [26], approximately 15% of the observed quasars in DR2 were originally targeted as emission line galaxies (ELGs). These have a different redshift-magnitude distribution and only have one observation (Lyα quasars in DESI can have up to four observations). These quasars were not included in the DR1 mocks presented in [27]. In the DR2 mocks we have included these extra quasars, using similar recipes to the ones used to generate the quasar targets and ignoring the fact that quasars targeted as ELGs could have different spectral properties than those targeted as quasars.

After generating the mock spectra, we create in post-process a quasar catalog where we additionally include a random error on the reported redshift that emulates the statistical uncertainties from quasar redshift estimation algorithms such as Redrock [40,41] or QuasarNET [42]. We add these errors following a Gaussian distribution with a dispersion of $\sigma_z = 400\ \mathrm{km\,s^{-1}}$, based on the measurements from [43]. Redshift errors cause a contamination in the measured correlations that was first identified in [34], and their impact on the measurement of BAO was partly discussed in Sec. IV.4 of the paper describing the validation of the DR1 Lyα analysis [27]. In this work, we include these redshift errors in the baseline configuration of the analysis. We discuss the results and propose an strategy to mitigate their effect in Sec. V.

III. METHODOLOGY

This work aims to validate the BAO analysis of Lyα forest measurements from DESI DR2 using the mocks described in Sec. II. In this section, we outline the entire process from simulated spectra to BAO measurements. We describe our method for extracting Lyα forest catalogs in Sec. III A, computing the Lyα flux overdensity field in Sec. III B, estimating Lyα correlation functions in Sec. III C, and modeling these correlations in Sec. III D. Finally, we explain how the BAO parameters are extracted by fitting the model to the measured correlation functions in Sec. III E. The first three steps are implemented using the publicly available PICCA,[9] while the modeling and fitting are conducted with the VEGA[10] package. The methodology followed here is similar to the one used in the analysis of the DR1 mocks [27], but we have introduced a couple of changes that we highlight below.

A. The Lyα forest catalog

In this section, we summarize the procedure for deriving a Lyα forest catalog from simulated quasar spectra. We begin by applying two observed-frame wavelength cuts to the spectra. The lower limit, $\lambda_{\mathrm{min}} = 3600$ Å, corresponds to the minimum wavelength detected by the blue arm of the DESI spectrographs, while the upper limit, $\lambda_{\mathrm{max}} = 5772$ Å, marks the midpoint of the overlap region between the blue and red arms of the spectrographs.

We continue by extracting the Lyα forest from two distinct regions of the spectra. As shown in Fig. 3, region A spans the rest-frame wavelength range $\lambda_{\mathrm{RF}}[\text{Å}] \in [1040, 1205]$, which is delimited by Lyα and Lyβ lines, while region B covers $\lambda_{\mathrm{RF}}[\text{Å}] \in [920, 1020]$. The rationale for using two regions is that region A contains absorption exclusively from the first line of the Lyman series, Lyα, while region B also includes higher-order absorption lines such as Lyβ and Lyγ. Although both regions are affected by absorption from other transition lines (metals), we assume all absorption is due to Lyα when measuring the correlation functions (Sec. III C) and later model their impact on these (Sec. III D).

We use "Lyα pixels" to refer to those pixels in the rest-frame wavelength regions defined above, and we refer to the collection of all Lyα pixels within a quasar spectrum as a forest. The combination of the observed- and rest-frame selection criteria constrain the redshift range of quasars contributing Lyα pixels in the region A to $z \in [2.1, 4.4]$ and in the region B to $z \in [2.6, 5.1]$. The number of quasars contributing to regions A and B in the mocks is approximately 800 000 and 360 000, respectively.

Before extracting the Lyα overdensity field, we apply corrections to mitigate the impact of two astrophysical contaminants: damped Lyα systems and broad absorption line quasars. We first identify pixels associated with DLAs detected at a signal-to-noise ratio greater than 2 and with column densities $N_{\mathrm{HI}} > 10^{20.3}\ \mathrm{cm^{-2}}$. The DLA detection algorithm achieves a completeness of 75% for such systems [44]. To replicate the impact of DLAs in observations, we randomly mask the same fraction of DLAs that satisfy these criteria in the mocks. Specifically, we mask all DLA pixels where the transmitted flux decreases by 20% or more and correct the transmitted flux of the remaining DLA pixels using a Voigt profile [45] (see also Appendix B). For the validation of the DR1 measurement using mocks, we

[7] https://github.com/imcgreer/simqso.
[8] https://github.com/desihub/specsim.
[9] https://github.com/igmhub/picca.
[10] https://github.com/andreicuceu/vega.

assumed 100% completeness for the DLA finder, thereby masking all DLAs in the mocks [27].

We mask the expected locations of all potential BAL features, regardless of whether the absorption is apparent [46]. These features include Lyα, N IV, C III, Si IV, and P V in region A and O vi, O i, Lyβ, Lyγ, N iii, and Lyδ in region B. Finally, we discard forests with fewer than 150 pixels, which is a threshold required by the continuum fitting procedure (see next section).

B. The flux transmission field

In this section, we provide a brief overview of our method for measuring the Lyα flux transmission field. We follow the same approach used for analyzing DR1 measurements; see [20,27], for more details.

We begin by calculating the Lyα flux overdensity field, given by

$$\delta_q = \frac{f_q(\lambda)}{\bar{F}(z)C_q(\lambda)} - 1, \tag{2}$$

where f_q is the observed flux density for a quasar q, C_q is the unabsorbed flux density (also referred to as the quasar continuum), $\bar{F}$ is the mean transmission of the intergalactic medium (IGM) at the absorber redshift $z = \lambda/\lambda_\alpha - 1$, and λ_α is the rest-frame wavelength of Lyα.

Measuring the flux overdensity field requires estimating the product $\bar{F}(z)C_q(\lambda)$ for each quasar. This step is known as continuum fitting, which is described in detail in [47,48]. Here, we provide a brief overview of the process, noting that we perform it separately for regions A and B. First, we approximate the product $\bar{F}(z)C_q(\lambda)$ for each quasar as the product of $\bar{C}$, which is the same for all quasars, and a quasar-specific first-degree polynomial in $\Lambda = \log\lambda$,

$$E_q^X = \bar{C}^X(\lambda_{\rm RF})\left(a_q^X + b_q^X \frac{\Lambda - \Lambda_{\rm min}}{\Lambda_{\rm max} - \Lambda_{\rm min}}\right), \tag{3}$$

where X refers to regions A and B, $\lambda_{\rm RF} = \lambda/(1+z_q)$, and $\Lambda_{\rm max}$ and $\Lambda_{\rm min}$ correspond to the maximum and minimum observed-frame wavelengths of the forest for the quasar q. We fit a_q^X and b_q^X by maximizing the likelihood function

$$\ln\mathcal{L} = -0.5\sum_i \left[\frac{f_q(\lambda_i) - E_q^X(\lambda_i)}{\sigma_q^X(\lambda_i, a_q^X, b_q^X)}\right]^2 - \sum_i \ln\sigma_q^X(\lambda_i, a_q^X, b_q^X), \tag{4}$$

where $(\sigma_q^X)^2 = \eta_{\rm pip}^X(\lambda)\sigma_{{\rm pip},q}^2(\lambda) + [\sigma_{\rm LSS}^X(\lambda)E_q^X(\lambda)]^2$ is the flux variance of each pixel, $\sigma_{{\rm pip},q}^2$ is the pipeline estimate for the flux variance, $\eta_{\rm pip}^X$ is a correction factor for inaccuracies in this estimate, and $\sigma_{\rm LSS}^X$ represents the intrinsic standard deviation of fluctuations in the Lyα forest. Continuum fitting is an iterative process that begins by assuming initial values for $\bar{C}^X$, $\eta_{\rm pip}^X$, and $\sigma_{\rm LSS}^X$. The coefficients a_q^X and b_q^X are then computed for each region of all quasars. Next, we estimate the Lyα flux overdensity field for all forests, calculate its variance, and adjust the values of $\eta_{\rm pip}^X$ and $\sigma_{\rm LSS}^X$. Finally, we measure $\bar{C}^X$ and repeat the whole process until convergence, which typically takes five steps.

As first discussed in [49,50], when fitting the a_q^X and b_q^X coefficients for each quasar we are suppressing very long wavelength fluctuations in the Lyα forest, distorting the measured correlations. In order to make the modeling of the distortion easier, we follow [51] and we explicitly subtract the mean and first moment from each forest. This is done by applying the projection operator ζ to the fluctuation field where

$$\zeta_{ij}^q = \delta_{ij}^K - \frac{w_j}{\Sigma_k w_k} - \frac{w_j(\Lambda_i - \bar{\Lambda}_q)(\Lambda_j - \bar{\Lambda}_q)}{\Sigma_k w_k(\Lambda_k - \bar{\Lambda}_q)^2} \tag{5}$$

and redefining the fluctuations as

$$\tilde{\delta}^q(\lambda_i) = \Sigma_j \zeta_{ij}^q \delta^q(\lambda_j) \tag{6}$$

where δ^K is the Kronecker delta and $\bar{\Lambda}_q$ is the mean of Λ for quasar q. The weights are identical to those used in the correlation measurement [see Eq. (10)]. For simplicity we will refer to this projected density fluctuation as δ in the rest of the text.

C. Measuring the correlation functions

There are six possible correlations involving quasar positions and Lyα fluctuations in regions A and B. However, following recent Lyα BAO analyses [20,47], we focus on those with the highest signal-to-noise ratio: the autocorrelation of Lyα fluctuations in region A, $\xi^{\rm AA}$; the cross-correlation of Lyα fluctuations between regions A and B, $\xi^{\rm AB}$; and the cross-correlation of quasar positions with Lyα fluctuations in regions A and B, denoted as $\xi^{\rm QA}$ and $\xi^{\rm QB}$, respectively. In this section, we briefly outline our methodology for measuring these correlations, which is the same as the one used in the DR1 analysis and its validation. See [27,52] for further details.

We compute the correlation functions on a grid of separations, $r_\parallel$ and $r_\perp$, corresponding to distances along and across the line of sight, respectively. To define these components we adopt the angular bisector LOS convention as explained in [53]. This computation requires assuming a fiducial cosmology to convert angular and redshift differences into comoving separations

$$r_\parallel = [D_C(z_i) - D_C(z_j)]\cos\left(\frac{\theta_{ij}}{2}\right), \tag{7}$$

$$r_\perp = [D_M(z_i) + D_M(z_j)] \sin\left(\frac{\theta_{ij}}{2}\right), \quad (8)$$

where i and j index pixel-pixel or pixel-quasar pairs, pixel redshifts are determined by assuming Lyα absorption as $z_i = \lambda/\lambda_\alpha - 1$, θ_{ij} refers to the angular separation, and D_C and D_M denote the comoving and angular comoving distances, respectively, which are identical in a flat Universe. We use the best-fitting flat ΛCDM model from Planck 2015 results [54] as our fiducial cosmology, the same model used to generate the LyaCoLoRe mocks.[11]

The algorithms to compute the auto- and cross-correlations can be expressed in the following compact form:

$$\hat{\xi}_M^{XY} = \frac{\sum_{i,j\in M} w_i^X w_j^Y \hat{\delta}_i^X \hat{\delta}_j^Y}{\sum_{i,j\in M} w_i^X w_j^Y}, \quad (9)$$

where X and Y correspond to the auto- and cross-correlation cases defined at the beginning of this section, M represents a two-dimensional bin with widths $(\Delta r, \Delta r)$, and $\delta^Q = 1$. The weight assigned to quasar i is given by $w_i^Q = [(1+z_i)/(1+z_Q)]^{\gamma_Q - 1}$, where $z_Q = 2.25$ and $\gamma_Q = 1.44$ are derived from previous analyses [56]. The weights for Lyα fluctuations in regions A and B are

$$w_i^X = \frac{[(1+z_i)/(1+z_{\rm fid})]^{\gamma_\alpha - 1}}{\eta_{\rm pip}^X(\lambda)\sigma_{{\rm pip},q}^2(\lambda)[E_q^X(\lambda)]^{-2} + \eta_{\rm LSS}[\sigma_{\rm LSS}^X(\lambda)]^2}, \quad (10)$$

where $\gamma_\alpha = 2.9$ [57]. The term $\eta_{\rm LSS} = 7.5$ enhances the contribution of the intrinsic Lyα forest variance relative to the pipeline noise, which improves the precision of the correlation function [48].

We measure the correlation functions using bins of $\Delta r = 4h^{-1}$ Mpc in both the radial and perpendicular separations. For the autocorrelations, we use 50 bins spanning separations from 0 to $200h^{-1}$ Mpc in both directions. For the cross-correlations, we differentiate between pixels located in front of ($r_\parallel < 0$) and behind ($r_\parallel > 0$) quasars. With this distinction, we measure the cross-correlations using 50 bins for the perpendicular separation (0 to $200h^{-1}$ Mpc) and 100 bins for the parallel separation (from -200 to $200h^{-1}$ Mpc).

We compute the covariance matrix associated with the correlations using the same approach as in previous DESI Lyα analyses [27]. First, we divide the mock survey footprint into HEALPix pixels [58] with $N_{\rm side} = 16$, each covering approximately $3.7 \times 3.7 = 13.4$ deg^2 of the sky. We then compute the correlation functions independently for each of the 1028 HEALPix pixels spanned by the DESI DR2 dataset. After that, we estimate the covariance matrix from these measurements as follows:

$$C_{MN} = \frac{1}{W_M W_N} \sum_s W_M^s W_N^s (\hat{\xi}_M^s \hat{\xi}_N^s - \hat{\xi}_M \hat{\xi}_N), \quad (11)$$

where $W_M^{XY,s} = \sum_{ij\in(M,s)} w_i^X w_j^Y$ represents the summed weight for a correlation in subsample s, $W^s = \{W^{{\rm AA},s}, W^{{\rm AB},s}, W^{{\rm QA},s}, W^{{\rm QB},s}\}$ is a vector containing the summed weights for all correlations, $\hat{\xi} = \{\xi^{\rm AA}, \xi^{\rm AB}, \xi^{\rm QA}, \xi^{\rm QB}\}$ refers to a vector containing the correlations, and $W_M = \sum_s W_M^s$ denotes the total weight.

The covariance matrix in our fiducial analysis has dimensions $15\,000 \times 15\,000$, making its estimation inherently noisy due to the limited number of subsamples. To mitigate this noise, we smooth the covariance matrix of each mock using the same procedure as in previous analyses [27]. This smoothing method is also applied when computing the covariance of the stack of multiple mocks.

In Fig. 4, we compare the four correlation functions derived from DESI DR2 observations (data points) with the averages of 300 CoLoRe-QL mocks (solid lines) and 100 Saclay mocks (dashed lines). The measurements are presented as a function of radial separation, $r = (r_\parallel^2 + r_\perp^2)^{1/2}$, and the wedges defined by $\mu = r^{-1} r_\parallel$. Overall, the mock measurements agree well with the observational data across the entire range of scales, supporting the use of these mocks for assessing the performance of the correlation model (see Sec. III D) in estimating the BAO position. However, we can readily see minor discrepancies between the observed correlations and the best-fit models obtained from the stack of mocks. For example, in the autocorrelations, particularly in the line-of-sight wedge, both CoLoRe-QL and the Saclay mocks deviate from observed data. We attribute these differences primarily to contaminants: DLAs, redshift errors, and metal contamination, which are known to affect the shape of the correlation functions significantly.

First, DLAs introduce a scale-dependent excess in the Lyα correlation functions [59–61]. In our mocks we mask 75% of the $\log(\frac{N_{\rm HI}}{cm^{-2}}) > 20.3$ DLAs informed by finding algorithms applied in a mock dataset (see Appendix B and [44]). However, the real performance in observed data is difficult to determine, thus the shape of the observed correlation functions is affected by a possible discrepancy in the assumed and the real performance. Second, redshift errors introduce spurious correlations along the line of sight [34,62], the distribution of redshift errors in observed data is undetermined, and therefore the errors introduced into our mock catalogs might not be representative of the reality. Finally, metals introduce spurious correlations along the line of sight (see [37] for a brief discussion). These were tuned in our mocks to reach the linear bias values found in DESI DR1 [20,27]. Although the metal biases found in DR1 and DR2 are statistically consistent they differ and therefore the amplitude of the spurious features introduced by metals might differ with the observed correlations.

[11]Note that for the analysis of observational measurements we use as fiducial cosmology Planck 2018 [55].

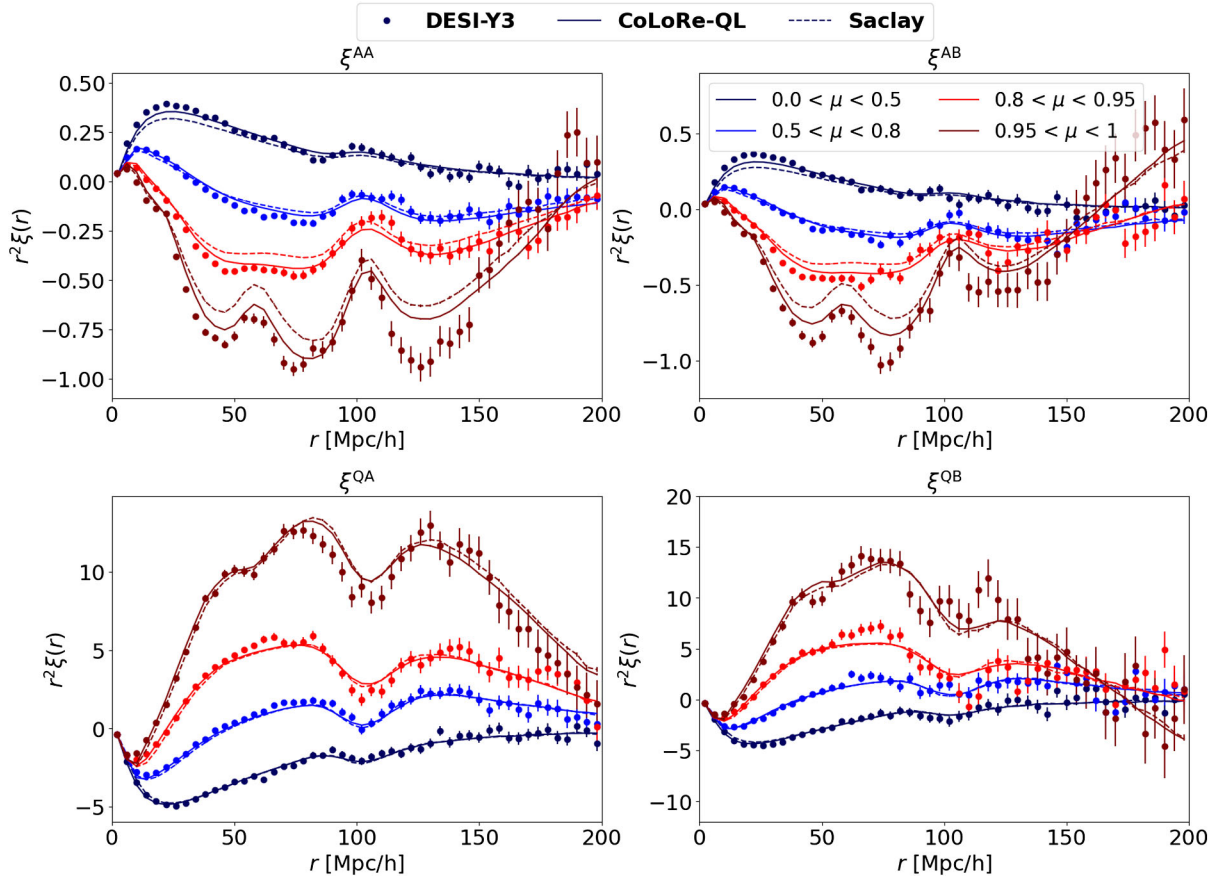

FIG. 4. Comparison between DESI DR2 observational measurements and mock predictions. The panels, in clockwise order, display the following: the autocorrelation of the Lyα forest in region A, the autocorrelation of the Lyα forest in regions A and B, the cross-correlation of the Lyα forest in region B with quasar positions, and the cross-correlation of the Lyα forest in region A with quasar positions. Data points and error bars represent the observational measurements and their associated uncertainties, while the solid and dashed lines represent the results from the average of 300 CoLoRe-QL and 100 Saclay mocks, respectively. The colors of the lines indicate the average correlations in different wedges. Overall, the mocks show good agreement with the observational measurements; however, their accuracy diminishes for the line-of-sight wedge in the autocorrelations.

Finally, other differences, seen in other wedges, or in the cross-correlations can be attributed to the clustering model assumed in our mocks. CoLoRe-QL mocks were created to reach bias values found in eBOSS DR12 (see [30]), while Saclay mocks were created to match eBOSS DR16 (see [31]), moreover the small-scale clustering of quasars in both mocks differ between each other (see Sec. II) and are yet to be refined in future realizations of mocks.

D. Modeling the correlation function

Our goal is to measure the BAO scale both along and across the line of sight. To this end, we use a model first introduced for the analysis of the data release 9 of BOSS [4–6] and progressively improved over the analysis of other data releases of BOSS [7,51,56,63], eBOSS [47,64,65], and DESI [20,27,52]. In what follows, we progressively build this model and detail some differences compared to DESI DR1.

We begin with the isotropic linear matter power spectrum $P_{\rm fid}$ computed evaluating CAMB [66] for the same fiducial cosmology used to convert angles and redshifts into comoving coordinates in the previous section. Next, we decompose $P_{\rm fid}$ into an oscillatory component $P_{\rm BAO}$, which contains the BAO signal, and a smooth component $P_{\rm smooth}$, using the algorithm described in [6]. To account for nonlinear effects, we apply an anisotropic Gaussian damping to the oscillatory component

$$P_{\rm QL}(k_\parallel, k_\perp, z) = P_{\rm BAO}(k, z) \exp\left[-\frac{k_\parallel^2\Sigma_\parallel^2 + k_\perp^2\Sigma_\perp^2}{2}\right] + P_{\rm smooth}(k, z) \tag{12}$$

where $k = (k_\parallel^2 + k_\perp^2)^{1/2}$ is the wave vector, with $k_\parallel$ and $k_\perp$ as the parallel and perpendicular components, and where $\Sigma_\parallel$ and $\Sigma_\perp$ control the nonlinear broadening of BAO along and perpendicular to the line of sight.

Building on the previous expression, we construct a model for the anisotropic power spectrum of the autocorrelations

($X = Y = \alpha$) and the cross-correlations ($X = \mathrm{Q}$, $Y = \alpha$) as follows:

$$\begin{aligned} P_{XY}(k,\mu_k,z_{\mathrm{eff}}) &= b_X(z_{\mathrm{eff}})[1+\beta_X(z_{\mathrm{eff}})\mu_k^2] \\ &\times b_Y(z_{\mathrm{eff}})[1+\beta_Y(z_{\mathrm{eff}})\mu_k^2] \\ &\times G(k,\mu_k)F_{\mathrm{NL}}^{XY}(k,\mu_k) \\ &\times F_{\mathrm{SM}}(k,\mu_k)P_{\mathrm{QL}}(k,\mu_k,z_{\mathrm{eff}}), \end{aligned} \tag{13}$$

where $\mu_k = k^{-1}k_\parallel$ denotes the cosine of the angle between the wave vector and the line of sight and the factor $b(1+\beta\mu_k^2)$, evaluated at the effective redshift of the measurement ($z_{\mathrm{eff}} = 2.33$), accounts for the impact of linear bias and large-scale redshift-space distortions [67].

As discussed later in this section, the model correlation function $\xi(r_\parallel, r_\perp)$ is calculated on a two-dimensional grid in configuration space, where each bin $r_\parallel$, $r_\perp$ has a slightly different weighted mean redshift. Following the procedure outlined in Sec. 4.2 of [20], we start by defining the power spectrum at an effective redshift z_{eff} [see Eq. (13)]; we then compute its inverse Fourier transform to obtain the correlation function at the same effective redshift; finally, the amplitude of the correlation function is rescaled for each measured bin according to the redshift evolution of the bias parameters: $b_X(z) = b_X(z_{\mathrm{eff}})[(1+z)/(1+z_{\mathrm{eff}})]^{\gamma_X}$. We do not include an explicit redshift evolution for β. Since the mean redshifts of the bins differ only slightly from z_{eff}, this approximation is well justified. Note that we describe the correlations involving both Lyα regions (A and B) with the same model.

The term $G(k,\mu_k) = \mathrm{sinc}(k_\parallel \Delta r/2)\mathrm{sinc}(k_\perp \Delta r/2)$ accounts for the binning of the correlation function on a grid [51]. Following [27], in order to account for the finite resolution of the mocks we multiply the power spectra by a Gaussian anisotropic smoothing factor: $F_{\mathrm{SM}} = \exp -[(k_\parallel\sigma_\parallel)^2 + (k_\perp\sigma_\perp)^2]/2$. F_{NL} accounts for nonlinear corrections that are limited to relatively small scales. When modeling the cross-correlations, we include the terms $F_{\mathrm{NL}}^{\mathrm{Q}\alpha} = \exp[-(k_\parallel\sigma_v)^2/2]$ to model the combined impact of quasar redshift errors and nonlinear peculiar velocities. When modeling the autocorrelation from observational data, we introduce a term $F_{\mathrm{NL}}^{\alpha\alpha}$ to account for the effects of nonlinear growth, peculiar motions, and thermal broadening on the autospectrum, using a model as specified by Eq. (3.6) in [68],[12] that accounts for nonlinear growth of structure and thermal effects from the IGM. These effects are not included in the mocks, and therefore we ignore them when modeling the correlations measured from mocks.

Finally, we account for potential systematic errors in quasar redshift measurements by introducing a systematic shift ($\Delta r_\parallel$) in the line-of-sight separation between quasars and Lyα absorption [43,69].

Next, we transform the anisotropic power spectra from Fourier to configuration space. This is done by first performing a multipole decomposition of the anisotropic power spectrum up to $\ell = 6$, followed by a Hankel transform using the FFTLog algorithm [70] implemented in the MCFIT package[13] to obtain the correlation function multipoles. From these multipoles, we then compute the two-dimensional correlation function ξ_M^{XY} on a grid with spacing $\Delta r = 2h^{-1}$ Mpc. We verified that including higher-order multipoles ($\ell > 6$) does not affect BAO measurements. Additionally, we use a finer grid than the one used for measuring the correlation functions, $\Delta r = 4h^{-1}$ Mpc, to improve the precision of the model.

Next, we model the impact of astrophysical contaminants and systematics on the correlation functions. Given the range of separations considered in our analysis, we account for the influence of four metal absorption lines: SiII(1190), SiII(1193), SiIII(1207), and SiII(1260). This requires computing correlation function models for all Lyα -metal, QSO-metal, and metal-metal correlations using nearly the same framework as for Lyα-only correlations. Each metal line is characterized by parameters b_i and β_i, where i runs over the four metal lines. Following [52], we fix $\beta_i = 0.5$ and allow b_i to vary freely in the fits. This leads to the following expressions for the auto- and cross-correlations:

$$\tilde{\xi}^{\alpha\alpha} = \xi^{\alpha\alpha} + \sum_i \xi^{\alpha i} + \sum_{ij} \xi^{ij} + \xi^{\mathrm{inst}}, \tag{14}$$

$$\tilde{\xi}^{\mathrm{Q}\alpha} = \xi^{\mathrm{Q}\alpha} + \sum_i \xi^{\mathrm{Q}i} + \xi^{\mathrm{TP}}, \tag{15}$$

where ξ^{inst} and ξ^{TP} represent the contributions from DESI instrumental systematics and quasar radiation effects, respectively. However, since these effects are not included in the mocks, we chose not to account for them in the analysis (see also [27]).

For the auto- and cross-correlations with metals, we account for the misestimation of pixel redshifts due to the assumption that all absorption originates from Lyα. Specifically, we compute the redshift of each Lyα pixel as $z = \lambda/\lambda_\alpha - 1$, whereas some absorptions are actually caused by metal lines, meaning their true redshifts should be $z = \lambda/\lambda_{\mathrm{metal}} - 1$. As a result, these pixels are assigned to incorrect bins in the correlation function [51,59]. We model this contamination with the same method used in the DR1 analysis [20], with a minor modification discussed in the companion paper [26].

As discussed in Sec. III A, we mask pixels that are most contaminated by DLAs. However, as discussed in [44],

[12]Since these only affect the correlations on small scales, we do not vary the parameters of this model in BAO analyses [26].

[13]https://github.com/eelregit/mcfit.

we only mask DLAs detected in spectra with relatively high signal-to-noise ratio (SNR > 2), where the efficiency of the DLA finders is high, finding roughly 75% of the DLA systems.[14]

Additionally, there are intermediate column density systems not identified by the DLA-finding algorithm that contribute to contamination. Following [50,60], we account for the impact of these contaminants by adding a scale dependence to the bias and redshift-space distortion parameters of Lyα fluctuations to include the contributions from both Lyα absorption and HCDs,

$$\tilde{b}_\alpha = b_\alpha + b_{\rm HCD} F_{\rm HCD}(k_\parallel), \tag{16}$$

$$\tilde{b}_\alpha \tilde{\beta}_\alpha = b_\alpha \beta_\alpha + b_{\rm HCD} \beta_{\rm HCD} F_{\rm HCD}(k_\parallel), \tag{17}$$

where $b_{\rm HCD}$ and $\beta_{\rm HCD}$ represent the contributions from HCD systems, while $F_{\rm HCD}(k_\parallel) = \exp(-L_{\rm HCD} k_\parallel)$ depends on the column density distribution of the HCDs present in the data.

The continuum fitting process (see Sec. III B) notably alters the shape of the measured correlation functions [49,50]. We use the formalism introduced in [51], that multiplies the modeled correlations $\tilde{\xi}$ by a "distortion matrix" D^{XY},

$$\hat{\xi}^{XY}_M = \sum_N D^{XY}_{MN} \xi^{XY}_N. \tag{18}$$

The distortion matrix can be computed from the geometry of the dataset and the distribution of weights, and in the DR2 analysis we have implemented an improved computation that takes into account the redshift evolution of the signal (see [26] for a detailed explanation and a discussion on the impact of this change).

E. Fitting the correlation function

By measuring the position of the BAO feature along and across the line of sight, we can constrain the distance ratios D_M/r_d and D_H/r_d, where r_d is the sound horizon at recombination, and $D_H(z) = c/H(z)$. In order to isolate this information from the rest of the measured correlations, in BAO analyses it is common to introduce two BAO parameters ($\alpha_\perp$, $\alpha_\parallel$) that only shift the position of the peak, without affecting the smooth component of the model,

$$\xi^{XY} = \xi^{XY}_{\rm smooth}(r_\parallel, r_\perp) + \xi^{XY}_{\rm BAO}(\alpha_\parallel r_\parallel, \alpha_\perp r_\perp), \tag{19}$$

where $\alpha_\parallel = [r_{\rm d}^{-1} D_{\rm H}(z)]/[r_{\rm d}^{-1} D_{\rm H}(z)]_{\rm fid}$ and $\alpha_\perp = [r_{\rm d}^{-1} D_{\rm M}(z)]/[r_{\rm d}^{-1} D_{\rm M}(z)]_{\rm fid}$ are defined with respect to the distances in the fiducial cosmology, denoted by "fid."

Besides the two BAO parameters, our model to describe the correlations in mocks uses 12 nuisance parameters: three for Lyα and quasar biases and redshift-space distortions (b_α, β_α, and $b_{\rm Q}$), three for HCD systems ($b_{\rm HCD}$, $\beta_{\rm HCD}$, and $L_{\rm HCD}$), four for the linear biases of each metal line ($b_{\rm SiII(1190)}$, $b_{\rm SiII(1193)}$, $b_{\rm SiIII(1207)}$, and $b_{\rm SiII(1260)}$), and two for redshift errors ($\Delta r_\parallel$ and σ_v). We determine the best-fitting value of these parameters using IMINUIT.[15] While the BAO parameters have flat, noninformative priors, we use informative priors for some of the less constrained nuisance parameters as shown in Appendix C.

In the results presented in the following section, we hold fixed the value of the smoothing parameters [$\sigma_\parallel$, $\sigma_\perp$ in the $F_{\rm SM}$ term of Eq. (13)] that capture numerical artifacts in the simulations (these parameters are not included in the analysis of real data). In order to choose a value for these parameters, we do a preliminary fit of the mean of the 300 CoLoRe-QL and 100 SACLAY mocks, including separations from 10 to 180h^{-1} Mpc. We show the results from the fit in Fig. 5. We find the values $\sigma_\parallel = 2.0$ and $\sigma_\perp = 1.7h^{-1}$ Mpc for the CoLoRe-QL mocks, and $\sigma_\parallel = 2.2$ and $\sigma_\perp = 2.2h^{-1}$ Mpc for the SACLAY mocks. In the case of the CoLoRe-QL mocks, we also use this preliminary fit to infer the value of the nonlinear BAO broadening parameters, finding $\Sigma_\perp = 3.1$ and $\Sigma_\parallel = 4.3h^{-1}$ Mpc. These parameters are held fixed to these values in later analyses of the CoLoRe-QL mocks.[16] While the broadening measured in the transverse direction is consistent with the theoretical expectation at this redshift ($\Sigma_\perp = 3.26h^{-1}$ Mpc), the line-of-sight broadening is a bit lower than the expected one ($\Sigma_\parallel = 6.4h^{-1}$ Mpc). As discussed in Sec. II A 2, this might be caused by the simple method that we have followed to add nonlinear broadening of the BAO in the CoLoRe-QL mocks.

Even though the model is able to describe fairly well the correlations on mocks down to $r = 10h^{-1}$ Mpc, in the rest of the analysis presented here we only fit correlations in the range of separations of $r = 30$–$180h^{-1}$ Mpc (corresponding to 1590 and 3180 bins for the auto- and cross-correlations). This is the configuration that was used in the DESI DR2 Lyα BAO analysis [26] that we would like to validate.

IV. RESULTS

In this section, we describe and discuss the results obtained from performing the BAO analysis on the mocks. The mocks are used to validate the BAO analysis in two

[14] As discussed in Appendix B, the impact of the exact SNR cut on the BAO parameters is minor.

[15] We verified that the difference between the best-fitting parameters and error estimates obtained from IMINUIT (https://github.com/scikit-hep/iminuit) and the PolyChord nested sampler [71] is minimal (see also [72]).

[16] For the SACLAY mocks, generated from a linear power spectrum, these parameters are set to zero.

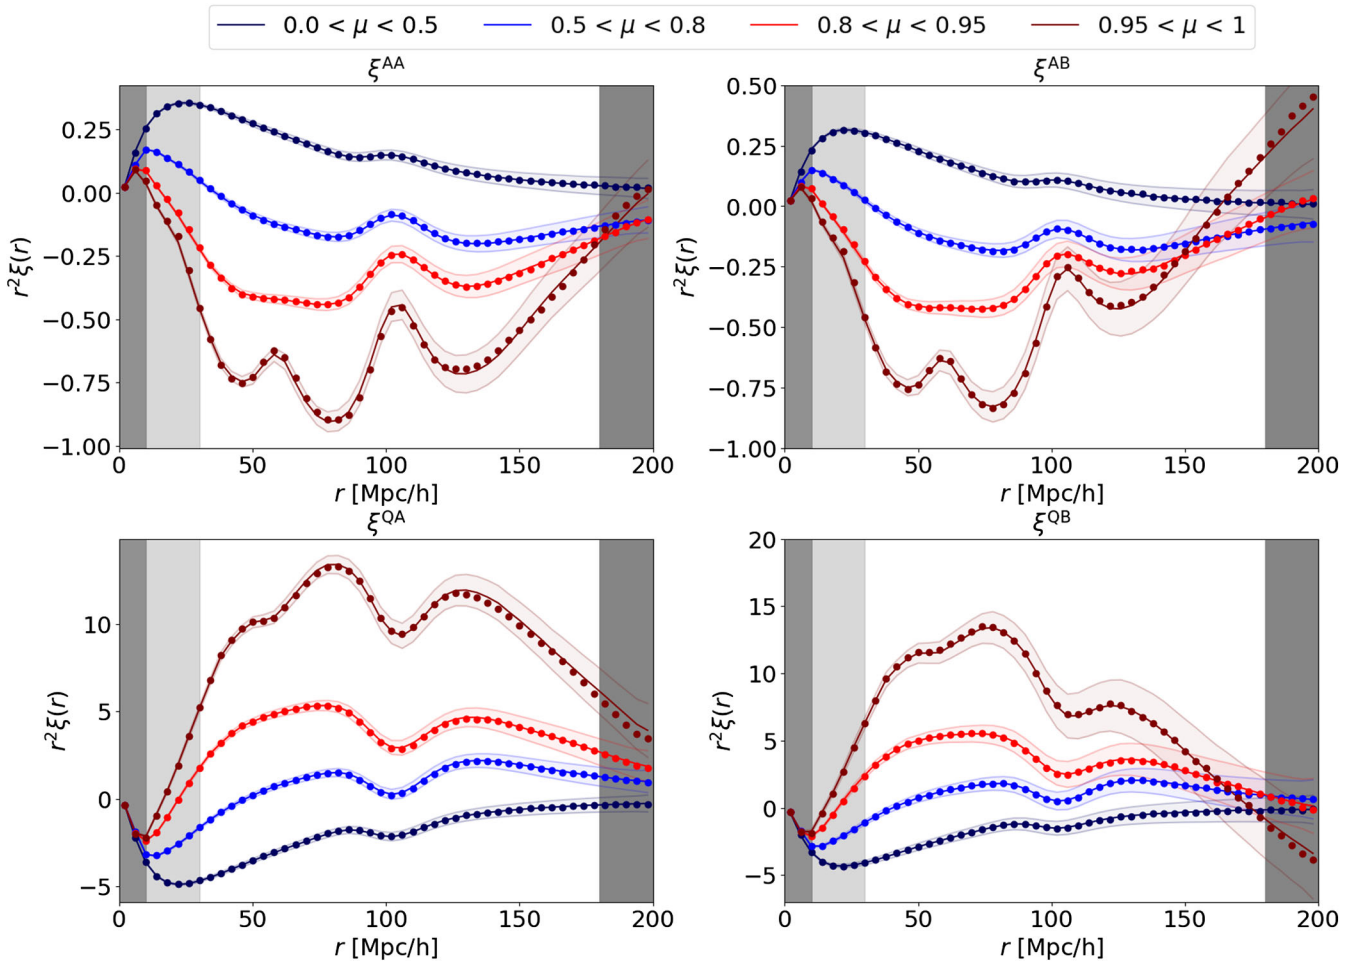

FIG. 5. Best-fitting model to the average correlation of 300 CoLoRe-QL mocks. Dots represent the simulation measurements, lines correspond to the best-fitting model, and the shaded areas around the lines indicate the diagonal elements of the covariance matrix estimated from DESI DR2 measurements. The dark-shaded areas highlight the scales not used in this particular fit, while the light-shaded are show the scales that are not used in the main BAO analysis of DR2.

ways: in Sec. IV A we use them to validate that our analysis pipeline is unbiased,[17] and in Sec. IV B we use them to validate that the reported uncertainties on the BAO parameters are representative of the scatter between the measurement in different realizations.

A. Investigating potential systematic biases in the BAO measurements

In order to validate that our analysis pipeline produces unbiased BAO measurements, we combine the measurement from all our synthetic datasets to obtain correlation functions with very high signal-to-noise ratio. Similar to the other validation tests discussed in [26], our requirement was to recover the true BAO parameters ($\alpha_\perp = \alpha_\parallel = 1$) within a third of the statistical uncertainty from DESI DR2. Biases exceeding this threshold must be investigated, corrected, or incorporated into the systematic errors.

In the top panel of Fig. 6 we show the BAO measurements from the combined (or "stacked") correlations from 100 SACLAY mocks (red) and 300 CoLoRe-QL mocks (blue), compared to the threshold of a third of the statistical uncertainty from DESI DR2 (dashed, black line). It is clear in the figure and in Table I that $\alpha_\parallel$ measurements are biased at high significance, with a bias near the threshold criterion we imposed, while $\alpha_\perp$ satisfies the criteria.

In order to distinguish between a bias in our analysis pipeline from a bias in our simulated boxes, in the bottom panel of Fig. 6 we show a simplified BAO analysis where we measure the correlations directly from the raw mocks introduced in Sec. II A. This analysis, detailed in Appendix A, does not require any continuum fitting and is not affected by instrumental noise or astrophysical contaminants. The fact that these results are unbiased relative to our threshold of $\sigma_{\rm DESI}/3$ rules out intrinsic issues with the simulated boxes. Instead, as discussed in detail in Sec. V,

[17] SACLAY mocks were generated with a cosmology slightly different than the fiducial cosmology used in the analysis, but in terms of BAO parameters the differences are smaller than 0.03%.

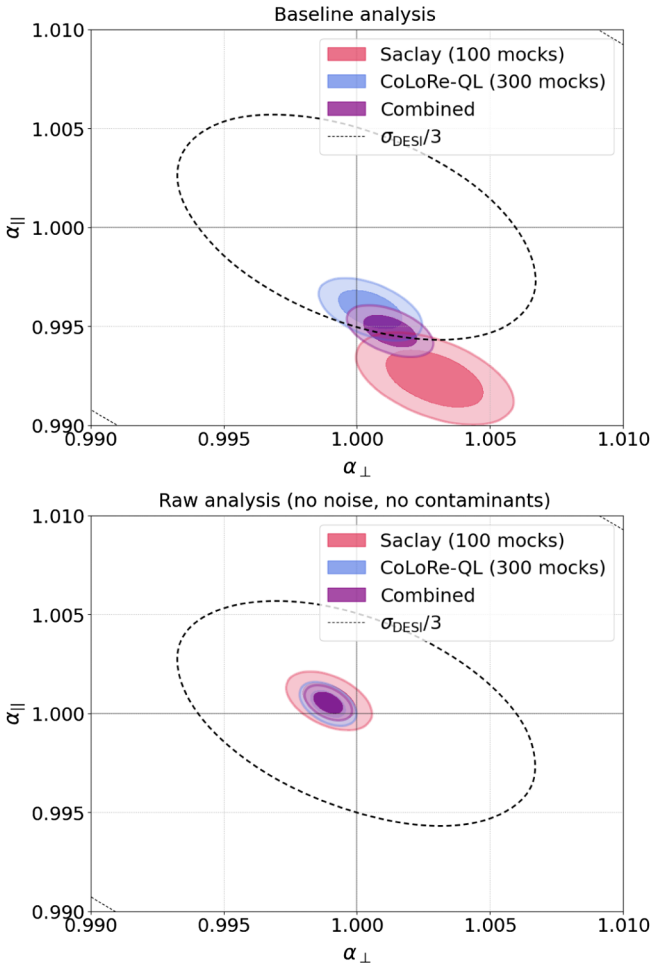

FIG. 6. Top: best fit of the stacks using the and their combination (purple). The main difference with respect to DR1 is that we do continuum fitting after adding quasar redshift errors. As pointed out in [34], this introduces a spurious correlation that can result in a small bias on the BAO results. $\sigma_{\mathrm{DESI}/3}$ represents one-third of the DR2 uncertainty. This corresponds to 0.0038 for $\sigma_{\parallel}$ and 0.0045 for $\sigma_{\perp}$. Bottom: BAO analysis on raw mocks without instrumental noise, contaminants, or continuum fitting. Both the SACLAY and CoLoRe-QL mocks show no intrinsic biases, confirming that the observed shift is induced by the analysis.

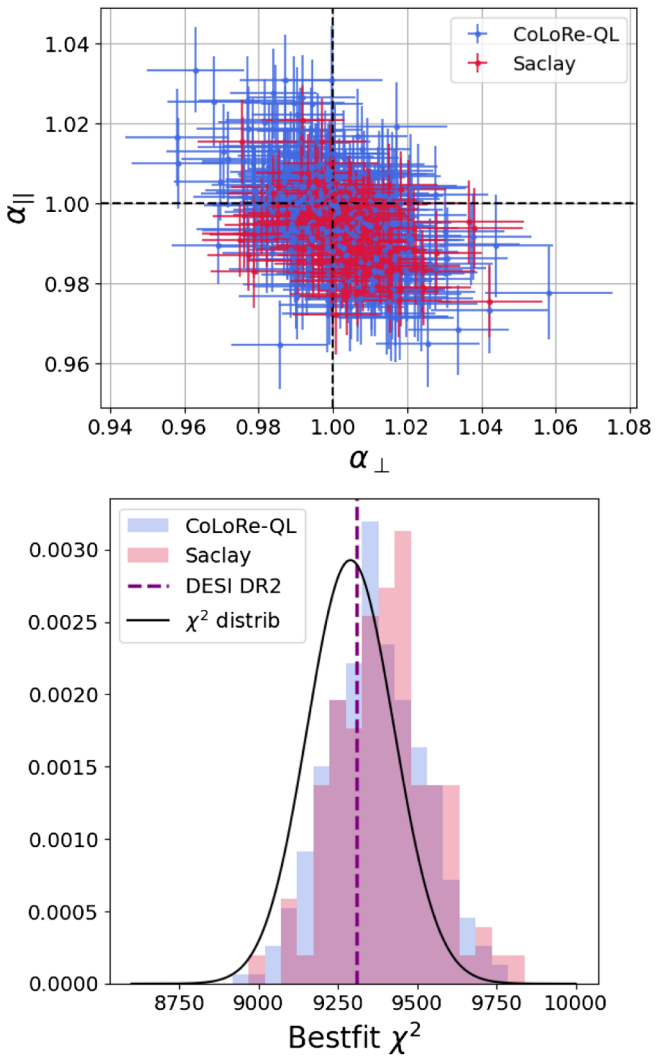

FIG. 7. Top: scatter plot of the best-fit BAO measurements obtained for each of the 100 SACLAY mocks (red) and 300 CoLoRe-QL mocks (blue). Bottom: distribution of χ^2 values from the fits, compared to the expected distribution given the number of degrees of freedom (black). The dashed purple line represents the χ^2 of the best fit to the data.

we have identified that the bias in the baseline analysis (bottom panel) arises from the method used to incorporate redshift errors in the mocks catalogs.

B. Validating the BAO uncertainties

In this section we present the BAO results obtained by individually fitting each of the DESI DR2 mocks. In the top

TABLE I. Biases on the BAO parameters ($\alpha_{\parallel}$, $\alpha_{\perp}$) from the combined analyses of 300 CoLoRe-QL and 100 SACLAY mocks, as well as their combination. The last two columns show the biases on the isotropic BAO parameter, defined here as $\alpha_{\mathrm{iso}} = \alpha_{\perp}^{9/20}\alpha_{\parallel}^{11/20}$ and the anisotropic Alcock-Paczyński parameter $\phi = \alpha_{\perp}/\alpha_{\parallel}$. As discussed in Sec. V, these small biases can be attributed to the method used to add redshift errors in the mocks.

	$\Delta\alpha_{\parallel}$	$\Delta\alpha_{\perp}$	$\Delta\alpha_{\mathrm{iso}}$	$\Delta\alpha_{AP}$
CoLoRe-QL (300)	-0.0042 ± 0.0006	0.0005 ± 0.0008	-0.0020 ± 0.0004	0.0047 ± 0.0012
SACLAY (100)	-0.0074 ± 0.0010	0.0031 ± 0.0012	-0.0027 ± 0.0005	0.0106 ± 0.0019
Combined (400)	-0.0052 ± 0.0005	0.0012 ± 0.0006	-0.0023 ± 0.0003	0.0065 ± 0.0010

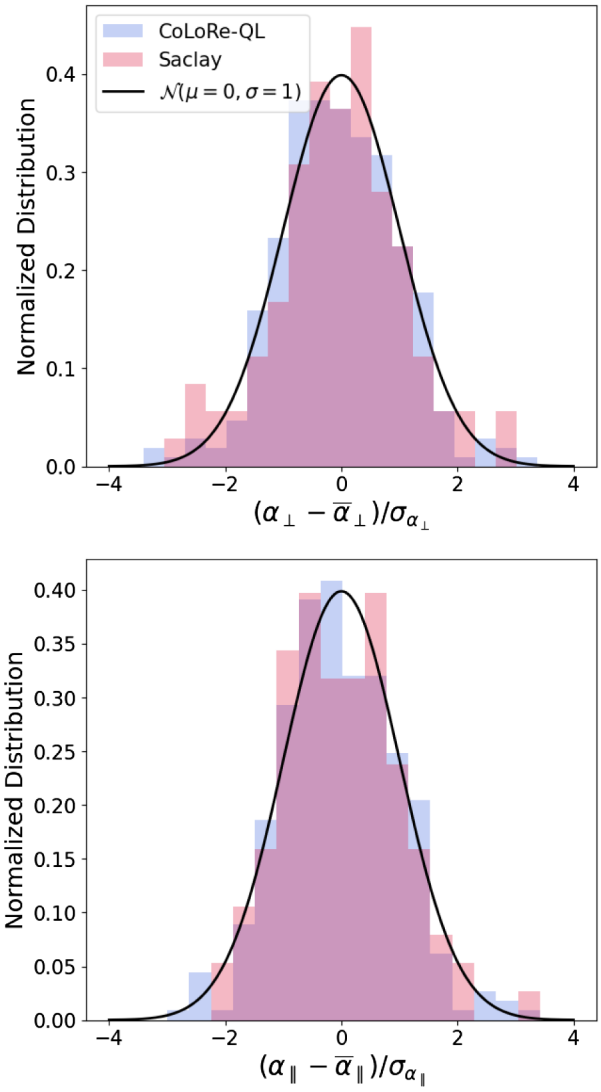

FIG. 8. Histograms of BAO residuals, defined as $\Delta\alpha_\parallel/\sigma_{\alpha_\parallel}$ and $\Delta\alpha_\perp/\sigma_{\alpha_\perp}$, where the uncertainties are the Gaussian ones reported by IMINUIT in each individual fit. The results for the 100 SACLAY mocks (red) and 300 CoLoRe-QL mocks (blue) are in good agreement with a normal distribution (black lines), indicating that the posteriors are Gaussian and the uncertainties are well estimated.

panel of Fig. 7 we display the scatter of the best-fit BAO parameters across the 400 mocks (300 CoLoRe-QL and 100 SACLAY). While there are no specific outliers, the measurements are not perfectly centered around the true values ($\alpha_\perp = \alpha_\parallel = 1$), particularly in the parallel direction, in agreement with the results from the stack of mocks (see Table. I).

The bottom panel of Fig. 7 shows the distribution of the best-fit χ^2 values obtained from the fits of individual mocks. The purple dashed line on this plot represents the χ^2 of the data best fit and the black curve represents the expected distribution based on the degrees of freedom in the fit. The plot shows that the distribution of χ^2 values

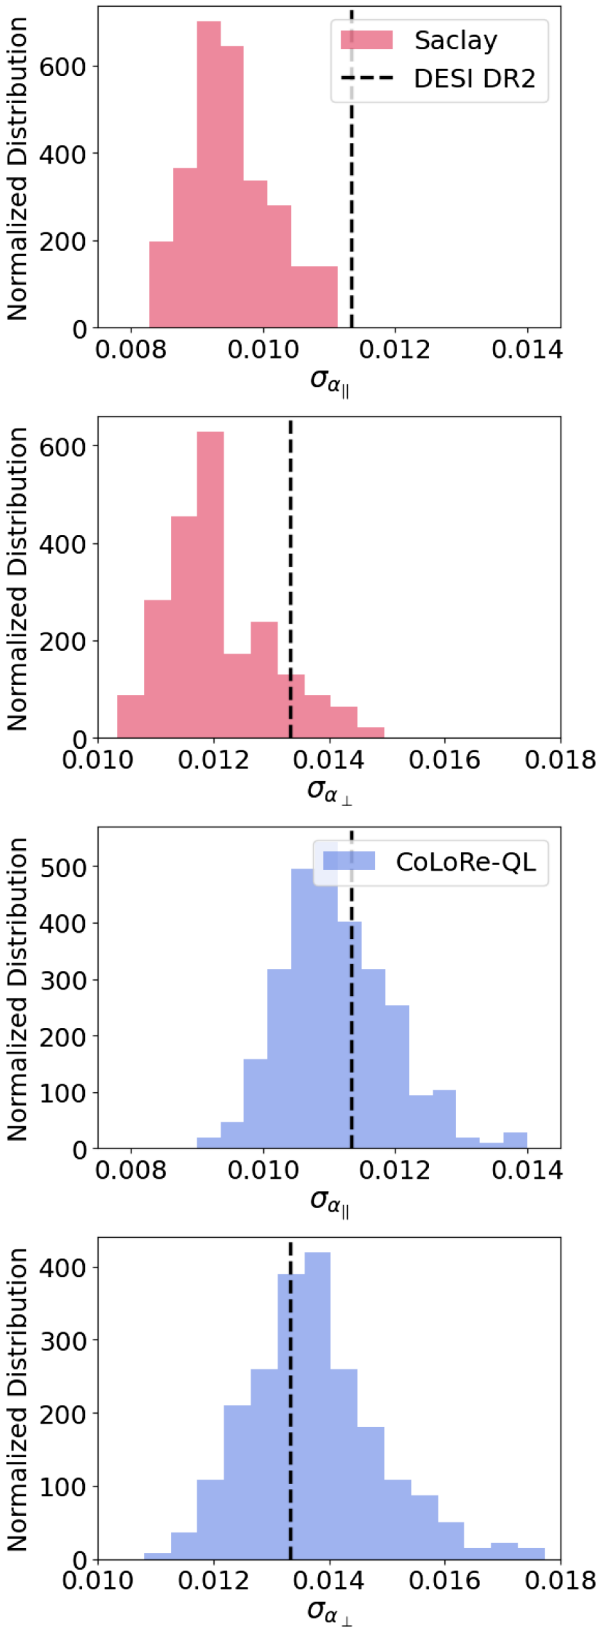

FIG. 9. Distribution of uncertainties on $\alpha_\parallel$ and $\alpha_\perp$ for 100 SACLAY mocks (top) and 300 CoLoRe-QL mocks (bottom). The black dashed line represents the uncertainties derived from the data, which align more closely with the CoLoRe-QL mocks, as the SACLAY mocks lack nonlinear broadening of the BAO peak.

from the individual mocks is higher than expected, although the situation is significantly better than in the mocks used in the validation of DESI DR1 (see Fig. 8 of [27]). In [27], Monte Carlo simulations starting from either the best-fit model or the stacked correlation function measurements demonstrate that this shift toward higher χ^2 values arises from the model's difficulty in fitting the correlation functions and not from the estimates of their covariance. We have tested that this disagreement remains even when limiting the analysis to larger scales ($r > 50h^{-1}$ Mpc), meaning that other effects that are also present on large scales must be responsible for the somewhat large values of χ^2. Examples of these could include the impact of continuum fitting distortions [73], the anisotropy caused by DLA wings even on large separations 61]], or the difficulties to model the nonstandard BAO broadening in the CoLoRe-QL mocks.

We validate the method to estimate the BAO uncertainties in Fig. 8, where we show the BAO residuals from the fit of each mock realization, defined as $\Delta\alpha_{\parallel}/\sigma_{\alpha_{\parallel}}$ and $\Delta\alpha_{\perp}/\sigma_{\alpha_{\perp}}$. The residual distributions closely follow a normal distribution, represented by the black curve, presenting evidence that the BAO uncertainties are accurate and approximately Gaussian.

Finally, in Fig. 9, the uncertainties on BAO parameters of the mocks are shown alongside the statistical uncertainties from DESI data (vertical dashed line). The top panel shows the results for SACLAY mocks in red, while the bottom panel shows CoLoRe-QL mocks in blue. As discussed in [27], nonlinear broadening of the BAO peak modestly yet noticeably degrades the BAO measurements. Consequently, mocks that do not include this effect, such as the SACLAY mocks, exhibit uncertainties that are smaller than those observed in the data. Specifically, the uncertainties in SACLAY mocks are 10%–20% smaller than those in the DESI results. In contrast, the CoLoRe-QL mocks demonstrate better agreement with the data, as they include a certain degree of nonlinear broadening, as discussed in Sec. II A 2.

V. DISCUSSION

In Sec. IV A we presented the BAO measurements from the combined correlations measured in 300 CoLoRe-QL and 100 SACLAY mocks. As shown in the top panel of Fig. 6 and reported in Table. I, we detect a small but significant bias from the expected value ($\alpha_{\perp} = \alpha_{\parallel} = 1$). In this section we demonstrate that this bias is caused by the method used to simulate redshift errors in the mocks (Sec. V A), and propose a mitigation strategy to avoid most of the contamination when analyzing real data (Sec. V B).

A. Analysis of mocks without redshift errors contamination

As discussed in Sec. II B, once we have already simulated the quasar spectra we add a random shift to the redshifts in our quasar catalogs, creating a mismatch between the reported redshift of the simulated quasar and the positions of the emission lines in their spectra. This allows us to capture the impact of redshift errors in the real catalogs, without the need to run the redshift fitter Redrock [40] for hundreds of mocks.

The first clear impact of random quasar redshift errors in our analysis is that they smear the measured cross-correlations along the line of sight [69], similar to the impact of nonlinear velocities (or fingers of God) in the clustering of quasars. As discussed in Sec. III D, we model this impact by multiplying the power spectrum of the cross-correlation model by $F_{\rm NL}^{\rm Q\alpha} = \exp\left[-(k_{\parallel}\sigma_v)^2/2\right]$.

Recently, [34] described a second, more subtle impact of redshift errors that cause spurious correlations not only in the cross-correlation with quasars, but also in the autocorrelation. As discussed in Sec. III B, in order to compute the fluctuations in the Lyα forest we need to estimate the mean continua of all our quasars, as a function of rest-frame wavelength [$\bar{C}(\lambda_{\rm RF})$]. However, the random shifts added to the redshifts of our mock catalogs are translated into misestimations of the rest-frame wavelengths, and this causes a smoothing of the features in the mean continuum, in particular of the emission lines present in the Lyα forest regions. This effect, coupled to the clustering of the background quasars, causes spurious correlations that could bias our results. These biases are discussed in more detail in [62].

In order to test this hypothesis, in Fig. 10 we present an alternative analysis (labeled "no σ_z") where the random redshift errors are only added *after* the continuum fitting, so that they only impact the cross-correlation measurements by smoothing them along the line of sight with $F_{\rm NL}^{\alpha\rm Q}$. For both sets of mocks this alternative analysis is significantly less biased than our baseline configuration, bringing it below the threshold of one-third of the DESI DR2 statistical uncertainty (dashed black line ellipses).

B. Mitigating the BAO scale bias by removing close pairs of quasars

Unfortunately, when analyzing real data we cannot avoid having redshift errors affecting our continuum fitting. For this reason, in this section we propose an alternative analysis that should minimize its impact without losing a significant fraction of the data.

In a recent study, [62] presented a model that can describe the spurious correlations introduced by redshift errors. In that paper, the authors also show that most of the contamination in the cross-correlation comes from pixel-quasar pairs where the background quasar (whose spectrum contains the pixel) and the foreground quasar (whose position we are considering) are very close to each other. Similarly, they show that most of the contamination in the autocorrelation comes from pairs where one of the pixels is very close to the background quasar of the other pixel.

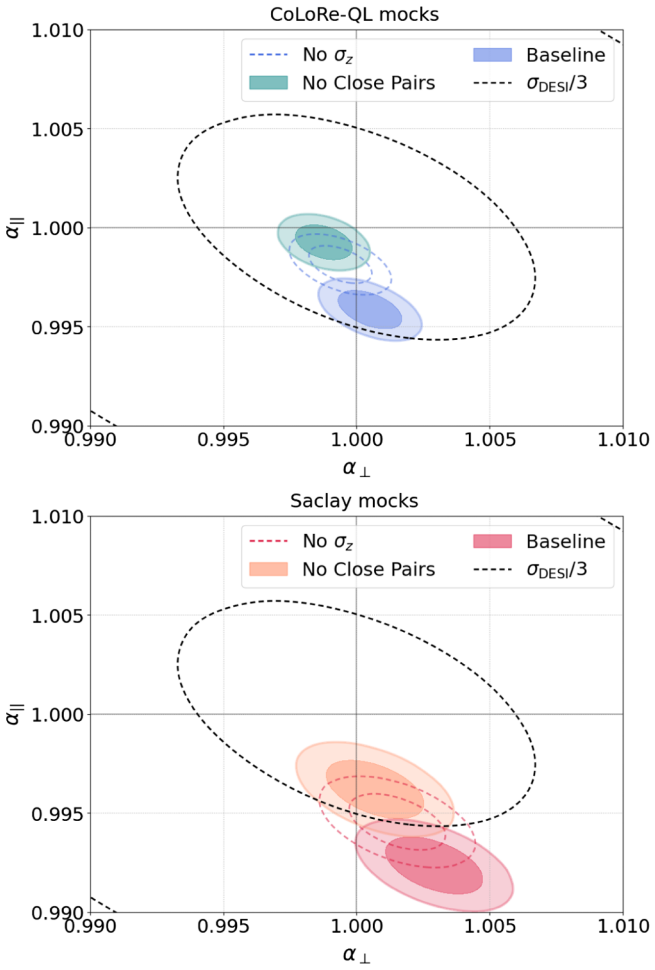

FIG. 10. Comparison of "No Close Pairs" and "no σ_z" analyses with the baseline. Both approaches aim to mitigate contamination from redshift errors prior to continuum fitting, with "No Close Pairs" excluding close quasar pairs (which are the most affected) and "no σ_z" removing estimation redshift errors entirely before continuum fitting.

Motivated by this, in Fig. 10 we also show an analysis (labeled "No close pairs") where we do not include these problematic pairs, in particular those with angular separations smaller than 20 arcmin and velocity separations smaller than 4000 km s^{-1}. This analysis is significantly less biased than our baseline configuration, and more importantly, it can be done when analyzing real data. In [26] we present the results of this alternative analysis on real data, and show that it has a negligible impact on our main BAO results. Also, to better understand the origin of the difference between the no close pairs and no σ_z analyses, we applied the same pair-removal cuts used in the no close pairs analysis to the no σ_z mocks. We found that the resulting BAO contours were essentially unchanged, indicating that these cuts do not affect the results furthermore than removing the redshift errors contamination. We therefore attribute the small differences between the analyses to statistical fluctuations.

While both alternative analyses show effectively unbiased results when analyzing the stack of 300 CoLoRe-QL mocks (Fig. 10, top panel), there seems to be a small bias when analyzing the stack of 100 SACLAY mocks, even though significantly smaller than in the baseline analysis and no longer larger than the tolerance limit of $\sigma_{\rm DESI}/3$. We have not been able to identify the cause for this residual bias in the analysis of the SACLAY mocks. However, we did test the SACLAY mocks in both their raw (see Appendix A) and uncontaminated forms, and found no evidence of bias in either case. This suggests that the bias arises only after adding contaminants. Given that metals and HCDs are included in both the SACLAY and CoLoRe-QL mocks using a similar prescription, it is likely that the observed bias stems from how the mocks were constructed, rather than from the analysis framework itself.

Finally, it is important to note that the current method used to add random errors to the mocks could have exaggerated the contamination. One of the main sources of redshift errors is that some of the quasar emission lines used to estimate their redshifts (like C III, Si IV, or C IV) can be affected by outflows or complex quasar physics. However, this should also affect the other high-ionization lines that are present in the Lyα region, meaning the forest would shift more uniformly. This would reduce the amount of smoothing of the mean quasar continuum, the level of the spurious correlations and therefore the bias on the BAO measurements.

VI. SUMMARY

In this work, we use synthetic data to validate the analysis of Lyα BAO measurements from the second data release of DESI, presented in [26]. DR2 includes spectra from nearly 1.2×10^6 quasars at redshift $z \geq 1.77$, nearly doubling the sample size of DESI DR1 [20]. As a result, DR2 achieves approximately a factor of 2 better statistical precision in Lyα forest BAO measurements compared to DR1, thereby necessitating a more rigorous validation of the cosmological inference pipeline.

The main differences between the validation of DR2 with respect to that of DR1 (presented in [27]) are the following. On the one hand, we have improved the mocks. We have increased the number of synthetic datasets used to validate the BAO analysis from 50 to 100 SACLAY mocks [31] and from 100 to 300 LyaCoLoRe mocks [30]. We have presented the CoLoRe-QL mocks, an improved version of the LyaCoLoRe mocks used in DR1 that include the nonlinear broadening of the BAO peak. On the other hand, we have improved the analysis to better mimic the analysis of real data. Instead of masking all the DLAs in the spectra (as done in the validation of DR1), we only mask DLAs in high signal-to-noise spectra (SNR > 2) and only for a randomly selected subset of 75% to emulate the completeness of the DLA catalog used in DR2 (see the discussion in Appendix B).

Using these mocks, we validated that the reported uncertainties on the BAO parameters are consistent with the scatter between the different realizations. Using the average measurement of correlations from all mocks, we have identified a small, but statistically significant bias in the line-of-sight BAO parameter. This bias is close to a third of the statistical uncertainty in DR2, the threshold that we had set to consider the analysis validated. We have shown that most of this bias is related to the redshift errors in the quasar catalog, a contamination first discussed in [34]. Following [62], we have presented an analysis that discards a small fraction of the data that is most contaminated, and we have shown that the residual bias is significantly smaller than our statistical uncertainty.

In the near future, we will present a cosmological analysis using the "full-shape" information contained in the Lyα correlations, not limited to the position of the BAO peak [74–76]. In order to validate these analyses, it will be important to count on improved Lyα mocks that simulate the nonlinear growth of structure, using perturbation theory models or similar. Nonlinearities affect BAO measurements in two ways: they broaden the peak, increasing the uncertainties, and they also induce a small shift in the peak position itself, as demonstrated in recent studies [77–79]. This is discussed further in [26], and obtaining nonlinear mocks will be a key goal for future work. It will also be important to improve the method to add redshift errors to the simulated catalogs. In order to do this, it will be useful to have a dedicated study of the relative redshifts of the different emission lines in the quasar spectra.

ACKNOWLEDGMENTS

The authors are honored to be permitted to conduct scientific research on Iolkam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation. This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science, and Technology of Mexico (CONAHCYT); the Ministry of Science, Innovation, and Universities of Spain (MICIU/AEI/10.13039/501100011033), and by the DESI Member Institutions [80]. L. C., J. C. M., A. F. R. and M. L. acknowledge support from the European Union's Horizon Europe research and innovation program (COSMO-LYA, Grant Agreement No. 101044612). A. F. R. acknowledges financial support from the Spanish Ministry of Science and Innovation under the Ramon y Cajal program (RYC-2018-025210) and the PGC2021-123012NB-C41 project. I. F. A. E. is partially funded by the CERCA program of the Generalitat de Catalunya. D. A. acknowledges support from the Beecroft Trust. A. C. and C. G. Q. acknowledge support from a NASA fellowship awarded through the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. National Science Foundation, the U.S. Department of Energy, or any of the listed funding agencies.

DATA AVAILABILITY

The data that support the findings of this article are openly available [81,82].

APPENDIX A: ANALYSES ON RAW MOCKS

We refer to a BAO analysis procedure applied on the raw mocks, described in Sec. II A, without astrophysical contaminants, instrumental noise, or quasar continuum templates added as the "raw analysis." This analysis allows us to verify that we are able to recover the position of the BAO peak specified by the input cosmology and discard an intrinsic bias on the raw mocks.

In this analysis, we do not have to perform the astrophysical contaminants masking or the continuum fitting procedures described in Secs. III A, and III B, respectively, as they are not present in raw mocks. In other words, we compute the δ field from the transmitted flux fraction boxes directly by

$$\delta_q(\lambda) = \frac{F(\lambda)}{\bar{F}(z)} - 1. \tag{A1}$$

The correlation functions are computed in the same way as in the standard analysis, described in Sec. III C. However, in this case we do not smooth the covariance matrix since it is already positive definite given the high SNR of the raw analysis. Moreover, smoothing this covariance produces a nonpositive definite matrix which induces a bias on the cross-correlation measurement on both SACLAY and CoLoRe-QL mocks.

In the case of the model, described in Sec. III D, we do not apply a distortion matrix or exclude astrophysical contaminants as these effects do not included in the raw analysis. Furthermore, we perform additional fits for the auto- and cross-correlations individually. We free the BAO

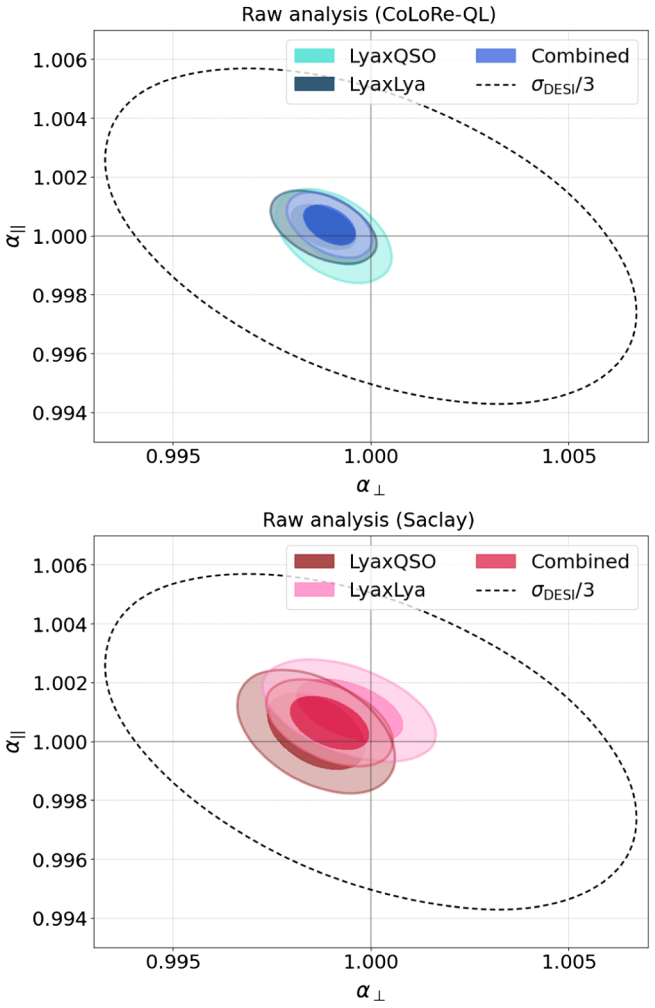

FIG. 11. Top: fit results for the autocorrelation alone, cross-correlation alone, and combined fit for CoLoRe-QL raw mocks. In the raw analysis there is no noise, no contaminants, and no continuum fitting. Bottom: same fits for SACLAY raw mocks. All fits are compatible and unbiased.

parameters ($\alpha_\parallel$ and $\alpha_\perp$), the smoothing parameters ($\sigma_\parallel$ and $\sigma_\perp$), and Lyα linear bias and RSD parameters ($b_{\mathrm{Ly}\alpha}$ and $\beta_{\mathrm{Ly}\alpha}$). For the cross-correlation fit, we additionally fix the quasar bias (b_q) value to each type of mocks truth value ($b_q = 3.123$ for CoLoRe-QL and $b_q = 3.3$ for SACLAY).

Fig. 11 shows the results for CoLoRe-QL mocks (top panel) and SACLAY mocks (bottom panel). The combined fit results for $\alpha_{||}, \alpha_\perp$ are

$$\alpha_{||}^{\mathrm{QL}} = 1.0004 \pm 0.0005, \qquad \alpha_\perp^{\mathrm{QL}} = 0.9990 \pm 0.0004, \tag{A2}$$

and

$$\alpha_{||}^{\mathrm{Saclay}} = 1.0009 \pm 0.0006, \qquad \alpha_\perp^{\mathrm{Saclay}} = 0.9992 \pm 0.0007, \tag{A3}$$

for CoLoRe-QL and SACLAY mocks, respectively. All the results presented in Fig. 11 are within a third of the statistical uncertainty from DESI DR2 threshold and close to the $\alpha_{||} = \alpha_\perp = 1$ target value, confirming that the raw mocks are suitable to be used in the analysis validation presented in this work.

APPENDIX B: IMPACT OF DLA MASKING

Damped Lyman-α absorbers are found and masked in the BAO analysis to optimize the statistical precision of the measurement, and any residual DLA contamination is accounted for via free parameters in the model fit. In this section, we use mocks to study the impact on BAO results when changing the DLA-finding algorithm and to find the optimal DLA catalog configuration for the DR2 BAO analysis on real data.

The Lyα BAO analysis from DESI DR1 [20] masked those DLAs that were found by both a Gaussian process (GP) finder [83] and a convolutional neural network (CNN) finder [84]. Tests on one mock found the combination to produce unbiased BAO results compared to a catalog of all input DLAs [27]. In preparation for DR2 analysis, several improvements have been made to the GP and CNN finders, and an additional finder (referred to as the template, or TMP, finder) has been developed that fits each spectrum with a template for quasars and one or more DLA models [44].

To test the performance of these finders for the DR2 BAO analysis, we ran all three finders on one realization of the LyaCoLoRe mocks.[18] After each DLA finder was run on the mock data, several quality cuts were made to the raw output catalogs. In a first series of analyses, all three finders were reduced to only include DLAs found in spectra for which the mean signal-to-noise ratio measured on the red side of the Lyα forest was SNR > 3. Furthermore, only DLAs with high column density, $\log(\frac{N_{\mathrm{HI}}}{\mathrm{cm}^{-2}}) > 20.3$ for the N_{HI} calculated by each algorithm were kept. We included additional quality cuts, for the CNN confidence parameter CONF > 0.5; for the GP probability parameter P_DLA > 0.9; and for the TMP enforcement of DLAFLAG $= 0$ (see [44,83,84] for further details on these parameters). The performance of the different algorithms, in terms of purity and completeness of their DLA catalogs, is discussed in the companion paper [44], and we refer the interested reader to that publication for more detail. Here we complement this study by looking at the impact on the BAO parameters when using different DLA finders.

We ran different end-to-end BAO analyses on this particular LyaCoLoRe mock, but using different catalogs to mask DLAs, including an analysis without masking any DLA (labeled "no masking") and an analysis that masked

[18]This mock was generated before the improvements described in Sec. II A 2 and does not have nonlinear broadening of the BAO.

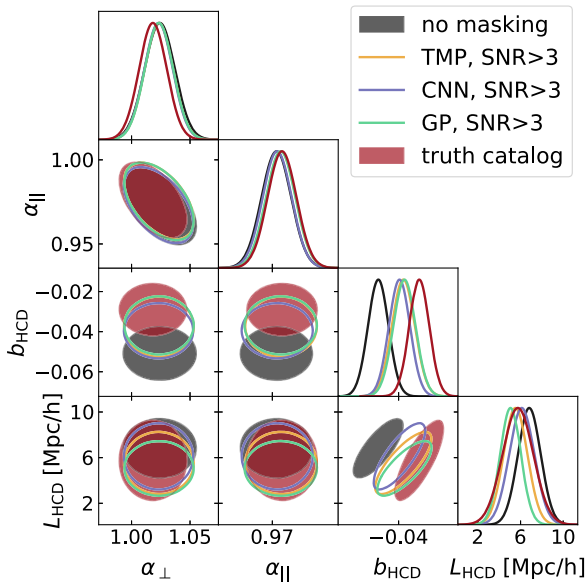

FIG. 12. Contours for HCD parameters and BAO parameters under different DLA masking schemes. Only 95% confidence levels are shown. The results from the three DLA finders are shown in open contours to facilitate visualization, while the two extremes (no masking and truth catalog) are shaded. For $b_{\rm HCD}$, the finder-based analyses produce values lying between the two extremes as expected, and the $L_{\rm HCD}$ varies depending on the typical length scale masked. However, the BAO parameters (upper left) are stable to variations. The BAO parameters are in mild tension with the true input cosmology values because this is an analysis on only a single mock; the stack of mocks is unbiased as discussed in Sec. IV A.

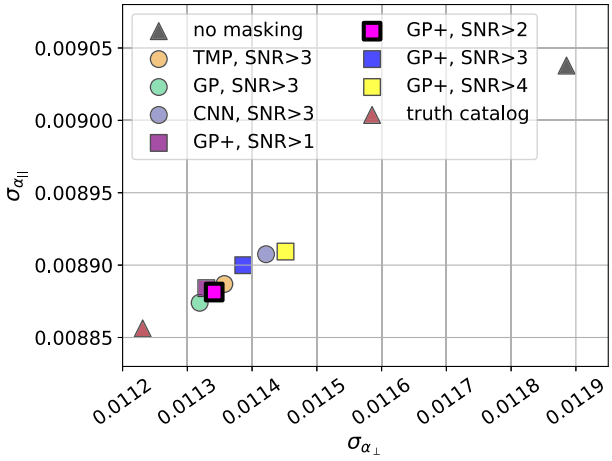

FIG. 13. Forecasted uncertainties on $\alpha_\parallel$ and $\alpha_\perp$ for the one-finder (circles) and two-finder (squares) DLA catalogs compared to the case of no masking and masking all objects in the truth catalog (triangles). The final choice for data analysis is the GP+, SNR > 2 catalog (pink square with bold borders), which has an optimal completeness-purity tradeoff in addition to performing with high precision, as shown here.

all the DLAs added to the mocks (labeled "truth catalog"). Fig. 12 shows that the BAO parameters are very robust to the DLA catalog used, and the small differences can be explained by statistical fluctuations. However, each analysis has a different amount of residual contamination from DLAs that were not identified, and this translates into different values for the nuisance parameters describing the contamination by high column density systems: $L_{\rm HCD}$, $\beta_{\rm HCD}$, and $b_{\rm HCD}$, presented in Sec. III D. As expected, $b_{\rm HCD}$ varies between a large negative value for the no masking case, where a large contribution to the bias due to DLAs must be taken into account in the model, to a small negative value when all DLAs are removed and only small HCD contributions remains (truth catalog masking). The three DLA finders result in very similar values for $b_{\rm HCD}$ and all lie between the two extremes, having found most but not all DLAs in the fiducial catalog. $L_{\rm HCD}$, the typical length scale of DLA systems that appear in the data, is large when none are masked (because even large DLAs remain in the spectra) and decreases with any form of masking, because only smaller systems remain. The model succeeds in capturing the correlation function variations through the HCD parameters such that the choice of DLA masking does not bias the BAO parameters, as evidenced by the stability in $\alpha_\parallel$ and $\alpha_\perp$ parameters.

In order to decide on the optimal method to mask DLAs in the DR2 analysis, [44] studied possible ways to combine the catalogs constructed by the different DLA finders. Combined catalogs containing DLAs detected by at least two finders, when limited to spectra with SNR > 2, resulted in similar performances in terms of completeness and purity of the catalogs, both around 75%. The recommendation from [44], that was adopted in the Lyα BAO measurement from DR2 of [26], was to mask DLAs that were identified by the GP finder and one of the other two finders,[19] a combination that we label as "GP+".

After the decision was made to use the GP+ catalog, we revisited the impact of different SNR cuts. Masking DLAs detected in lower SNR spectra would trade a lower purity of the sample for a higher completeness. In Fig. 13 we compare the forecasted BAO uncertainties for different DLA masking strategies, in particular when varying the SNR threshold in the GP+ catalog. However, as the different finders mask different Lyα data, there is some natural sample variation in both the signal and the uncertainty of the resulting BAO parameters. This makes it challenging to directly compare the BAO uncertainties to make the optimal choice. To more robustly quantify the performance of each masking catalog, we make a BAO forecast using VEGA. The forecast works by creating a vector that represents the correlation function but is purely equal to the (noiseless) best-fit model for that analysis.

[19]We use the values of redshift and column density reported by the GP finder.

Using the covariance matrix of the mock data in combination with that vector, the BAO fit is performed, obtaining BAO uncertainties that are not affected by statistical fluctuations.

Figure 13 shows that the error on $\alpha_\parallel$ and $\alpha_\perp$ is highest when no DLAs are masked and moves lower for the various finders, approaching the truth catalog case with smallest error bars. Although GP on its own results in the smallest errors of the single-finder cases (circles), as previously stated, we deemed it more robust to require detection in at least one of the other finders. Of the SNR cuts tested for the GP+ catalog, the SNR > 2 and SNR > 1 versions perform best, so we choose SNR > 2 for its higher purity for the BAO DR2 data analysis.

Given the trend to lower $\sigma_{\alpha_\perp}$ and $\sigma_{\alpha_\parallel}$ for increasing numbers of masked DLAs, it is interesting to question whether masking increasingly smaller DLAs would continue to reduce the error bars. Given that masking also removes the Lyα signal, however, there must be a point at which the gains in precision due to removal of these contaminants are outweighed by the loss in signal. We explored this idea by masking not only all the DLAs added to the mocks, but also all the HCD systems with increasingly lower column density. By running VEGA forecasts, we find that the uncertainties continue to decrease until a threshold of $\sim\log(N_{\rm HI}) > 19$. However, beyond this, continuing to mask systems with even lower column density causes the uncertainties to increase again due to excess loss in the Lyα signal. This investigation is not relevant to the DESI DR2 analysis, as the DLA finders are not expected to perform well at such low column densities; however, it demonstrates that with future improved algorithms, the analysis could benefit from masking lower column density systems.

APPENDIX C: DETAILS ON THE NUISANCE PARAMETERS

As discussed in Secs. III E and IV, we perform various types of analyses on the mocks presented throughout this work. Table II provides a description of the free parameters used to model the correlation functions from our mocks in the range $30 < r < 180h^{-1}$ Mpc, with the corresponding priors listed in the second column.

The last two rows of this table specify the priors used for the finite-grid Gaussian smoothing parameters $(\sigma_\parallel, \sigma_\perp)$ and the nonlinear BAO peak broadening parameters $(\Sigma_\parallel, \Sigma_\perp)$. These priors were applied exclusively in a fit on the high-precision BAO measurements described in Sec. IV A over $10 < r < 180h^{-1}$ Mpc to determine these parameter values. For all other analyses, these parameters remain fixed at $\sigma_\parallel = 2.0$, $\sigma_\perp = 1.7$, $\Sigma_\parallel = 4.3$, and $\Sigma_\perp = 3.1h^{-1}$ Mpc for CoLoRe-QL mocks, while for SACLAY mocks, they are fixed at $\sigma_\parallel = \sigma_\perp = 2.2h^{-1}$ Mpc and $\Sigma_\parallel = \Sigma_\perp = 0$.

For the raw analyses presented in Appendix A, we only allow the BAO scale parameters $(\alpha_\parallel, \alpha_\perp)$, the Ly$\alpha$ linear bias and RSD parameter (b_α, β_α), the quasar linear bias $(b_{\rm q})$, the finite-grid Gaussian smoothing parameters $(\sigma_\parallel, \sigma_\perp)$, and the nonlinear BAO peak broadening parameters $(\Sigma_\parallel, \Sigma_\perp)$ to vary. All other parameters are excluded from the model in this specific type of analyses.

TABLE II. Free parameters and priors used in the $30 < r < 180h^{-1}$ Mpc model fitting presented in this work. $U(\min, \max)$ denotes flat priors in the [min, max] range. $\mathcal{N}(\mu, \sigma)$ denotes Gaussian priors of mean μ and σ standard deviation. The priors set to the parameters on the last two rows are only used in a $10 < r < 180h^{-1}$ Mpc fit of the correlation functions to determine these parameter values and remain fixed for all other analyses.

Parameter	Prior	Description
$\alpha_\parallel, \alpha_\perp$	$U(0.01, 2.0)$	BAO scale position parameters
$b_{Ly\alpha}$	$U(-2.0, 0.0)$	Lyman-α linear bias
$\beta_{Ly\alpha}$	$U(0.0, 5.0)$	Lyman-α RSD parameter
$b_{\rm q}$	$U(0.0, 10.0)$	Quasar linear bias
$\sigma_v [h^{-1}$ Mpc$]$	$U(0.0, 15.0)$	Statistical QSO redshift estimate errors amplitude
$\Delta r_{\rm shift} [h^{-1}$ Mpc$]$	$\mathcal{N}(0.0, 1.0)$	Systematic QSO redshift estimate shift
$b_{\rm HCD}$	$U(-0.2, 0.0)$	High column density systems linear bias
$\beta_{\rm HCD}$	$\mathcal{N}(0.5, 0.09)$	HCD RSD parameter
$L_{\rm HCD}$	$\mathcal{N}(5.0, 1.0)$	Typical length scale of unmasked HCDs
$b_{\rm Si\,II(1190)}$	$U(-0.5, 0.5)$	Linear bias of the Si II (1190) transition
$b_{\rm Si\,II(1193)}$	$U(-0.5, 0.5)$	Linear bias of the Si II (1193) transition
$b_{\rm Si\,III(1207)}$	$U(-0.5, 0.5)$	Linear bias of the Si III (1207) transition
$b_{\rm Si\,II(1260)}$	$U(-0.5, 0.5)$	Linear bias of the Si II (1260) transition
$\sigma_\parallel, \sigma_\perp [h^{-1}$ Mpc$]$	$U(0.0, 10.0)$	Finite-grid Gaussian smoothing parameters
$\Sigma_\parallel, \Sigma_\perp [h^{-1}$ Mpc$]$	$U(0.0, 20.0)$	Nonlinear broadening of the BAO peak parameters

[1] S. Cole, W. J. Percival, J. A. Peacock, P. Norberg, C. M. Baugh, C. S. Frenk, I. Baldry, J. Bland-Hawthorn, T. Bridges, R. Cannon *et al.*, Mon. Not. R. Astron. Soc. **362**, 505 (2005).

[2] D. J. Eisenstein, I. Zehavi, D. W. Hogg, R. Scoccimarro, M. R. Blanton, R. C. Nichol, R. Scranton, H.-J. Seo, M. Tegmark, Z. Zheng *et al.*, Astrophys. J. **633**, 560 (2005).

[3] K. S. Dawson, D. J. Schlegel, C. P. Ahn, S. F. Anderson, É. Aubourg, S. Bailey, R. H. Barkhouser, J. E. Bautista, A. Beifiori, A. A. Berlind*et al.*, Astron. J. **145**, 10 (2013).

[4] N. G. Busca, T. Delubac, J. Rich, S. Bailey, A. Font-Ribera, D. Kirkby, J.-M. Le Goff, M. M. Pieri, A. Slosar, É. Aubourg *et al.*, Astron. Astrophys. **552**, A96 (2013).

[5] A. Slosar, V. Iršič, D. Kirkby, S. Bailey, N. G. Busca, T. Delubac, J. Rich, É. Aubourg, J. E. Bautista, V. Bhardwaj *et al.*, J. Cosmol. Astropart. Phys. 04 (2013) 026.

[6] D. Kirkby, D. Margala, A. Slosar, S. Bailey, N. G. Busca, T. Delubac, J. Rich, J. E. Bautista, M. Blomqvist, J. R. Brownstein *et al.*, J. Cosmol. Astropart. Phys. 03 (2013) 024.

[7] A. Font-Ribera, D. Kirkby, N. Busca, J. Miralda-Escudé, N. P. Ross, A. Slosar, J. Rich, É. Aubourg, S. Bailey, V. Bhardwaj *et al.*, J. Cosmol. Astropart. Phys. 05 (2014) 027.

[8] M. Levi, C. Bebek, T. Beers, R. Blum, R. Cahn, D. Eisenstein, B. Flaugher, K. Honscheid, R. Kron, O. Lahav *et al.* (Representing the DESI Collaboration), arXiv:1308.0847.

[9] A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland, O. Ballester *et al.* (DESI Collaboration), arXiv:1611.00036.

[10] B. Abareshi, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, R. Alfarsy, L. Allen, C. Allende Prieto, O. Alves, J. Ameel *et al.* (DESI Collaboration), Astron. J. **164**, 207 (2022).

[11] A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland, O. Ballester *et al.* (DESI Collaboration), arXiv:1611.00037.

[12] J. H. Silber, P. Fagrelius, K. Fanning, M. Schubnell, J. N. Aguilar, S. Ahlen, J. Ameel, O. Ballester, C. Baltay, C. Bebek *et al.*, Astron. J. **165**, 9 (2023).

[13] T. N. Miller, P. Doel, G. Gutierrez, R. Besuner, D. Brooks, G. Gallo, H. Heetderks, P. Jelinsky, S. M. Kent, M. Lampton *et al.* , Astron. J. **168**, 95 (2024).

[14] C. Poppett, L. Tyas, J. Aguilar, C. Bebek, D. Bramall, T. Claybaugh, J. Edelstein, P. Fagrelius, H. Heetderks, P. Jelinsky *et al.*, Astron. J. **168**, 245 (2024).

[15] E. F. Schlafly, D. Kirkby, D. J. Schlegel, A. D. Myers, A. Raichoor, K. Dawson, J. Aguilar, C. Allende Prieto, S. Bailey, S. BenZvi *et al.*, Astron. J. **166**, 259 (2023).

[16] J. Guy, S. Bailey, A. Kremin, S. Alam, D. M. Alexander, C. Allende Prieto, S. BenZvi, A. S. Bolton, D. Brooks, E. Chaussidon *et al.*, Astron. J. **165**, 144 (2023).

[17] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, G. Aldering, D. M. Alexander, R. Alfarsy, C. Allende Prieto, M. Alvarez, O. Alves *et al.* (DESI Collaboration), Astron. J. **167**, 62 (2024).

[18] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, G. Aldering, D. M. Alexander, R. Alfarsy, C. Allende Prieto, M. Alvarez, O. Alves *et al.* (DESI Collaboration), Astron. J. **168**, 58 (2024).

[19] M. A. Karim, A. G. Adame, D. Aguado, J. Aguilar, S. Ahlen, S. Alam, G. Aldering, D. M. Alexander, R. Alfarsy, L. Allen *et al.* (DESI Collaboration), arXiv:2503.14745.

[20] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.* (DESI Collaboration), J. Cosmol. Astropart. Phys. 01 (2025) 124.

[21] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.* (DESI Collaboration), J. Cosmol. Astropart. Phys. 04 (2025) 012.

[22] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.* (DESI Collaboration), J. Cosmol. Astropart. Phys. 09 (2025) 008.

[23] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.* (DESI Collaboration), J. Cosmol. Astropart. Phys. 07 (2025) 017.

[24] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.* (DESI Collaboration), J. Cosmol. Astropart. Phys. 02 (2025) 021.

[25] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, C. Allende Prieto, M. Alvarez, O. Alves, A. Anand, U. Andrade *et al.* (DESI Collaboration), J. Cosmol. Astropart. Phys. 07 (2025) 028.

[26] M. A. Karim, J. Aguilar, S. Ahlen, C. Allende Prieto, O. Alves, A. Anand, U. Andrade, E. Armengaud, A. Aviles, S. Bailey *et al.* (DESI Collaboration), Phys. Rev. D **112**, 083514 (2025).

[27] A. Cuceu, H. K. Herrera-Alcantar, C. Gordon, P. Martini, J. Guy, A. Font-Ribera, A. X. Gonzalez-Morales, M. Abdul Karim, J. Aguilar, S. Ahlen *et al.*, J. Cosmol. Astropart. Phys. 01 (2025) 148.

[28] J. J. Givans, A. Font-Ribera, A. Slosar, L. Seeyave, C. Pedersen, K. K. Rogers, M. Garny, D. Blas, and V. Iršič, J. Cosmol. Astropart. Phys. 09 (2022) 070.

[29] M. A. Karim, J. Aguilar, S. Ahlen, S. Alam, L. Allen, C. Allende Prieto, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.* (DESI Collaboration), Phys. Rev. D **112**, 083515 (2025).

[30] J. Farr, A. Font-Ribera, H. du Mas des Bourboux, A. Muñoz-Gutiérrez, F. J. Sánchez, A. Pontzen, A. Xochitl González-Morales, D. Alonso, D. Brooks, P. Doel *et al.*, J. Cosmol. Astropart. Phys. 03 (2020) 068.

[31] T. Etourneau, J.-M. Le Goff, J. Rich, T. Tan, A. Cuceu, S. Ahlen, E. Armengaud, D. Brooks, T. Claybaugh, A. de la Macorra *et al.*, J. Cosmol. Astropart. Phys. 05 (2024) 077.

[32] C. Ramírez-Pérez, J. Sanchez, D. Alonso, and A. Font-Ribera, J. Cosmol. Astropart. Phys. 05 (2022) 002.

[33] R. A. C. Croft, D. H. Weinberg, N. Katz, and L. Hernquist, Astrophys. J. **495**, 44 (1998).

[34] S. Youles, J. E. Bautista, A. Font-Ribera, D. Bacon, J. Rich, D. Brooks, T. M. Davis, K. Dawson, A. de la Macorra, G. Dhungana *et al.*, Mon. Not. R. Astron. Soc. **516**, 421 (2022).

[35] S. Yuan, H. Zhang, A. J. Ross, J. Donald-McCann, B. Hadzhiyska, R. H. Wechsler, Z. Zheng, S. Alam, V. Gonzalez-Perez, J. N. Aguilar *et al.*, Mon. Not. R. Astron. Soc. **530**, 947 (2024).

[36] P. Laurent *et al.* (eBOSS Collaboration), J. Cosmol. Astropart. Phys. 07 (2017) 017.
[37] H. K. Herrera-Alcantar, A. Muñoz-Gutiérrez, T. Tan, A. X. González-Morales, A. Font-Ribera, J. Guy, J. Moustakas, D. Kirkby, E. Armengaud, A. Bault *et al.*, J. Cosmol. Astropart. Phys. 01 (2025) 141.
[38] I. McGreer, J. Moustakas, and J. Schindler, simqso: Simulated quasar spectra generator, Astrophysics Source Code Library, record ascl:2106.008 (2021).
[39] D. Kirkby, S. Bailey, J. Guy, and B. A. Weaver, Quick simulations of fiber spectrograph response v0.5 (2016).
[40] A. Anand *et al.*, Astron. J. **168**, 124 (2024).
[41] A. Brodzeller, K. Dawson, S. Bailey, J. Yu, A. J. Ross, A. Bault, S. Filbert, J. Aguilar, S. Ahlen, D. M. Alexander *et al.*, Astron. J. **166**, 66 (2023).
[42] N. Busca and C. Balland, arXiv:1808.09955.
[43] A. Bault, D. Kirkby, J. Guy, A. Brodzeller, J. Aguilar, S. Ahlen, S. Bailey, D. Brooks, L. Cabayol-Garcia, J. Chaves-Montero *et al.*, J. Cosmol. Astropart. Phys. 01 (2025) 130.
[44] A. Brodzeller, M. Wolfson, D. M. Santos, M. Ho, T. Tan, M. M. Pieri, A. Cuceu, M. A. Karim, J. Aguilar, S. Ahlen *et al.*, Phys. Rev. D **112**, 083510 (2025).
[45] P. Noterdaeme, P. Petitjean, W. C. Carithers, I. Pâris, A. Font-Ribera, S. Bailey, E. Aubourg, D. Bizyaev, G. Ebelke, H. Finley *et al.*, Astron. Astrophys. **547**, L1 (2012).
[46] L. Ennesser, P. Martini, A. Font-Ribera, and I. Pérez-Ràfols, Mon. Not. R. Astron. Soc. **511**, 3514 (2022).
[47] H. du Mas des Bourboux, J. Rich, A. Font-Ribera, V. de Sainte Agathe, J. Farr, T. Etourneau, J.-M. Le Goff, A. Cuceu, C. Balland, J. E. Bautista *et al.*, Astrophys. J. **901**, 153 (2020).
[48] C. Ramírez-Pérez, I. Pérez-Ràfols, A. Font-Ribera, M. A. Karim, E. Armengaud, J. Bautista, S. F. Beltran, L. Cabayol-Garcia, Z. Cai, S. Chabanier *et al.*, Mon. Not. R. Astron. Soc. **528**, 6666 (2024).
[49] A. Slosar, A. Font-Ribera, M. M. Pieri, J. Rich, J.-M. Le Goff, É. Aubourg, J. Brinkmann, N. Busca, B. Carithers, R. Charlassier *et al.*, J. Cosmol. Astropart. Phys. 09 (2011) 001.
[50] A. Font-Ribera, J. Miralda-Escudé, E. Arnau, B. Carithers, K.-G. Lee, P. Noterdaeme, I. Pâris, P. Petitjean, J. Rich, E. Rollinde *et al.*, J. Cosmol. Astropart. Phys. 11 (2012) 059.
[51] J. E. Bautista, N. G. Busca, J. Guy, J. Rich, M. Blomqvist, H. du Mas des Bourboux, M. M. Pieri, A. Font-Ribera, S. Bailey, T. Delubac *et al.*, Astron. Astrophys. **603**, A12 (2017).
[52] C. Gordon, A. Cuceu, J. Chaves-Montero, A. Font-Ribera, A. X. González-Morales, J. Aguilar, S. Ahlen, E. Armengaud, S. Bailey, A. Bault *et al.*, J. Cosmol. Astropart. Phys. 11 (2023) 045.
[53] F. Beutler, E. Castorina, and P. Zhang, J. Cosmol. Astropart. Phys. 03 (2018) 040.
[54] P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo *et al.* (Planck Collaboration), Astron. Astrophys. **594**, A13 (2016).
[55] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak *et al.* (Planck Collaboration), Astron. Astrophys. **641**, A6 (2020).
[56] H. du Mas des Bourboux, J.-M. Le Goff, M. Blomqvist, N. G. Busca, J. Guy, J. Rich, C. Yèche, J. E. Bautista, É. Burtin, K. S. Dawson *et al.*, Astron. Astrophys. **608**, A130 (2017).
[57] P. McDonald, U. Seljak, S. Burles, D. J. Schlegel, D. H. Weinberg, R. Cen, D. Shih, J. Schaye, D. P. Schneider, N. A. Bahcall *et al.*, Astrophys. J. Suppl. Ser. **163**, 80 (2006).
[58] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, Astrophys. J. **622**, 759 (2005).
[59] A. Font-Ribera and J. Miralda-Escudé, J. Cosmol. Astropart. Phys. 07 (2012) 028.
[60] K. K. Rogers, S. Bird, H. V. Peiris, A. Pontzen, A. Font-Ribera, and B. Leistedt, Mon. Not. R. Astron. Soc. **476**, 3716 (2018).
[61] T. Tan, J. Rich, E. Chaussidon, J. M. Le Goff, C. Balland, E. Armengaud, J. Aguilar, S. Ahlen, D. Bianchi, D. Brooks *et al.*, J. Cosmol. Astropart. Phys. 11 (2025) 074.
[62] C. Gordon, A. Cuceu, A. Font-Ribera, H. K. Herrera-Alcantar, J. N. Aguilar, S. Ahlen, D. Bianchi, D. Brooks, T. Claybaugh, S. Cole *et al.*, arXiv:2505.08789.
[63] T. Delubac, J. E. Bautista, N. G. Busca, J. Rich, D. Kirkby, S. Bailey, A. Font-Ribera, A. Slosar, K.-G. Lee, M. M. Pieri *et al.*, Astron. Astrophys. **574**, A59 (2015).
[64] V. de Sainte Agathe, C. Balland, H. du Mas des Bourboux, N. G. Busca, M. Blomqvist, J. Guy, J. Rich, A. Font-Ribera, M. M. Pieri, J. E. Bautista *et al.*, Astron. Astrophys. **629**, A85 (2019).
[65] M. Blomqvist, H. du Mas des Bourboux, N. G. Busca, V. de Sainte Agathe, J. Rich, C. Balland, J. E. Bautista, K. Dawson, A. Font-Ribera, J. Guy *et al.*, Astron. Astrophys. **629**, A86 (2019).
[66] A. Lewis, J. Cosmol. Astropart. Phys. 08 (2025) 025.
[67] N. Kaiser, Mon. Not. R. Astron. Soc. **227**, 1 (1987).
[68] A. Arinyo-i-Prats, J. Miralda-Escudé, M. Viel, and R. Cen, J. Cosmol. Astropart. Phys. 12 (2015) 017.
[69] A. Font-Ribera, E. Arnau, J. Miralda-Escudé, E. Rollinde, J. Brinkmann, J. R. Brownstein, K.-G. Lee, A. D. Myers, N. Palanque-Delabrouille, I. Pâris *et al.*, J. Cosmol. Astropart. Phys. 05 (2013) 018.
[70] A. J. S. Hamilton, Mon. Not. R. Astron. Soc. **312**, 257 (2000).
[71] W. J. Handley, M. P. Hobson, and A. N. Lasenby, Mon. Not. R. Astron. Soc. **450**, L61 (2015).
[72] A. Cuceu, A. Font-Ribera, and B. Joachimi, J. Cosmol. Astropart. Phys. 07 (2020) 035.
[73] N. Busca, J. Rich, J. Bautista, A. Cuceu, A. Font-Ribera, J. Guy, H. K. Herrera-Alcantar, J. Stermer, C. Balland, J. Aguilar *et al.*, J. Cosmol. Astropart. Phys. 09 (2025) 020.
[74] A. Cuceu, A. Font-Ribera, B. Joachimi, and S. Nadathur, Mon. Not. R. Astron. Soc. **506**, 5439 (2021).
[75] A. Cuceu, A. Font-Ribera, S. Nadathur, B. Joachimi, and P. Martini, Phys. Rev. Lett. **130**, 191003 (2023).
[76] A. Cuceu, A. Font-Ribera, P. Martini, B. Joachimi, S. Nadathur, J. Rich, A. X. González-Morales, H. du Mas des Bourboux, and J. Farr, Mon. Not. R. Astron. Soc. **523**, 3773 (2023).
[77] R. de Belsunce, S.-F. Chen, M. M. Ivanov, C. Ravoux, S. Chabanier, J. Sexton, and Z. Lukic, Phys. Rev. D **111**, 063524 (2025).

[78] F. Sinigaglia, F.-S. Kitaura, K. Nagamine, and Y. Oku, Astrophys. J. Lett. **971**, L22 (2024).
[79] B. Hadzhiyska, R. de Belsunce, A. Cuceu, J. Guy, M. M. Ivanov, H. Coquinot, and A. Font-Ribera, arXiv:2503.13442.
[80] https://www.desi.lbl.gov/collaborating-institutions.
[81] L. Casas and Collaborators, Data to reproduce Figs. 1–13 of "Validation of the DESI DR2 Lyα BAO analysis using synthetic datasets" (2025), https://zenodo.org/records/17947139.
[82] DESI Collaboration, DESI Data Release 2 (2025), https://data.desi.lbl.gov/doc/releases/.
[83] M.-F. Ho, S. Bird, and R. Garnett, Mon. Not. R. Astron. Soc. **496**, 5436 (2020).
[84] B. Wang, J. Zou, Z. Cai, J. X. Prochaska, Z. Sun, J. Ding, A. Font-Ribera, A. Gonzalez, H. K. Herrera-Alcantar, V. Irsic *et al.*, Astrophys. J. Suppl. Ser. **259**, 28 (2022).