Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **553,** 1–16 (2026)
Advance Access publication 2026 October 29

https://doi.org/10.1093/mnras/stag2028

# Forecast for the detectability of patchy hydrogen reionization in WEAVE-QSO measurements of the Lyman-$\alpha$ forest power spectrum at redshift $z \geqslant 4$

Ke Ma [1]★, James S. Bolton [1], Vid Iršič [2], Prakash Gaikwad [3], Matthew M. Pieri [4], Trystyn A. M. Berg [5,6], Rajeshwari Dutta [7], Matteo Fossati [5,8], Michele Fumagalli [5,9], Emanuel Gafton [10], Ignasi Pérez Ràfols [11] and Francesco Pistis [8,12]

[1]*School of Physics and Astronomy, The University of Nottingham, University Park, Nottingham NG7 2RD, UK*
[2]*Centre for Astrophysics Research, Department of Physics, Astronomy and Mathematics, University of Hertfordshire, College Lane, Hatfield AL10 9AB, UK*
[3]*Department of Astronomy, Astrophysics and Space Engineering, Indian Institute of Technology Indore, Simrol, MP 453552, India*
[4]*Aix Marseille Université, CNRS, CNES, Laboratoire d'Astrophysique de Marseille, F-13388 Marseille, France*
[5]*Dipartimento di Fisica 'G. Occhialini', Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy*
[6]*Department of Physics and Astronomy, Camosun College, 3100 Foul Bay Road, Victoria, BC V8P 5J2, Canada*
[7]*IUCAA, Postbag 4, Ganeshkhind, Pune 411007, Maharashtra, India*
[8]*INAF – Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, via E. Bianchi 46, I-23807 Merate, Italy*
[9]*INAF – Osservatorio Astronomico di Trieste, via G.B. Tiepolo 11, I-34143 Trieste, Italy*
[10]*Isaac Newton Group of Telescopes, Apartado 321, E-38700 Santa Cruz de la Palma, Tenerife, Spain*
[11]*Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 10, E-08930 Barcelona, Spain*
[12]*National Centre for Nuclear Research, ul. Pasteura 7, PL-02-093 Warsaw, Poland*

Accepted 2026 October 21. Received 2026 October 16; in original form 2026 August 11

## ABSTRACT

We present the first detailed forecasts for the detectability of patchy hydrogen reionization in the one-dimensional Ly$\alpha$ forest power spectrum to be measured by the WEAVE-QSO survey. Using the Sherwood-Relics reionization simulations and a WEAVE-QSO survey configuration, we generate mock spectra in four redshift bins, $z = 4.0, 4.2, 4.4,$ and 4.6, in which relic ionization and temperature fluctuations from patchy hydrogen reionization enhance the Ly$\alpha$ forest power spectrum on large scales (i.e. at wavenumber $k \sim 10^{-3}\,\mathrm{s\,km^{-1}}$). Our Ly $\alpha$ forest pipeline forecasts the power spectrum covariance by considering sample size, spectral resolution, noise subtraction, continuum placement, metal contamination, and damping wings from high column density absorbers. Applying our covariance forecast within a Bayesian parameter inference framework, we find that the signature of patchy hydrogen reionization should be detectable at a significance of $\simeq 4.5\sigma$. The forthcoming WEAVE-QSO 1D power spectrum measurements should therefore be able to directly detect and characterize the large-scale relic imprint of patchy hydrogen reionization in the Ly$\alpha$ forest power spectrum at $z \geq 4$.

**Key words:** methods: numerical – intergalactic medium – quasars: absorption lines.

## 1 INTRODUCTION

The intergalactic medium (IGM), as the largest reservoir of baryons in the Universe, plays a crucial role in tracing large-scale cosmic structure. The IGM is primarily studied through Lyman-$\alpha$ (Ly $\alpha$) forest absorption (A. A. Meiksin 2009; M. McQuinn 2016) detected in the spectra of distant quasi-stellar objects (QSOs). Statistical analysis of the Ly $\alpha$ forest offers insights into the underlying matter density fluctuations (R. A. C. Croft et al. 2002; M. Viel, M. G. Haehnelt & V. Springel 2004; P. McDonald et al. 2005) as well as the thermal and ionization state of the IGM (M. Rauch et al. 1997; J. Schaye et al. 2000; J. S. Bolton et al. 2005; P. Gaikwad et al. 2021) across a range of scales and redshifts.

The statistical properties of the Ly $\alpha$ forest are often quantified by the one-dimensional flux power spectrum, $P_{\mathrm{Ly}\alpha}$, which measures the amplitude of Ly $\alpha$ absorption fluctuations as a function of wavenumber, $k$. Over the past two decades, $P_{\mathrm{Ly}\alpha}$ has been established as a flexible observable with distinct scientific applications spanning various scales and redshifts (N. Palanque-Delabrouille et al. 2013; V. Iršič et al. 2017b; N. G. Karaçaylı, A. Font-Ribera & N. Padmanabhan 2020; N. G. Karaçaylı et al. 2022; C. Ravoux et al. 2023; N. G. Karaçaylı et al. 2024). On small scales [wavenumbers $\log_{10}(k/\mathrm{s\,km^{-1}}) \gtrsim -1.5$], the shape of the power spectrum is sensitive to the Jeans smoothing of the IGM and the free-streaming length of dark matter particles. This sensitivity allows constraints to be placed on warm dark matter models (V. Iršič et al. 2017a; D. C. Hooper et al. 2022; B. Villasenor et al. 2023; V. Iršič et al. 2024; O. Garcia-Gallego et al. 2025; K. K. Rogers &

★ E-mail: ke.ma@nottingham.ac.uk

© The Author(s) 2026.
Published by Oxford University Press on behalf of Royal Astronomical Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

V. Poulin 2025), primordial magnetic fields (M. Pavičević et al. 2025), and isocurvature models (O. Garcia-Gallego et al. 2026).

Another key aspect of $P_{\mathrm{Ly}\alpha}$ at high redshifts ($z \gtrsim 4$) lies in its connection to the epoch of hydrogen reionization. During reionization, the first stars and galaxies photoionized and photoheated their surrounding environments, driving the formation and expansion of ionized bubbles. As these bubbles grew and overlapped, the photoheated gas slowly cooled and imprinted spatial inhomogeneities in the IGM temperature distribution (A. A. Meiksin 2009; M. McQuinn 2016). Due to the long timescales associated with adiabatic cooling (T. Theuns et al. 2002; L. Hui & Z. Haiman 2003), these thermal inhomogeneities persist well after the completion of reionization at $z \sim 6$, remaining detectable down to $z \sim 4$ (L. C. Keating, E. Puchwein & M. G. Haehnelt 2018; X. Wu et al. 2019; M. Molaro et al. 2022). At fixed gas density and for photoionization equilibrium, hotter gas has a lower hydrogen recombination coefficient and hence a lower neutral fraction, resulting in weaker Ly $\alpha$ absorption. These large-scale 'thermal relics' of reionization introduce additional fluctuations in the Ly $\alpha$ optical depth, thereby enhancing the large-scale [i.e. $\log_{10}(k/\mathrm{s\,km^{-1}}) \lesssim -1.5$] $P_{\mathrm{Ly}\alpha}$ at $z \gtrsim 4$ (R. Cen et al. 2009; A. D'Aloisio, M. McQuinn & H. Trac 2015; M. Molaro et al. 2023). In particular, M. Molaro et al. (2023) investigated this enhancement using the high resolution $P_{\mathrm{Ly}\alpha}$ measurements from N. G. Karaçaylı et al. (2022). Although M. Molaro et al. (2023) only found tentative ($2.7\sigma$) evidence for enhanced large scale power in $P_{\mathrm{Ly}\alpha}$, the limited precision of the N. G. Karaçaylı et al. (2022) measurements from the Keck Observatory Database of Ionized Absorption toward Quasars J. M. O'Meara et al. (2015), Spectral Quasar Absorption Database M. T. Murphy et al. (2019) and XQ-100 survey (S. López et al. 2016) (with a redshift path length of only $\Delta z = 17.2$ at $z > 4$) meant that the expected signature of inhomogeneous reionization could not be recovered at high significance. M. Molaro et al. (2023) concluded that improved statistical precision and good control of the systematic uncertainties at large scales were required for further progress.

In this context, the $P_{\mathrm{Ly}\alpha}$ measurements from massive spectroscopic surveys with resolution $1500 < R < 7000$, such as the Baryon Oscillation Spectroscopic Survey of the Sloan Digital Sky Survey (SDSS/BOSS; N. Palanque-Delabrouille et al. 2013), extended-BOSS (SDSS/eBOSS; S. Chabanier et al. 2019), and the Dark Energy Spectroscopic Instrument (DESI, C. Ravoux et al. 2023; N. G. Karaçaylı et al. 2024, 2025; J. Chaves-Montero et al. 2026) represent an exciting opportunity. Although restricted in their ability to access small scales due to spectral resolution ($\Delta v \sim 40$–$200\,\mathrm{km\,s^{-1}}$ for $R = 7000$–$1500$), the large sample sizes result in smaller statistical uncertainties compared to the high resolution $P_{\mathrm{Ly}\alpha}$ measurements. Ongoing surveys such as DESI (DESI Collaboration 2016), the William Herschel Telescope Enhanced Area Velocity Explorer (WEAVE; S. Jin et al. 2024), the BarYon Cycle project (R. Szakacs et al. 2023) on the 4-meter Multi-Object Spectroscopic Telescope (R. S. Jong et al. 2019), along with planned facilities including the Prime Focus Spectrograph (J. Greene et al. 2022), the MUltiplexed Survey Telescope (C. Zhao et al. 2026), and the wide-field spectroscopic telescope (V. Mainieri et al. 2024) are expected to further refine these measurements and provide new insights into the epoch of reionization through higher spectral resolution, larger sample sizes, and increased QSO number densities. Realizing this potential, however, requires robust synthetic data pipelines that include patchy reionization modelling to forecast systematic uncertainties, optimize survey strategies, and validate analysis frameworks.

In this paper, we therefore introduce a new pipeline for generating detailed Ly $\alpha$ forest mocks tailored to large-scale spectroscopic surveys. Our approach is based on the Sherwood-Relics simulation suite (J. S. Bolton et al. 2017; E. Puchwein et al. 2023) – a set of cosmological hydrodynamical simulations specifically designed to study the IGM and the Ly $\alpha$ forest during and after reionization. In contrast to other recent spectroscopic survey pipelines, we use Sherwood-Relics to model the impact of inhomogeneous reionization on the thermal and ionization history of the IGM. Furthermore, the mocks we construct here are designed to replicate the instrumental properties of the WEAVE-QSO survey (M. M. Pieri et al. 2016), as part of WEAVE (S. Jin et al. 2024), which is optimized for quasar absorption science. We use these mocks to obtain the first detailed forecast of the precision to which WEAVE-QSO will measure $P_{\mathrm{Ly}\alpha}$ at redshift $z \geq 4$, and to assess the detectability of large-scale power arising from patchy hydrogen reionization. We note that some recent work (P. Montero-Camacho & Y. Mao 2021; Y. Zheng et al. 2026) has also explored the Ly $\alpha$ forest power spectrum as a probe of reionization and cosmic-dawn astrophysics. However, these studies focus on temperature fluctuations seeded by the hydrodynamical response of gas during reionization that may remain detectable at $z < 4$ (C. M. Hirata 2018; C. Cain et al. 2024); here, we instead focus on the large-scale thermal relic imprint predicted by patchy reionization models at $z > 4$ (e.g. L. C. Keating et al. 2018; X. Wu et al. 2019; M. Molaro et al. 2022).

This paper is structured as follows: Section 2 details the method for generating QSO spectra mocks. Section 3 outlines the computation of $P_{\mathrm{Ly}\alpha}$ and describes how instrumental effects and absorption line contamination are modelled to construct the associated overall uncertainty. In Section 4, we present forecasts for detecting the relic signature of inhomogeneous reionization using our $P_{\mathrm{Ly}\alpha}$ pipeline at redshifts $z \geq 4$. Finally, a summary is provided in Section 5.

## 2 CONSTRUCTION OF MOCK LY $\alpha$ FOREST SPECTRA

### 2.1 Hydrodynamical simulation

In this work, we employ the Sherwood-Relics simulation suite (E. Puchwein et al. 2023) for generating mock Ly $\alpha$ forest spectra. Sherwood-Relics was performed with a customised version of the cosmological hydrodynamical code P-GADGET-3 (V. Springel 2005). The underlying cosmology in the simulation used here is flat $\Lambda$CDM, parametrized by $\Omega_\Lambda = 0.692$, $\Omega_\mathrm{m} = 0.308$, $\Omega_\mathrm{b} = 0.0482$, $\sigma_8 = 0.829$, $n_\mathrm{s} = 0.961$, $h = 0.678$, with a primordial helium abundance by mass $Y = 0.24$. The dark matter particle mass is $M_\mathrm{dm} = 5.37 \times 10^5\, h^{-1}\,\mathrm{M_\odot}$ and the gas particle mass is $M_\mathrm{gas} = 9.97 \times 10^4\, h^{-1}\,\mathrm{M_\odot}$, which are chosen to adequately resolve the small-scale structure of the Ly $\alpha$ forest (J. S. Bolton & G. D. Becker 2009). We use simulated Ly $\alpha$ optical-depth skewers extracted from a $40\,h^{-1}$ cMpc box with $2 \times 2048^3$ particles, following the procedure described in K. Ma et al. (2026).

### 2.2 The WEAVEify algorithm

We construct our mock spectra assuming survey properties that resemble those of WEAVE-QSO (M. M. Pieri et al. 2016), one of the eight core surveys of WEAVE (S. Jin et al. 2024). The



Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

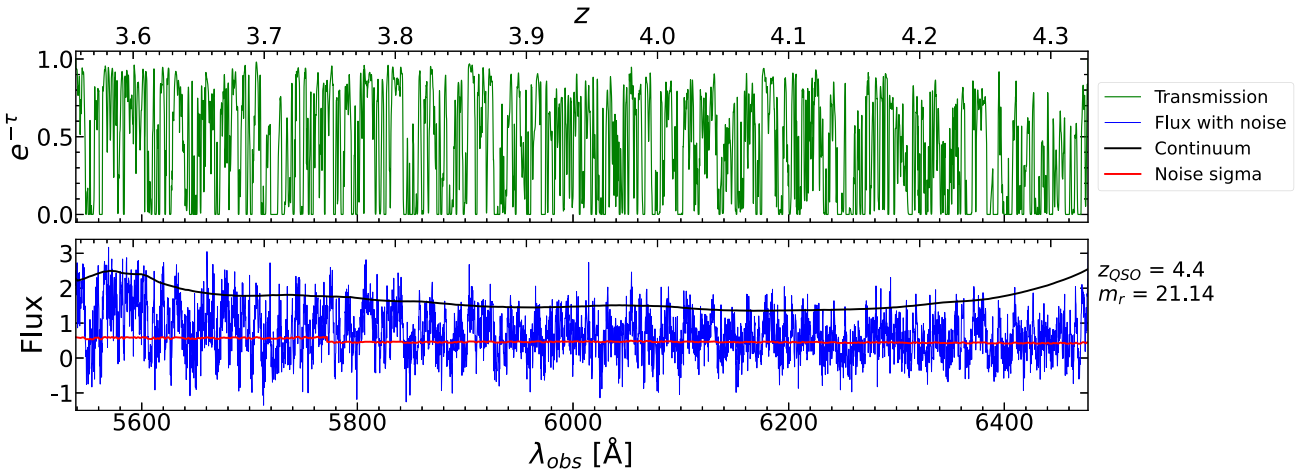



**Figure 1.** An example mock QSO spectrum constructed using Sherwood-Relics and WEAVEify. The top panel shows the Ly $\alpha$ forest transmission constructed from one of the Sherwood-Relics hydrodynamical simulations (E. Puchwein et al. 2023). The bottom panel shows the WEAVEified spectrum when taking the Ly $\alpha$ transmission in the top panel as input, for a QSO with systemic redshift $z_{\rm QSO} = 4.4$ and $r$-band magnitude 21.14. Note that the Ly $\alpha$ transmission is only modelled between the blue edge of the QSO proximity region of extent 4000 km s$^{-1}$ and the rest-frame Ly $\beta$ wavelength (see text for details). The inverse signal-to-noise per pixel, shown as $\sigma = (\mathrm{S/N})^{-1}$, is displayed by the red curve.

WEAVE-QSO survey will obtain spectra of approximately 450 000 QSOs at redshifts $z_{\rm QSO} > 2.2$, covering a nominal survey area of approximately 10 000 deg$^2$ to $m_r < 21.5$, with $z > 4$ QSOs additionally targeted over $21.5 < m_r < 22$. Also a further nested tier in the HETDEX spring field will provide 100 QSOs per deg$^2$.

We have developed a software package, WEAVEify, to fully forward model spectra representative of typical observations from the WEAVE-QSO survey. It has been rigorously tested with operation rehearsal (OpR) mock spectra of the WEAVE survey, and accurately models noise at the detector level. Further details on WEAVEify will be presented in a forthcoming publication (Gaikwad et al., in preparation). Similar WEAVE-like mock quasar spectra have also been used in recent WEAVE continuum-modelling work (F. Pistis et al. 2025). Here, we focus on describing the core functionality necessary for obtaining Ly $\alpha$ forest mocks.

Starting from input optical-depth skewers drawn from the Sherwood-Relics hydrodynamical simulation, WEAVEify maps the rest-frame optical depths to a common set of wavelength bins based on the systemic redshift of the QSO. Once aligned, the total optical depth is computed by summing the individual optical depths, $\tau$, from any of the transitions required in the mock spectrum (e.g. Ly $\alpha$, Si III). The transmission is computed as $F = \mathrm{e}^{-\tau}$. WEAVEify then applies wavelength-dependent instrumental broadening by convolving the transmission with a Gaussian function whose standard deviation varies with wavelength. WEAVEify then resamples the broadened transmission from the simulation wavelength grid onto the WEAVE wavelength solution using linear interpolation. For continuum placement, WEAVEify employs a principal component analysis (PCA)-based method, with PCA components derived from a large set of SDSS QSO continua as outlined in I. Pâris et al. (2011) (also see F. B. Davies et al. 2018). These components are used to project the continuum from the red to the blue side of the rest-frame Ly $\alpha$ wavelength.

In the final step, WEAVEify adds realistic noise to the spectrum. As WEAVE is a fibre-fed spectrograph, the spectra are photon-limited, with the noise governed by Poisson statistics. Photon counts in each fibre arise from four primary sources: (i) photons from the source within the fibre's field of view, (ii) sky photons, (iii) dark current, and (iv) detector readout noise. WEAVEify models the photon counts for the source and sky by incorporating key parameters such as source and sky brightness, exposure time, and the spatial extent of the source (point-like or extended). It further accounts for atmospheric observing conditions, including the seeing profile (with a Moffat function with FWHM 0.75 arcsec and $\beta = 2.5$) and an airmass of 1.107. The OpR3 noise-model defaults here use a fixed exposure time of 3060 s and a sky brightness of $m_B = 20.92$. Noise characteristics dependent on the telescope's engineering design, such as fibre diameter, fibre offset, CCD size, and telescope throughput, are also included to accurately estimate photon counts. Detector dark current and readout noise are included as separate noise terms. The overall noise is then calculated as a quadrature sum of these terms.

### 2.3 WEAVEifying the Ly α forest transmission

The mock Ly $\alpha$ optical depths used as the input to WEAVEify are obtained by combining different snapshots of the Sherwood-Relics simulation. The snapshots have a redshift cadence of $\Delta z = 0.1$ and contain Ly $\alpha$ optical depths along 5000 skewers drawn parallel to the simulation box boundaries. A full synthetic QSO line of sight is generated by connecting multiple such flux skewers end to end.

The top panel of Fig. 1 shows an example of the resulting Ly $\alpha$ forest transmission for a mock QSO with systemic redshift $z_{\rm QSO} = 4.4$. The optical depths are rescaled by a common multiplicative factor within each redshift interval, such that the mean transmission averaged over all skewers and pixels in that interval matches the effective optical depth measured by G. D. Becker et al. (2013). This reference model may underestimate $\tau_{\rm eff}$ at $z > 4$ (see the discussion in E. Boera et al. 2019); in our parameter inference, however, $\tau_{\rm eff}$ is varied independently in each redshift bin and marginalized over, rather than fixed to this prescription

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

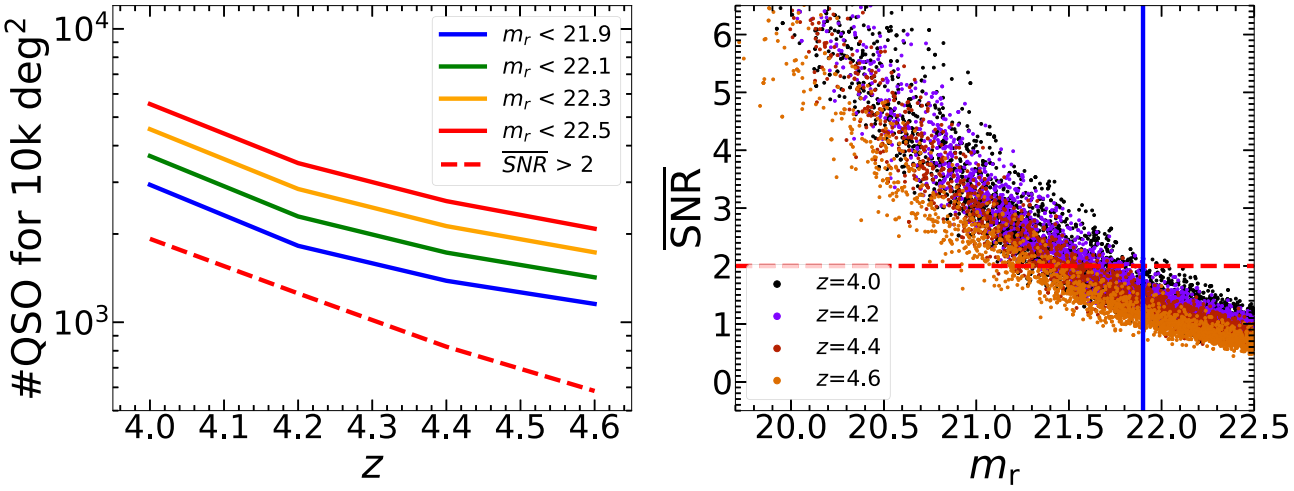



**Figure 2.** Left panel: The cumulative number of QSOs available in a 10 000 deg$^2$ survey for four Ly $\alpha$ forest power spectrum bins centred at $z =$ 4.0, 4.2, 4.4, and 4.6. The numbers are derived using the QSO luminosity function measurements presented by N. Palanque-Delabrouille et al. (2016). The results are shown for different $r$-band magnitude cuts, $m_r$ (solid curves), or a mean signal-to-noise ratio cut, $\overline{\text{SNR}}$ (dashed curve). Right panel: Scatter plot of $\overline{\text{SNR}}$ against the $r$-band magnitude for the QSOs in each redshift bin. The blue solid line corresponds to a $m_r$ cut at 21.9 and the red dashed line corresponds to an $\overline{\text{SNR}}$ cut at 2. We adopt an $r$-band magnitude cut $m_r < 21.9$ throughout this paper.

(Section 4.1). Each skewer is periodically shifted to ensure continuity at the points where the skewers are joined. For each mock QSO spectrum, the Ly $\alpha$ forest absorption is modelled only at wavelengths redward of the Ly $\beta$ emission line, below which contamination from higher order Lyman series lines becomes significant. Regions blueward of the Ly $\alpha$ emission line that are affected by the QSO proximity effect (R. F. Carswell et al. 1982; H. S. Murdoch et al. 1986) are also excluded. We account for this proximity effect over a velocity interval of 4000 km s$^{-1}$ (e.g. T. S. Kim et al. 2007). The Ly $\alpha$ transmission is then converted into a mock observed spectrum using the WEAVEify software. An example of the resulting mock QSO spectrum (blue) is shown in the lower panel of Fig. 1.

### 2.4 Generating the mock WEAVE-QSO survey sample

Using the procedure described in the previous section, for each of four $P_{\text{Ly}\alpha}$ redshift bins centred at $z = 4.0,\ 4.2,\ 4.4$, and 4.6 we generate 10 000 mock QSO spectra that have systemic redshifts selected to have full spectral coverage within the bin. These bins target the redshift range where relic fluctuations from patchy reionization are expected to remain observable while retaining sufficient WEAVE-QSO sightlines for a useful forecast. The distributions of the QSO $r$-band magnitude, $m_r$, and systemic redshift, $z_{\text{QSO}}$, are sampled using a Monte Carlo method based on the QSO number density derived from the N. Palanque-Delabrouille et al. (2016) luminosity function using their PLE + LEDE model predictions at $0 < z < 6$. Our reference forecast will therefore assume a complete magnitude-limited sample subject to the forest-coverage requirements above, without additional losses from target selection, fibre allocation or unsuccessful observations. Uncertainties in the luminosity function and departures from this completeness assumption may change the available QSO counts and their magnitude and redshift distributions, and hence the forecast precision. Note also that the $r$-band magnitude and systemic redshift are assigned to the simulated sightlines independently of the Ly $\alpha$ absorption in our mocks. We then scale this luminosity-function-based number density distribution to a reference forecast sample covering 10 000 deg$^2$, applying the relevant $m_r$ and $z_{\text{QSO}}$ selections to obtain the expected QSO counts used below. For each set of 10 000 mock QSO spectra, we also determine the fraction of QSOs in each redshift bin satisfying specific magnitude and signal-to-noise criteria (described below). These fractions are then scaled to obtain the total expected number of suitable QSOs across the full 10 000 deg$^2$ survey, as illustrated in the left panel of Fig. 2.

The QSO sample cuts for $P_{\text{Ly}\alpha}$ are determined as follows. The $r$-band magnitude $m_r$ influences the noise level in the Ly $\alpha$ forest region, which directly affects the measurement of $P_{\text{Ly}\alpha}$. We quantify the noise in our mocks using $\overline{\text{SNR}}$, defined as the signal-to-noise ratio per Å averaged over all pixels within a given redshift bin. Previous studies of $P_{\text{Ly}\alpha}$ at $2 \lesssim z \lesssim 4$ (e.g. S. Chabanier et al. 2019; C. Ravoux et al. 2023) have often applied $\overline{\text{SNR}}$ cuts of 2 or 3 to mitigate noise. However, at the comparatively high redshifts considered in this work where the number of sufficiently bright QSOs is smaller, such cuts would drastically reduce the number of usable QSOs and increase the statistical uncertainty on $P_{\text{Ly}\alpha}$. Hence, to maintain sample size while controlling noise in our mocks, we therefore instead implement a cut in $r$-band magnitude. Specifically, we adopt $m_r < 21.9$ for all redshift bins, but without an additional SNR cut; the SNR cuts shown in Fig. 2 are included for comparison only. Note again that because the assigned magnitudes are independent of the simulated absorption, this selection does not preferentially retain higher-transmission sightlines in our mocks, as might be expected if using an observed $r$-band-selected sample at $z > 4$. In practice, rather than the simple $r$-band cut we use here, for WEAVE-QSO the planned selection strategy will use prior redshift estimates and the forest SNR will be optimized using a spectral or narrow-band information redward of Ly $\alpha$ emission at each redshift.

As shown in the left panel of Fig. 2, this magnitude cut retains over 1000 QSOs even in our highest redshift bin at $z = 4.6$. After applying the cut $m_r < 21.9$, we find the number of suitable QSOs in a 10 000 deg$^2$ survey is $N_{\text{QSO}} = 2942$, 1822, 1385, and 1155 for



Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

$P_{\mathrm{Ly}\alpha}$ redshift bins $z = 4.0,\ 4.2,\ 4.4$, and 4.6, respectively. We use these numbers throughout the rest of the paper, but in Section 4.2 we will briefly discuss how a smaller number of suitable QSOs (e.g. through incompleteness or reduced survey area) may impact on our forecast.

## 3 POWER SPECTRUM ESTIMATION

### 3.1 FFT estimator

We now discuss the calculation of the 1D power spectrum and its statistical and systematic uncertainties using our QSO survey mocks. The two approaches for computing the $P_{\mathrm{Ly}\alpha}$ commonly used in the literature are the fast Fourier transform (FFT; e.g. N. Palanque-Delabrouille et al. 2013; S. Chabanier et al. 2019; C. Ravoux et al. 2023) and the quadratic maximum likelihood estimator (QMLE; e.g. N. G. Karaçaylı et al. 2022, 2024). Although the latter benefits from being less sensitive to masking of the Ly α forest spectrum, the former method is more straightforward. In addition, C. Ravoux et al. (2023) and N. G. Karaçaylı et al. (2024) measured the $P_{\mathrm{Ly}\alpha}$ on the same DESI data sets using FFT and QMLE, respectively, and found that the results are consistent within 1 per cent up to the half Nyquist frequency. We therefore use the FFT approach to measure $P_{\mathrm{Ly}\alpha}$.

When performing power spectra analysis, the commonly used estimator is the flux contrast $\delta_{\mathrm{Ly}\alpha}(\lambda)$, where

$$\delta_{\mathrm{Ly}\alpha}(\lambda) = \frac{F_{\mathrm{Ly}\alpha}(\lambda)}{\overline{F}(\lambda)} - 1. \tag{1}$$

Here, $F_{\mathrm{Ly}\alpha}(\lambda)$ is the Ly α transmission and $\overline{F}(\lambda)$ is the mean transmission over some wavelength interval (which also depends on redshift). In this work we assume a mean transmission based on the Ly α forest effective optical depth measurement presented by G. D. Becker et al. (2013), where $\tau_{\mathrm{eff}} = -\ln \bar{F}$. The power spectrum $P_{\mathrm{Ly}\alpha}(k)$ is then given by

$$P_{\mathrm{Ly}\alpha}(k) = |\mathcal{F}(\delta_{\mathrm{Ly}\alpha}(\lambda))|^2, \tag{2}$$

where $\mathcal{F}$ represents the Fourier transform. In practice, $P_{\mathrm{Ly}\alpha}$ is obtained by averaging the individual power spectra obtained from each of the QSO spectra within the redshift bin under consideration.

However, $\delta_{\mathrm{Ly}\alpha}(\lambda)$ is not directly observable. Instead, the observed flux contrast, $\delta_{\mathrm{obs}}(\lambda)$, is given by

$$\delta_{\mathrm{obs}}(\lambda) = \frac{F_{\mathrm{obs}}(\lambda)}{C_{\mathrm{QSO}}(\lambda)\,\overline{F}(\lambda)} - 1, \tag{3}$$

where $C_{\mathrm{QSO}}(\lambda)$ is the intrinsic QSO emission obtained from the continuum fitting procedure. For now, we shall assume the QSO continuum is perfectly known, but we will return to modelling continuum fitting uncertainties later.

The underlying Ly α flux contrast, $\delta_{\mathrm{Ly}\alpha}(\lambda)$, differs from the observed flux contrast, $\delta_{\mathrm{obs}}(\lambda)$, due to instrumental resolution effects as well contamination by metal lines present in the Ly α forest. The combination of these effects along with noise yields the observed flux. The relation between $\delta_{\mathrm{Ly}\alpha}(\lambda)$ and $\delta_{\mathrm{obs}}(\lambda)$ can be expressed as

$$\delta_{\mathrm{obs}}(\lambda) = [\delta_{\mathrm{Ly}\alpha}(\lambda) + \delta_{\mathrm{metals}}(\lambda)]W(\lambda, R, \Delta\lambda) + \delta_{\mathrm{noise}}(\lambda), \tag{4}$$

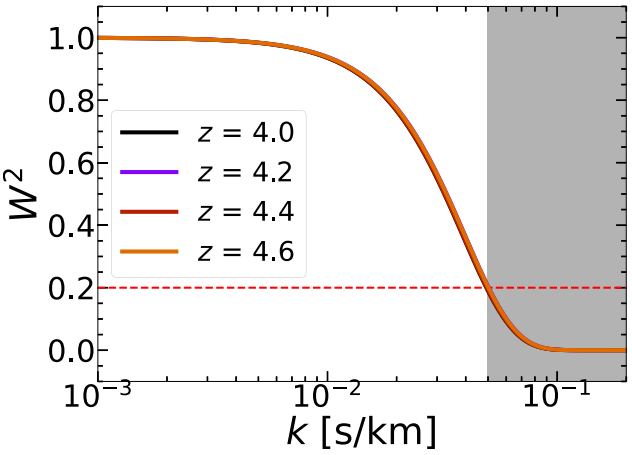



**Figure 3.** The instrumental resolution correction factor, $W^2(k, R, \Delta\lambda)$, for the four $P_{\mathrm{Ly}\alpha}$ redshift bins considered in this work. Note the curves are almost indistinguishable. The red dashed horizontal line marks the threshold $W^2 = 0.2$, and the grey shaded region shows the maximum wavenumber for the $P_{\mathrm{Ly}\alpha}$ measurements that we consider in this work, $k_{\mathrm{max}} = 0.05$ s km$^{-1}$, where the resolution correction factor falls below this value. This choice follows the threshold used by the DESI Collaboration (C. Ravoux et al. 2023).

where $W$ is the spectrograph line spread function as a function of wavelength λ, spectral resolution $R$, and the pixel size $\Delta\lambda$. The terms $\delta_{\mathrm{metals}}$ and $\delta_{\mathrm{noise}}$ represent the contribution to the observed flux contrast by intervening metal lines and noise, respectively. The contribution from high column density systems is treated separately in Section 3.5.

After Fourier transforming equation (4) using equation (2), we obtain the relation between the observed power spectrum, $P_{\mathrm{obs}}$, and the Ly α forest power spectrum, $P_{\mathrm{Ly}\alpha}$:

$$P_{\mathrm{obs}}(k) = [P_{\mathrm{Ly}\alpha}(k) + P_{\mathrm{metals}}(k)]W^2(k, R, \Delta\lambda) + P_{\mathrm{noise}}(k). \tag{5}$$

Once again, $P_{\mathrm{metals}}$ and $P_{\mathrm{noise}}$ represent the contribution to the observed power spectrum by metals and noise. We now turn to consider each of these components and their contribution to the measurement uncertainties in turn.

### 3.2 Spectrograph resolution correction

The WEAVE spectrograph low-resolution mode covers 366–959 nm at a representative resolving power of $R \sim 5000$, while the high-resolution mode provides restricted wavelength coverage at $R \sim 20\,000$ (M. M. Pieri et al. 2016; S. Jin et al. 2024). The majority of QSO spectra in WEAVE-QSO will be obtained using the low-resolution mode. We assume $R \sim 5000$ in this work.

The WEAVEify package characterizes the effect of the spectrograph resolution using the line spread function matrix $W(\lambda, R, \Delta\lambda)$. As a forecast-level approximation, a Gaussian kernel is assumed with a standard deviation that varies as a function of wavelength, as detailed in Section 2.2. The square of the Fourier transform of this resolution matrix corresponds to the resolution correction factor $W^2(k, R, \Delta\lambda)$ that is applied to the power spectrum in equation (5). The corresponding resolution windows for the four redshift bins we consider are shown in Fig. 3; these remain very similar in all four bins. We use these correction factors to determine the $k$-range impacted by instrument resolution and to estimate the contribution of the instrumental

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

resolution correction to the final systematic uncertainty estimate. The maximum wavenumber, $k_{\rm max}$, for the $P_{{\rm Ly}\,\alpha}$ measurements that we consider in this work is selected following the same criterion as the DESI Collaboration (C. Ravoux et al. 2023), where $k_{\rm max}$ is defined as the point beyond which the resolution correction falls below 0.2. This yields $k_{\rm max} = 0.05\,{\rm s\,km^{-1}}$.

To estimate the uncertainty associated with applying this resolution correction, we draw 1000 bootstrap realizations of $N_{\rm QSO}$ sightlines from our mock spectra. For each realization, we compare the resolution corrected estimate with the underlying model and measure the scatter in their ratio across the bootstrap samples. The resulting term quantifies the residual uncertainty in the correction procedure. However, we note that this uncertainty quantifies the robustness of the correction for a fixed, known resolution window function. In a real WEAVE-QSO analysis, the resolving power will vary with wavelength and observing conditions, and uncertainty in $R$ will propagate into uncertainty in $W^2(k, R, \Delta\lambda)$. This effect can become important near the high-$k$ limit of the measurement and should ultimately be treated using a wavelength-dependent resolution matrix or window-function uncertainty model. We discuss this further in Section 3.7.

### 3.3 Noise subtraction

The noise level is quantified by $\sigma_{\rm f}(\lambda)$, the standard deviation of the flux. In observational studies, this quantity is typically estimated from the data reduction pipeline or using methods such as the exposure difference technique (e.g. S. Chabanier et al. 2019; C. Ravoux et al. 2023). For our mock data, however, $\sigma_{\rm f}(\lambda)$ is directly provided as an output of the WEAVEify software. The standard deviation of the noise term $\delta_{\rm noise}(\lambda)$ in equation (4) is thus defined as

$$\sigma_\delta(\lambda) = \frac{\sigma_{\rm f}(\lambda)}{C_{\rm QSO}(\lambda)\,\overline{F}(\lambda)}. \tag{6}$$

In each redshift bin, for each pixel at wavelength $\lambda$ we generate 100 realizations of $\delta_{\rm noise}(\lambda)$ by drawing from a normal distribution with standard deviation $\sigma_\delta(\lambda)$. The noise power spectrum $P_{\rm noise}(k)$ in equation (5) is then estimated as $P_{\rm noise}(k) = \langle \|\mathcal{F}(\delta_{\rm noise}(\lambda))|^2 \rangle$, where the average is computed over all QSO spectra and $N_{\rm noise}$ realizations within the redshift bin. Note that the use of a normal distribution for generating $\delta_{\rm noise}(\lambda)$ reflects the assumption of Gaussian noise in WEAVEify. However, the noise model can be adapted as needed to match specific instrumental characteristics or data reduction pipelines.

Fig. 4 displays the observed power spectra and the estimated noise power spectra for the four $P_{{\rm Ly}\,\alpha}$ redshift bins considered in this work. The noise component becomes dominant beyond $k_{\rm max} = 0.05\,{\rm s\,km^{-1}}$, which is also the wavenumber limit adopted for the spectral resolution correction. In a similar manner to the resolution correction uncertainty, we again estimate the uncertainty associated with noise subtraction using 1000 bootstrap realizations of $N_{\rm QSO}$ sightlines. For each realization, we compare the noise-subtracted power spectrum estimate with the corresponding noiseless mock estimate and take the standard deviation of this correction ratio across bootstrap samples.

### 3.4 Metal absorption lines

The metal flux contrast, $\delta_{\rm metals}$, consists of two components: one originating from intervening metal lines with rest-frame wavelengths $\lambda > \lambda_{{\rm Ly}\,\alpha}$, and the other from metal lines that are correlated with the Ly $\alpha$ forest with rest-frame wavelengths $\lambda \sim \lambda_{{\rm Ly}\,\alpha}$. The first category includes C IV ($\lambda\lambda$1548, 1551), Si IV ($\lambda\lambda$1394, 1403), and Mg II ($\lambda\lambda$2796, 2804), which we will refer to as uncorrelated metals. The second category is dominated by Si III ($\lambda$1207) in our modelling, which we refer to as correlated metals. We have also tested Si II $\lambda\lambda$1190, 1193 in this modelling framework, but found its contribution to be sub-per cent over the $k < k_{\rm max}$ range considered here; see K. Ma et al. (2026) for further details.

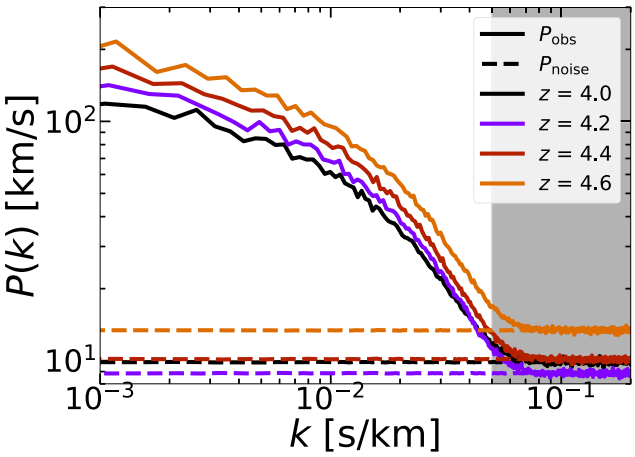



**Figure 4.** Comparison between the estimated noise power spectra $P_{\rm noise}$ from WEAVEify (dashed lines) and the observed power spectra $P_{\rm obs}$ [solid lines, see equation (5) for further details]. The noise will make a significant contribution to the power spectrum uncertainty at the smallest scales (i.e. large $k$). The grey shading shows the excluded region above the maximum wavenumber, $k_{\rm max} = 0.05\,{\rm s\,km^{-1}}$, considered in this work.

In practice, the uncorrelated metal power spectrum, $P_{\rm ucm}$, is typically estimated using sidebands – spectral regions redward of the Ly $\alpha$ rest-frame wavelength that are devoid of Ly $\alpha$ absorption (N. Palanque-Delabrouille et al. 2013; S. Chabanier et al. 2019; C. Ravoux et al. 2023). For this forecast, we generate the baseline Ly $\alpha$ mock spectra without uncorrelated metal absorption and then add synthetic uncorrelated metal features separately to estimate $P_{\rm ucm}$ and its contribution to the uncertainty budget. We do this following the approach of M. Viel et al. (2013). We first estimate the number of absorption lines for each metal line species per unit redshift path length by integrating the appropriate column density distribution function (CDDF; V. D'Odorico et al. 2013; N. L. Mathes, C. W. Churchill & M. T. Murphy 2017; V. D'Odorico et al. 2022) over the column density range $10^{12}\,{\rm cm^{-2}} < N < 10^{15}\,{\rm cm^{-2}}$. Using a Monte Carlo method, we draw column densities and line widths [with $b$-parameter distributions from L. Hui & R. E. Rutledge (1999) assuming $b_\sigma = 10\,{\rm km\,s^{-1}}$] until the required number of metal lines is reached. These synthetic absorption features are then randomly inserted into the Ly $\alpha$ forest region of each mock QSO spectrum. The CDDFs for C IV, Si IV, and Mg II doublets are modelled as power laws using the measurements from V. D'Odorico et al. (2013), V. D'Odorico et al. (2022), and N. L. Mathes et al. (2017), respectively. We use this to estimate the scale-dependent uncertainty associated with uncorrelated metal subtraction. The relative magnitude of $P_{\rm ucm}$ compared with $P_{\rm obs}$ is shown in Fig. 5.Although the contribution of $P_{\rm ucm}$ can become relevant at very small scales, $k > 0.1\,{\rm s\,km^{-1}}$ (M. Viel et al. 2013), at the larger scales of interest in this work it represents only a very small contribution to the total uncertainty (see Fig. 9 later).



Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

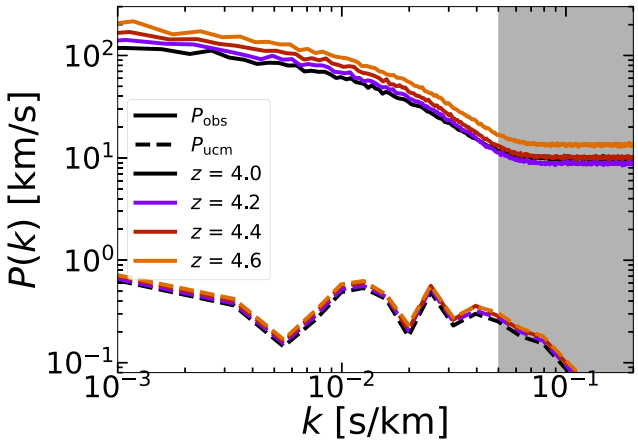



**Figure 5.** Comparison between the uncorrelated metal power spectrum, $P_{ucm}$, predicted by Monte Carlo sampling the observed column density distribution functions for C IV, Si IV, and Mg II (dashed lines) to the total observed power spectrum, $P_{obs}$. The grey shading shows the excluded region above the maximum wavenumber, $k_{max} = 0.05\,s\,km^{-1}$, considered in this work.

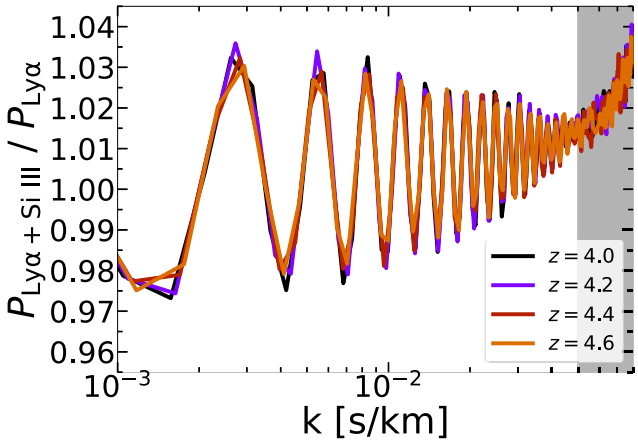



**Figure 6.** The ratio of the combined Ly α and Si III power spectrum, $P_{Ly\alpha+SiIII}$, to the Ly α power spectrum, $P_{Ly\alpha}$. The grey shading shows the excluded region above the maximum wavenumber, $k_{max} = 0.05\,s\,km^{-1}$, considered in this work.

For the correlated metals, we generate Si III absorption spectra following the method described in K. Ma et al. (2026). We assume the Si III optical depths scale as $10^{[Si/H]}$, where the silicon abundance [Si/H] is derived from the measurements of J. Schaye et al. (2003) and A. Aguirre et al. (2004) and the ionization fractions are computed using the photoionization code CLOUDY (G. J. Ferland et al. 1998; M. Chatzikos et al. 2023). Fig. 6 illustrates the impact of Si III (λ1207) on $P_{Ly\alpha}$ and the characteristic oscillations that result (see also P. McDonald et al. 2006). Note, however, that correlated metals like Si III are typically not removed from $P_{Ly\alpha}$ analyses but are instead incorporated as additional fitting parameters. This is because their power spectrum, denoted here as $P_{clm}$, cannot be directly isolated from observations as their absorption lines are blended with the Ly α forest. In this work, correlated metals are therefore not directly included in the systematic uncertainty estimate; instead, we assess in Section 4.1 whether omitting an explicit Si III model affects our parameter inference. For this purpose, we use the four parameter model for Si III absorption recently introduced by K. Ma et al. (2026). As we will discuss later, however, we find that the inclusion of the Si III in our mocks makes very little difference to our final results.

### 3.5 High column density systems

For the purpose of estimating the systematic uncertainty associated with incomplete absorber removal, we model high column density systems (HCDs), defined here as absorbers with neutral hydrogen column densities satisfying $N_{HI} > 10^{19}\,cm^{-2}$. These systems are highly relevant to this study because enhancements to the power spectrum on large scales can readily arise from column densities $N_{HI} > 10^{19}\,cm^{-2}$ due to their extended Lorentzian damping wings. However, the observed incidence of these absorption lines is not captured in our hydrodynamical simulations; self-shielded gas with normalized density $\rho/\langle\rho\rangle > 10^3$ and $T < 10^5$ K is ignored due to the 'Quick-Ly α' implementation (M. Viel et al. 2004) used in Sherwood-Relics (see e.g. fig. 1 in J. S. A. Miller, J. S. Bolton & N. Hatch 2019). We therefore incorporate these HCDs into our mock QSO spectra using a similar method to that employed for the uncorrelated metal lines. The column density distribution function from J. X. Prochaska et al. (2014) is used to estimate the expected number of HCDs per unit redshift path length, which we then Monte Carlo sample over the range $17.2 \leq \log_{10}(N_{HI}/cm^{-2}) \leq 23.0$ defined by J. X. Prochaska et al. (2014) assuming a fixed Doppler parameter $b = 20\,km\,s^{-1}$. The resulting synthetic HCD absorbers are inserted as Voigt profiles at locations where the Ly α optical depth in our mock spectra is largest.

In practice, HCDs cannot be perfectly identified and removed from observational data. To assess the potential bias in $P_{Ly\alpha}$ resulting from incomplete removal, we use the completeness function from B. Wang et al. (2022) for systems with $N_{HI} > 10^{19}\,cm^{-2}$. For each $\log_{10} N_{HI}$ bin, we randomly remove HCDs according to this completeness function – considering our mean signal-to-noise ratio – and retain the remaining systems in the Ly α forest region. The impact of HCDs on $P_{Ly\alpha}$, both with and without incomplete removal, is illustrated in Fig. 7. The systematic error introduced by the incomplete removal (solid curves) is added into the total uncertainty estimate in quadrature. For the incomplete-removal model shown in Fig. 7, the residual HCD contribution to $P_{Ly\alpha}$ is at most ∼5 per cent at $10^{-3}\,s\,km^{-1} \lesssim k < 0.05\,s\,km^{-1}$, with the maximum occurring at the largest scales. Our forecast does not separately propagate uncertainty in the assumed HCD column density distribution. However, steepening and flattening its high column density slope by ±0.2 dex, with the variation amplitude motivated by the observational uncertainties compiled by J. X. Prochaska et al. (2014), changes the residual HCD contribution by less than 1 per cent of the underlying Ly α power across all four redshift bins and the fitted wavenumber range.

### 3.6 Continuum shape

In our mock QSO spectra, a continuum template is assigned to each line of sight by WEAVEify and is subsequently divided out during the computation of $P_{Ly\alpha}$ (as described in Section 3.1). However, in practice the QSO continuum is not known and may be misplaced. As a conservative forecast-level test of this systematic effect, instead of dividing by the true continuum (which varies with wavelength) we divide the WEAVEified mocks by a flat, wavelength independent continuum and evaluate the resulting impact on $P_{Ly\alpha}$. The flat continuum is defined as the mean

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

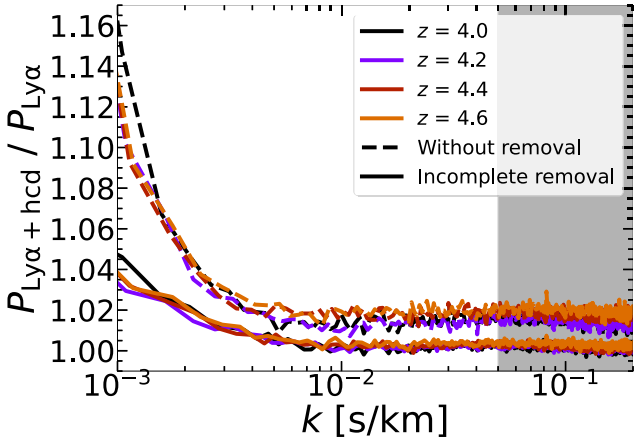



**Figure 7.** The ratio of the combined power spectrum for Ly $\alpha$ and HCDs, $P_{\mathrm{Ly}\alpha+\mathrm{hcd}}$ to the Ly $\alpha$ power spectrum, $P_{\mathrm{Ly}\alpha}$. The dashed curves show the ratio including all HCDs, whereas the solid curves give the ratio after removing a fraction of the HCDs based on the completeness function reported by B. Wang et al. (2022). The grey shading shows the excluded region above the maximum wavenumber, $k_{\mathrm{max}} = 0.05\,\mathrm{s\,km^{-1}}$, considered in this work.

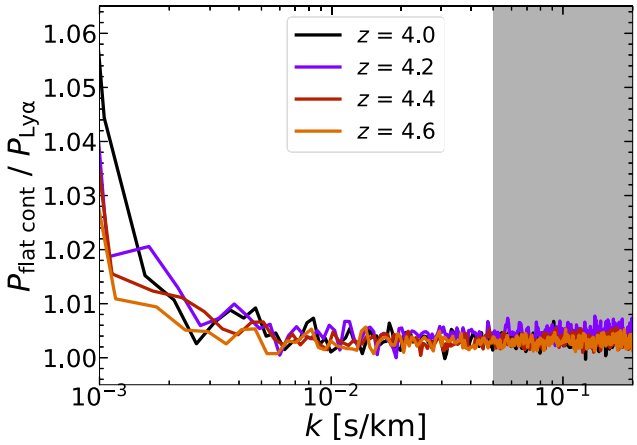



**Figure 8.** The ratio of the Ly $\alpha$ power spectrum measured after replacing the true continuum with a flat, wavelength-independent continuum, $P_{\mathrm{flat\,cont}}$, to the Ly $\alpha$ power spectrum, $P_{\mathrm{Ly}\alpha}$, measured assuming perfect knowledge of the continuum. The grey shading shows the excluded region above the maximum wavenumber, $k_{\mathrm{max}} = 0.05\,\mathrm{s\,km^{-1}}$, considered in this work.

value of the true continuum over the redshift bin covered by each mock QSO spectrum. This approach isolates shape related biases from normalization effects, which are instead marginalized over by varying the effective optical depth, $\tau_{\mathrm{eff}}$, in our parameter inference procedure.

The impact of an incorrect continuum shape on $P_{\mathrm{Ly}\alpha}$ is shown in Fig. 8. The associated systematic error is again added in quadrature to the total uncertainty estimate. Note, however, that this test does not model a specific continuum-fitting procedure. Underfitting or overfitting may introduce a systematic excess or suppression of large-scale power that is not explicitly accounted for here and will ideally need to be assessed for the specific continuum-fitting method adopted in the WEAVE-QSO survey. We also remark that if a smaller observational sample were obtained by selecting brighter QSOs, we might expect the associated improvements in continuum fitting accuracy (and also HCD identification) to partly offset the increase in statistical uncertainty from the reduced sample size. However, assessing this trade-off requires modelling the performance of these procedures as a function of data quality, which is beyond the scope of the present forecast.

### 3.7 Final uncertainty estimate

We now combine our systematic uncertainty estimates with the expected statistical uncertainty from the WEAVE-QSO survey. Statistical uncertainties originate from finite sample sizes and are estimated relative to the mean $P_{\mathrm{Ly}\alpha}$. We use the predicted number of QSOs (denoted $N_{\mathrm{QSO}}$) expected in a 10 000 deg$^2$ survey after applying an $r$-band cut of $m_r < 21.9$, as described in Section 2.4. For each of the four redshift bins considered here, we then generate 1000 bootstrap realizations of the mean $P_{\mathrm{Ly}\alpha}$ by randomly sampling $N_{\mathrm{QSO}}$ sightlines with replacement. The statistical errors are derived as the square root of the diagonal elements of the resulting covariance matrix.

Next, for systematics associated with spectral resolution, noise subtraction, and uncorrelated metals, we use the mock spectra to construct bootstrap estimates of the residual uncertainty associated with each correction. The corresponding quantities shown in Fig. 9 are the standard deviations of the corresponding correction ratios across 1000 bootstrap realizations of $N_{\mathrm{QSO}}$ randomly sampled sightlines. For HCDs and continuum placement, we estimate the uncertainty from the fractional change in $P_{\mathrm{Ly}\alpha}$. The HCD contribution is modelled using the ratio between power spectra with incomplete HCD removal and the pure Ly $\alpha$ spectrum (solid curves in Fig. 7). The continuum placement contribution is evaluated analogously (solid curves in Fig. 8). The total systematic uncertainty is computed as the quadratic sum of all individual contributions. We note that the forecast uncertainty estimate assumes perfect knowledge of the spectral-resolution window and of the noise model. In real WEAVE-QSO data, uncertainties in both quantities will need to be tested carefully. We have tested the effect of including a 10 per cent uncertainty in $R$, finding that the resulting fractional systematic uncertainty is roughly 25–50 per cent near $k_{\mathrm{max}}$, much larger than the current bootstrap resolution term and also larger than the noise term near $k_{\mathrm{max}}$. However, this contribution remains below 1 per cent at $\log_{10}(k/\mathrm{s\,km^{-1}}) \lesssim -2.0$ [and is $\sim$0.01 per cent at $\log_{10}(k/\mathrm{s\,km^{-1}}) \sim -3.0$] and therefore does not affect the large-scale modes that dominate the patchy reionization detection forecast presented in this paper.

Fig. 9 shows the resulting estimate for the $1\sigma$ relative uncertainties. The statistical and systematic contributions are indicated by dotted and dashed lines, respectively, while the solid lines represent their quadratic sum. As before, the grey shaded region shows the maximum wavenumber for the $P_{\mathrm{Ly}\alpha}$ measurements that we consider in this work, $k_{\mathrm{max}} = 0.05\,\mathrm{s\,km^{-1}}$. The total systematic uncertainty at large $k$ reflects the dominance of the noise subtraction term on small scales, whereas the increase in the statistical error with redshift is due to the smaller number of QSOs in the higher redshift bins. At large scales ($k \lesssim 0.01\,\mathrm{km\,s^{-1}}$), the total uncertainties remain below 10 per cent for all redshift bins. The average uncertainty at $k < k_{\mathrm{max}}$ over all four redshift bins is 5.2 per cent. Over the range $-2.6 \lesssim \log_{10}(k/\mathrm{s\,km^{-1}}) \lesssim -1.7$, the relative uncertainties on $P_{\mathrm{Ly}\alpha}$ are dominated by the statistical contribution (i.e. where the dotted line is above the dashed line in Fig. 9). Within this range, we forecast an average relative uncer-

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

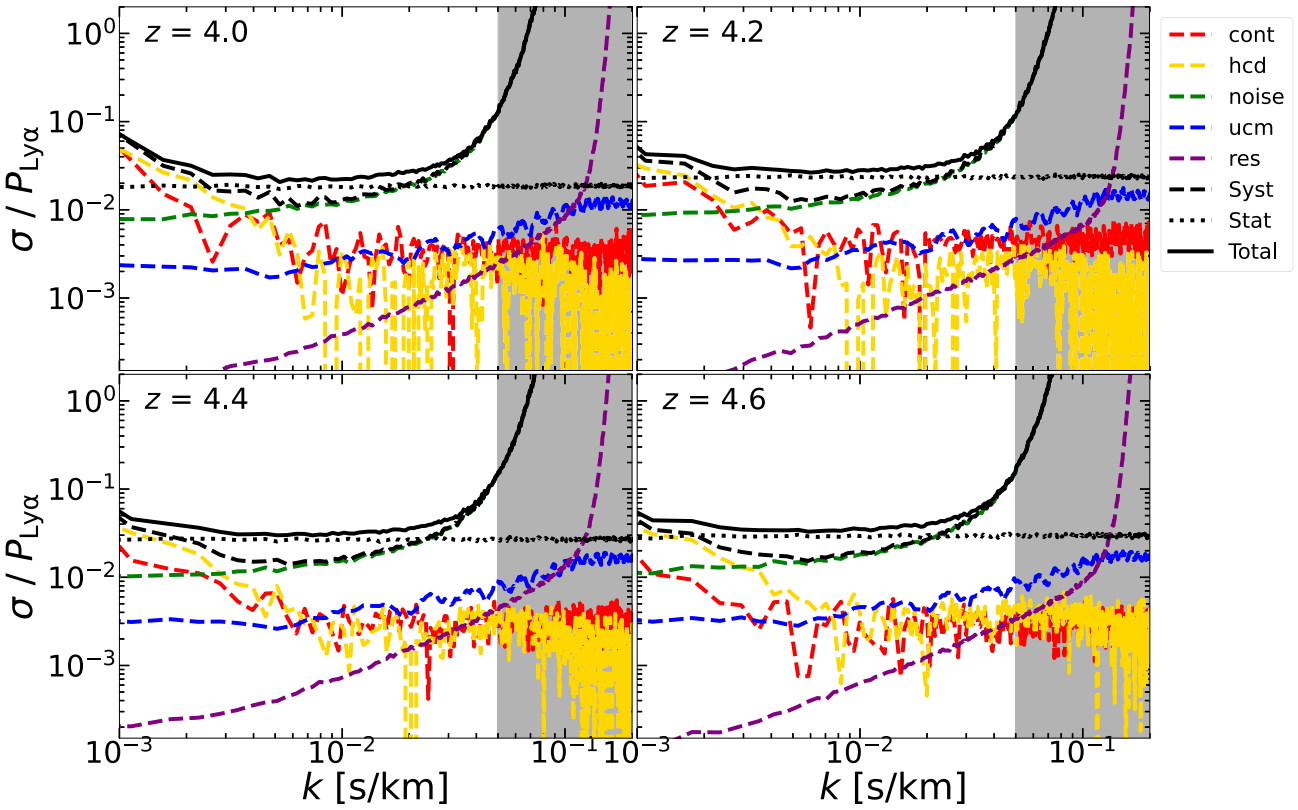



**Figure 9.** Forecast for the total relative uncertainty (solid black curves) on a WEAVE-QSO measurement of the 1D Ly $\alpha$ forest power spectrum at $z = 4.0$–4.6. Each panel represents a different bin of width $\Delta z = 0.2$ centred at the redshift indicated. The statistical errors (dotted black curves) are calculated by bootstrapping with replacement. The systematic uncertainty contributions are modelled separately for each observational or astrophysical effect and are added together in quadrature (black dashed curve). The red, yellow, green, blue, and purple dashed curves represent the contributions from continuum placement, incomplete HCD identification, noise subtraction, uncorrelated-metal contamination, and spectral resolution, respectively. The total uncertainty is obtained from the quadratic sum of the statistical and systematic uncertainty. The grey shading shows the excluded region above the maximum wavenumber, $k_{\rm max} = 0.05\,{\rm s\,km^{-1}}$, considered in this work. For $k < k_{\rm max}$, the relative errors at large scales are always within 10 per cent.

tainty of 3.1 per cent. These estimates are used to construct the covariance matrix for a WEAVE-QSO $P_{{\rm Ly}\alpha}$ measurement, which we now use as an input within our Markov chain Monte Carlo (MCMC) parameter inference framework.

## 4 THE IMPRINT OF PATCHY REIONIZATION AT $z > 4$

### 4.1 Parameter inference framework

We now use our mock 1D Ly $\alpha$ power spectrum covariance matrix to test whether a WEAVE-QSO measurement can recover the large-scale imprint of patchy reionization. Our MCMC parameter inference framework follows the approach used in earlier work, where the measured $P_{{\rm Ly}\alpha}$ is compared against a grid of simulated power spectra while marginalizing over astrophysical parameters (e.g. M. Viel et al. 2013; V. Iršič et al. 2017b; M. Molaro et al. 2023). For each redshift bin, our Ly $\alpha$ forest model is described by the effective optical depth $\tau_{\rm eff}$, the temperature at mean density $T_0$, the slope of the temperature–density relation $\gamma$, and the cumulative thermal energy per unit mass $u_0$ (F. Nasir, J. S. Bolton & G. D. Becker 2016). These parameters control the mean transmission and thermal state of the IGM, and hence the amplitude and shape of the intrinsic $P_{{\rm Ly}\alpha}$. We do not vary any cosmological parameters in this analysis (see M. Molaro et al. 2023, for further discussion of this point).

In addition, we include an extra parameter that describes the effect of patchy reionization on the Ly $\alpha$ forest at $z > 4$. For this purpose we adopt the patchy reionization correction used by M. Molaro et al. (2023), based on the tabulated data presented in table 2 of M. Molaro et al. (2022). These corrections were calculated using the Sherwood-Relics simulation suite (E. Puchwein et al. 2023), by taking the ratio of $P_{{\rm Ly}\alpha}$ obtained from hybrid radiative transfer simulations of patchy reionization to models using a spatially uniform UV background with *the same volume averaged reionization history* as the patchy model. On the scales relevant for this work the correction enhances the large-scale power at $\log_{10}(k/{\rm s\,km^{-1}}) \lesssim -1.5$, primarily due to spatial variations in the IGM ionization and thermal state that persist following the completion of reionization (R. Cen et al. 2009; A. D'Aloisio et al. 2015; M. Molaro et al. 2022).

Following M. Molaro et al. (2023), we parametrize this patchy reionization correction using the dimensionless parameter, $A_{\rm p}$, where

$$A_{\rm p} = \frac{z - 3.45}{2.55}, \tag{7}$$

and $z$ is the redshift at which the correction is linearly interpolated from table 2 of M. Molaro et al. (2022). We allow $A_{\rm p}$ to vary

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

over the range $0 \leq A_{\rm p} \leq 1$ by linearly extrapolating to a lower and upper redshift limit of $z = 3.45$ (i.e. for no correction) and $z = 6$, respectively. Note that $A_{\rm p}$ is used only as a numerical parameter controlling the amplitude of the patchy correction and should not be interpreted as a measurement of the redshift of reionization. In our MCMC analysis, we allow one independent $A_{\rm p}$ parameter in each redshift bin.

We compare two classes of fit to the same mock data set, which we construct by applying our covariance matrix estimate to $P_{{\rm Ly}\alpha}$ from the Sherwood-Relics model used in this work. To facilitate a direct comparison to the earlier work of M. Molaro et al. (2023), throughout we adopt the wavenumber binning introduced by N. G. Karaçaylı et al. (2022) and add a Gaussian noise realization to $P_{{\rm Ly}\alpha}$ drawn from the forecast covariance matrix. The first fit includes $A_{\rm p}$ and therefore allows the large-scale patchy reionization correction to vary in each $P_{{\rm Ly}\alpha}$ redshift bin. We add this to our simulated mocks such that $A_{\rm p} = 0.216,\ 0.294,\ 0.373,$ and 0.451 in the redshift bins at $z = 4.0,\ 4.2,\ 4.4,$ and 4.6, respectively. The second fit removes this degree of freedom by disabling the patchy reionization correction. In both cases, we apply a broad top-hat prior on $\tau_{\rm eff}$ spanning 0.4–1.6 times the reference effective optical depth, following E. Boera et al. (2019). The thermal parameters $T_0$, $\gamma$, and $u_0$ are assigned Gaussian priors centred on the reference thermal history of the Sherwood-Relics simulation used here, with a 10 per cent width in each redshift bin. We additionally impose a hard $T_0$–$u_0$ consistency prior, analogous to the simulation-grid prior used by M. Molaro et al. (2023), to exclude thermal histories outside the modelled region. As both fits are applied to the same mock data the difference in the goodness of fit provides a test of the significance that a WEAVE-QSO $P_{{\rm Ly}\alpha}$ measurement would favour an enhancement of large-scale power from reionization.

Finally, we also include the impact of correlated Si III absorption following K. Ma et al. (2026). As shown in Fig. 6, the correlated Si III contribution has an oscillatory dependence on wavenumber. For the relatively sparse large-scale $k$-binning that we adopt here, however, this oscillatory structure is always weakly sampled. Consequently, we find negligible bias in the recovered $A_{\rm p}$ posteriors when the Si III absorption model is included in our inference framework. We therefore do not consider this issue further.

### 4.2 Recovery of the large-scale power enhancement

Fig. 10 shows the mock $P_{{\rm Ly}\alpha}$ data in the four redshift bins considered in this work, together with the best-fitting model predictions obtained from the MCMC analysis. Note again the uncertainties are derived from the uncertainty estimate summarized in Section 3.7. This comparison is performed using the same mock data set and covariance matrix for all fits, such that any difference in the likelihood is driven only by the model assumptions. The red solid curves show the fiducial fit in which one independent $A_{\rm p}$ parameter is varied in each redshift bin. The blue dashed curves show the best-fitting model excluding the $A_{\rm p}$ parameter while all other IGM parameters are allowed to vary as before (the posterior distributions for all parameters are provided in Appendix A).

For the combined fit across all four redshift bins, the fiducial model provides an acceptable fit to the mock, with $\chi^2/{\rm dof} = 36.07/28$. By contrast, disabling the patchy correction degrades the fit to $\chi^2/{\rm dof} = 65.28/32$. The improvement is $\Delta\chi^2 = 29.21$ for four additional patchy reionization amplitude parameters. Interpreting this likelihood ratio statistic using Wilks' theorem gives $p \simeq 7.1 \times 10^{-6}$, corresponding to a two-sided Gaussian-equivalent significance of $\simeq 4.5\sigma$ for the fiducial patchy reionization model. Thus, even after allowing the no-$A_{\rm p}$ fit to adjust the thermal and ionization parameters the mock data strongly prefer the additional large-scale correction associated with patchy reionization. This demonstrates the patchy reionization imprint inserted into the mock is not fully absorbed by variations in $\tau_{\rm eff}$, $T_0$, $\gamma$, or $u_0$, and remains significant even after a detailed accounting of WEAVE-QSO survey uncertainties.

The posterior distributions for $A_{\rm p}$ are displayed in Fig. 11 and Table 1. In addition to the fiducial fit (red curves) here we consider two further cases. First, we repeat the analysis with broader thermal priors to test whether the inferred patchy reionization amplitudes are driven by restrictive assumptions about the IGM thermal state (orange curves). In this relaxed-prior run, we replace the fiducial Gaussian priors on $T_0$, $\gamma$, and $u_0$ with broad top-hat priors over the simulation grid ranges, while keeping the same $\tau_{\rm eff}$ prior and the same $T_0$–$u_0$ consistency cut. Specifically, we use $T_0 \in [0.5, 1.5] \times 10^4$ K and $\gamma \in [1.0, 1.7]$, with $u_0$ allowed to vary over the corresponding simulation-grid range in each redshift bin. Second, we remove the lowest-$k$ bin at $k \sim 10^{-3}\,{\rm s\,km^{-1}}$ from each redshift bin to test the sensitivity of the $A_{\rm p}$ recovery to the largest scale mode included in the fit (green curves).

The fiducial analysis recovers the expected patchy amplitudes in all bins. In each case, the expected value lies within the 68 per cent credible interval of the posterior distribution. Relaxing the thermal parameter priors leads to minor shifts and slightly broader posteriors, but the expected $A_{\rm p}$ values remain consistent with the inferred distributions. This indicates that the recovery is not simply imposed by restrictive priors on $T_0$, $\gamma$, or the integrated thermal history. However, removing the lowest-$k$ bin from each redshift bin weakens the constraints. This is expected because the patchy reionization correction is largest on this scale. The posteriors broaden by factors of approximately 1.3–2.4 relative to the fiducial case. In the corresponding no-lowest-$k$ bin model comparison, the improvement from including $A_{\rm p}$ is reduced to $\Delta\chi^2 = 5.40$ for four additional parameters, corresponding to a two-sided Gaussian-equivalent preference of only $\simeq 1.2\sigma$. In this case, there is no longer a preference for large-scale power enhancement from patchy reionization. Nevertheless, the expected $A_{\rm p}$ values remain well recovered with the exception of the $z = 4.4$, where only an upper limit is obtained. This demonstrates that most of the constraining power on $A_{\rm p}$ comes from retaining modes with $k < 0.0033\,{\rm s\,km^{-1}}$. We also emphasize that the forecast uncertainties in the lowest-$k$ bins are dominated by the systematic uncertainty associated with damping wings from high column density absorbers and to a lesser extent the continuum placement. We have verified this remains the case even if lowering $N_{\rm QSO}$ in each redshift bin by a factor of $\sim$2 from our fiducial values. This implies that the WEAVE-QSO survey configuration of approximately 10 000 deg$^2$ with $m_r < 22$ should comfortably exceed the minimum requirement on QSO sample size.

In summary, our results indicate a WEAVE-QSO $P_{{\rm Ly}\alpha}$ measurement at $z \geq 4$ with the forecast covariance adopted here can recover the redshift dependent patchy reionization amplitude at a two-sided Gaussian-equivalent significance of $4.5\sigma$. This forecast significance is higher than the $2.7\sigma$ preference for additional large scale power reported by M. Molaro et al. (2023) from an analysis of the observed 1D power spectrum measurements presented by N. G. Karaçaylı et al. (2022). While our analysis is based on mock WEAVE-QSO data with a forecast covariance, this

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

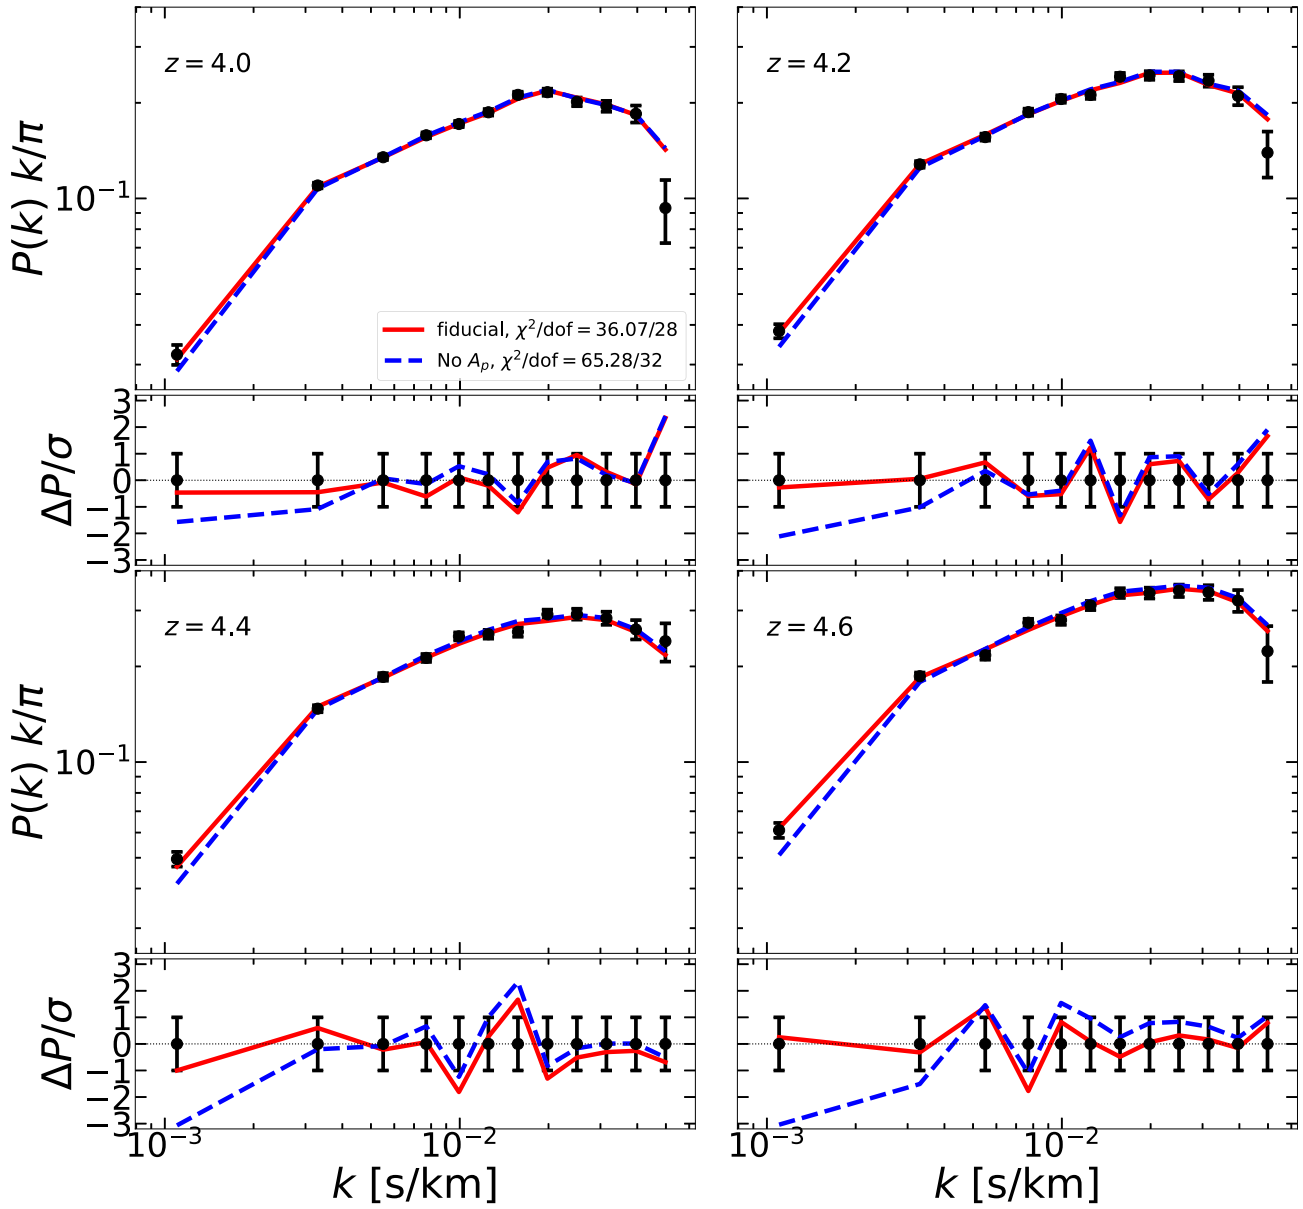



**Figure 10.** Comparison between the WEAVE-QSO mock $P_{\mathrm{Ly}\alpha}$ (black data points) and the best-fitting models from the MCMC parameter inference. The redshift bin is indicated in the upper left of each panel. The red solid curves show the fiducial fit, in which the patchy reionization parameter $A_{\mathrm{p}}$ is allowed to vary independently in each redshift bin. The blue dashed curves show the no-$A_{\mathrm{p}}$ comparison fit, where the patchy correction is disabled within the inference framework. The lower panels show the residuals relative to the mock data uncertainties, $(P_{\mathrm{fit}} - P_{\mathrm{data}})/\sigma_{\mathrm{data}}$. For the combined fit across all four redshift bins, allowing $A_{\mathrm{p}}$ to vary improves the fit from $\chi^2/\mathrm{dof} = 65.28/32$ to $36.07/28$, corresponding to $\Delta\chi^2 = 29.21$ for four additional parameters. This corresponds to a two-sided Gaussian-equivalent significance of $4.5\sigma$ for a large scale power enhancement due to patchy reionization.

comparison suggests that WEAVE-QSO should have sufficient statistical precision to turn the tentative evidence seen in current data into a significant detection of the large scale patchy reionization signature. However, the largest scale mode provides substantial constraining power: When the lowest-$k$ bin is removed, the model-comparison preference for including $A_{\mathrm{p}}$ is weakened to $\simeq 1.2\sigma$ and is no longer significant.

## 5 SUMMARY AND CONCLUSIONS

In this work, we have developed a framework to forecast the sensitivity of a massive spectroscopic survey to the large-scale imprint of patchy hydrogen reionization in the one-dimensional Ly $\alpha$ forest flux power spectrum, $P_{\mathrm{Ly}\alpha}$. We have opted to focus on a WEAVE-QSO survey configuration (M. M. Pieri et al. 2016; S. Jin et al. 2024), targeting four high-redshift bins of width $\Delta z = 0.2$ centred at $z = 4.0, 4.2, 4.4$, and 4.6, where relic temperature fluctuations from inhomogeneous reionization are expected to remain observable in $P_{\mathrm{Ly}\alpha}$ (X. Wu et al. 2019; M. Molaro et al. 2022, 2023). However, we anticipate that other current and forthcoming massive spectroscopic surveys that aim to constrain large-scale power in the Ly $\alpha$forest at wavenumbers $k \sim 10^{-3}\,\mathrm{s\,km^{-1}}$ at $z > 4$ should also be capable of this novel measurement.

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

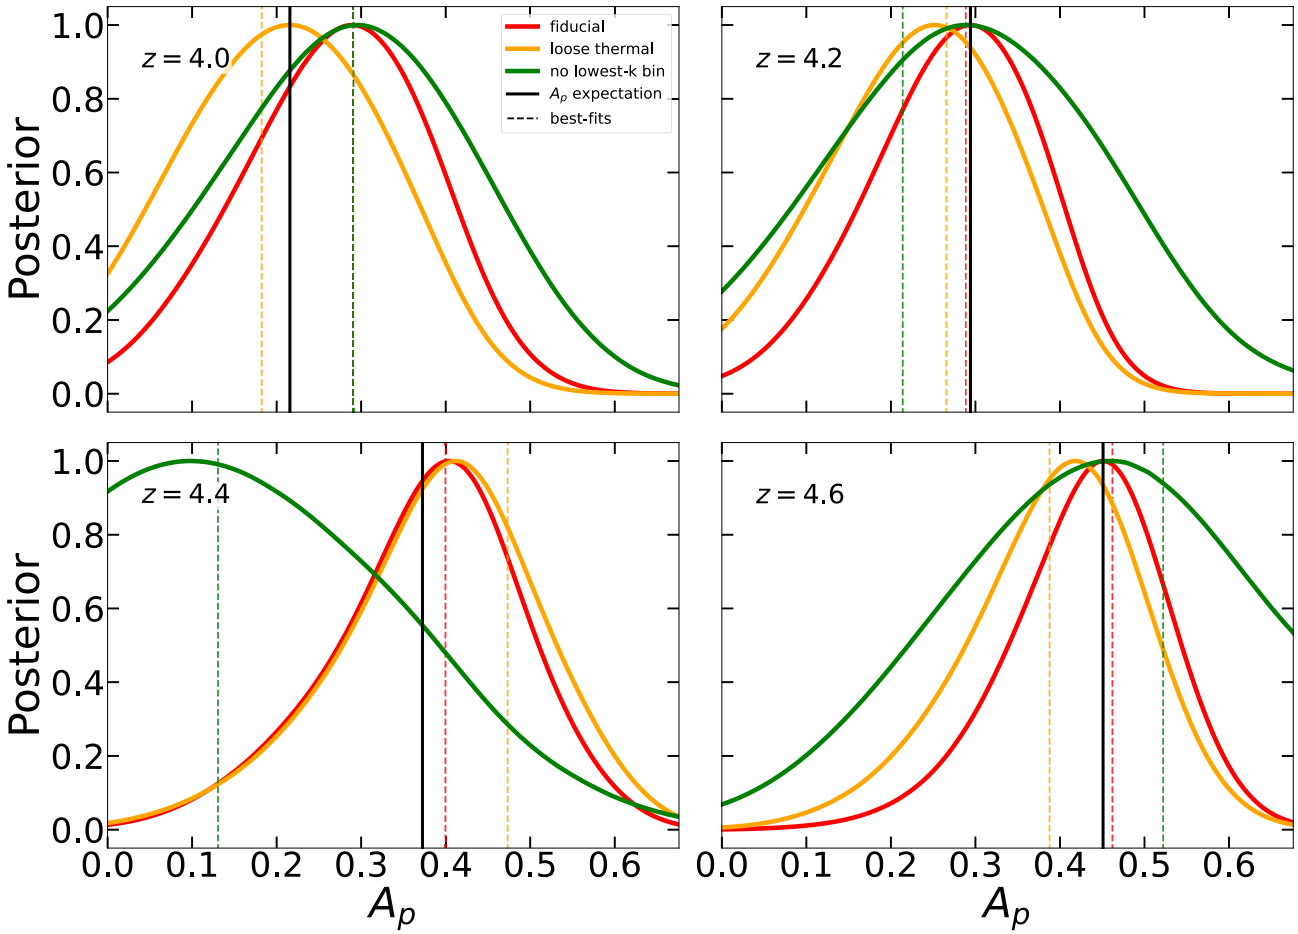



**Figure 11.** One dimensional posterior distributions for the patchy reionization parameter $A_p$ in the four $P_{Ly\alpha}$ redshift bins considered in this work. The vertical black lines show the expected values used in the mocks (see equation 7). The dashed vertical lines show the best-fitting values for each run. The fiducial analysis (red curves) recovers the amplitudes in all redshift bins. Relaxing the thermal-history priors (yellow curves) produces slightly broader posteriors but does not break the recovery. However, removing the lowest-$k$ bin from each redshift bin further broadens the constraints and leads to only an upper limit at $z = 4.4$. This confirms the importance of the largest-scale modes for detecting the enhanced large-scale power from patchy reionization.

**Table 1.** Recovered constraints on the patchy reionization amplitude parameter, $A_p$, for the three cases displayed in Fig. 11. The expected values are the amplitudes used in the mock data vector. The posterior peak column reports the marginalized one-dimensional posterior peak with the 68 per cent credible interval or one-sided upper limit, while the final column reports the best-fitting value.

| Run | $z$ | Expected $A_p$ | Posterior peak | Best fit |
|---|---|---|---|---|
| Fiducial | 4.0 | 0.216 | $0.289^{+0.105}_{-0.121}$ | 0.291 |
| | 4.2 | 0.294 | $0.294^{+0.094}_{-0.108}$ | 0.289 |
| | 4.4 | 0.373 | $0.404^{+0.100}_{-0.121}$ | 0.400 |
| | 4.6 | 0.451 | $0.451^{+0.084}_{-0.099}$ | 0.462 |
| Loose thermal | 4.0 | 0.216 | $0.216^{+0.116}_{-0.124}$ | 0.182 |
| | 4.2 | 0.294 | $0.252^{+0.104}_{-0.118}$ | 0.266 |
| | 4.4 | 0.373 | $0.412^{+0.111}_{-0.123}$ | 0.473 |
| | 4.6 | 0.451 | $0.418^{+0.093}_{-0.115}$ | 0.388 |
| No lowest-$k$ bin | 4.0 | 0.216 | $0.295^{+0.142}_{-0.151}$ | 0.291 |
| | 4.2 | 0.294 | $0.289^{+0.164}_{-0.158}$ | 0.214 |
| | 4.4 | 0.373 | $< 0.387$ | 0.131 |
| | 4.6 | 0.451 | $0.460^{+0.212}_{-0.223}$ | 0.522 |

Our mock spectra are based on the Sherwood-Relics simulation suite (E. Puchwein et al. 2023) and are processed with the WEAVEify pipeline to include the main observational ingredients relevant for WEAVE-QSO spectra. We construct a realistic high-redshift QSO population using luminosity-function-based Monte Carlo sampling (N. Palanque-Delabrouille et al. 2016) and adopt an $r$-band magnitude cut of $m_r < 21.9$. For a 10 000 deg$^2$ survey area, this selection yields expected QSO numbers of 2942, 1822, 1385, and 1155 in the four redshift bins considered here.

We use our pipeline to quantify the overall uncertainty for a WEAVE-QSO $P_{Ly\alpha}$ measurement. The covariance matrix estimate includes statistical errors and systematic contributions from spectral resolution, noise subtraction, continuum misplacement, uncorrelated metal contamination, and the incomplete identification of high column density absorbers with $N_{HI} > 10^{19}$ cm$^{-2}$. Within our adopted maximum wavenumber cut, $k < 0.05$ s km$^{-1}$, the estimated relative uncertainties remain below 10 per cent on the largest scales in all redshift bins. Averaged over all bins and scales below this cut, the total uncertainty is 5.2 per cent. Over the range $-2.6 \lesssim \log_{10}(k/\mathrm{s\,km^{-1}}) \lesssim -1.7$, the uncertainties are dominated by the statistical contribution and the average uncertainty is 3.1 per cent.

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

Using this forecast covariance, we construct a mock $P_{\mathrm{Ly}\alpha}$ data vector containing the fiducial patchy reionization signal and test whether it can be distinguished from an IGM that has been heated and ionized by a spatially uniform UV background. This inference test follows the patchy reionization parametrization and wavenumber binning used in earlier work (M. Molaro et al. 2023). Allowing one independent patchy reionization amplitude parameter, $A_{\mathrm{p}}$, in each redshift bin improves the fit from $\chi^2/\mathrm{dof} = 65.28/32$ to $\chi^2/\mathrm{dof} = 36.07/28$, corresponding to $\Delta\chi^2 = 29.21$ for four additional patchy amplitude parameters and a $\simeq 4.5\sigma$ preference for additional power at large scales from patchy reionization. This result remains robust to changes in the prior adopted for IGM thermal parameters and is driven almost entirely by the lowest-$k$ bin from each redshift bin.

Our results indicate that a WEAVE-QSO measurement of the high-redshift Ly $\alpha$ forest power spectrum should be capable of detecting and characterizing the large-scale relic imprint of patchy reionization, provided that the observational covariance is controlled at the level forecast here. More broadly, the framework developed in this work provides a route for assessing the sensitivity of current and future massive spectroscopic surveys to the timing of reionization through large-scale $P_{\mathrm{Ly}\alpha}$ measurements.

## ACKNOWLEDGEMENTS

KM gratefully acknowledges the financial support provided by the China Scholarship Council programme (Project ID: CSC NO. 202309110040). JSB was supported by Science and Technology Facilities Council (STFC) consolidated grant ST/X000982/1. The modelling in this work was performed using the Joliot Curie supercomputer at the Trés Grand Centre de Calcul (TGCC) and the DiRAC Data Intensive service (CSD3) at the University of Cambridge, managed by the University of Cambridge University Information Services on behalf of the STFC DiRAC HPC facility (www.dirac.ac.uk). We acknowledge the Partnership for Advanced Computing in Europe (PRACE) for awarding us time on Joliot Curie in the 16th call. The DiRAC component of CSD3 at Cambridge was funded by Department for Business, Energy and Industrial Strategy (BEIS), UK Research and Innovation (UKRI), and STFC capital funding and STFC operations grants. DiRAC is part of the National e-Infrastructure. We thank Volker Springel for making P-GADGET-3 available and Philip Parry for technical support.

## DATA AVAILABILITY

All data and analysis code used in this work are available from the first author on reasonable request. Further guidance on accessing the publicly available Sherwood-Relics simulation data may also be found at https://www.nottingham.ac.uk/astronomy/sherwood-relics/.

## REFERENCES

Aguirre A., Schaye J., Kim T.-S., Theuns T., Rauch M., Sargent W. L. W., 2004, ApJ, 602, 38
Becker G. D., Hewett P. C., Worseck G., Prochaska J. X., 2013, MNRAS, 430, 2067
Boera E., Becker G. D., Bolton J. S., Nasir F., 2019, ApJ, 872, 101
Bolton J. S., Becker G. D., 2009, MNRAS, 398, L26
Bolton J. S., Haehnelt M. G., Viel M., Springel V., 2005, MNRAS, 357, 1178
Bolton J. S., Puchwein E., Sijacki D., Haehnelt M. G., Kim T.-S., Meiksin A., Regan J. A., Viel M., 2017, MNRAS, 464, 897
Cain C., Scannapieco E., McQuinn M., D'Aloisio A., Trac H., 2024, MNRAS, 533, L100
Carswell R. F., Whelan J. A. J., Smith M. G., Boksenberg A., Tytler D., 1982, MNRAS, 198, 91
Cen R., McDonald P., Trac H., Loeb A., 2009, ApJ, 706, L164
Chabanier S. et al., 2019, J. Cosmol. Astropart. Phys., 2019, 017
Chatzikos M. et al., 2023, Rev. Mex. Astron. Astrofis., 59, 327
Chaves-Montero J. et al., 2026, J. Cosmol. Astropart. Phys., 2026, 040
Croft R. A. C., Weinberg D. H., Bolte M., Burles S., Hernquist L., Katz N., Kirkman D., Tytler D., 2002, ApJ, 581, 20
D'Aloisio A., McQuinn M., Trac H., 2015, ApJ, 813, L38
DESI Collaboration 2016, preprint (arXiv:1611.00036)
D'Odorico V. et al., 2013, MNRAS, 435, 1198
D'Odorico V. et al., 2022, MNRAS, 512, 2389
Davies F. B. et al. 2018, ApJ, 864, 143.
Ferland G. J., Korista K. T., Verner D. A., Ferguson J. W., Kingdon J. B., Verner E. M., 1998, PASP, 110, 761
Gaikwad P., Srianand R., Haehnelt M. G., Choudhury T. R., 2021, MNRAS, 506, 4389
Garcia-Gallego O., Iršič V., Haehnelt M. G., Viel M., Bolton J. S., 2025, Phys. Rev. D, 112, 043502
Garcia-Gallego O., Iršič V., Viel M., Haehnelt M. G., Bolton J. S., 2026, preprint (arXiv:2603.04401)
Greene J., Bezanson R., Ouchi M., Silverman J., the PFS Galaxy Evolution Working Group, 2022, preprint (arXiv:2206.14908)
Hirata C. M., 2018, MNRAS, 474, 2173
Hooper D. C., Schöneberg N., Murgia R., Archidiacono M., Lesgourgues J., Viel M., 2022, J. Cosmol. Astropart. Phys., 2022, 032
Hui L., Haiman Z., 2003, ApJ, 596, 9
Hui L., Rutledge R. E., 1999, ApJ, 517, 541
Iršič V. et al., 2017a, Phys. Rev. D, 96, 023522
Iršič V. et al., 2017b, MNRAS, 466, 4332
Iršič V. et al., 2024, Phys. Rev. D, 109, 043511
Jin S. et al., 2024, MNRAS, 530, 2688
de Jong R. S. et al., 2019, The Messenger, 175, 3
Karaçaylı N. G., Font-Ribera A., Padmanabhan N., 2020, MNRAS, 497, 4742
Karaçaylı N. G. et al., 2022, MNRAS, 509, 2842
Karaçaylı N. G. et al., 2024, MNRAS, 528, 3941
Karaçaylı N. G. et al. 2025, JCAP, 10, 004.
Keating L. C., Puchwein E., Haehnelt M. G., 2018, MNRAS, 477, 5501
Kim T. S., Bolton J. S., Viel M., Haehnelt M. G., Carswell R. F., 2007, MNRAS, 382, 1657
López S. et al., 2016, A&A, 594, A91
Ma K., Bolton J. S., Iršič V., Gaikwad P., Puchwein E., 2026, MNRAS, 546, staf2262
Mainieri V. et al., 2024, preprint (arXiv:2403.05398)
Mathes N. L., Churchill C. W., Murphy M. T., 2017, preprint (arXiv:1701.05624)
McDonald P. et al., 2005, ApJ, 635, 761
McDonald P. et al., 2006, ApJS, 163, 80
McQuinn M., 2016, ARA&A, 54, 313
Meiksin A. A., 2009, Rev. Mod. Phys., 81, 1405
Miller J. S. A., Bolton J. S., Hatch N., 2019, MNRAS, 489, 5381
Molaro M. et al., 2022, MNRAS, 509, 6119
Molaro M., Iršič V., Bolton J. S., Lieu M., Keating L. C., Puchwein E., Haehnelt M. G., Viel M., 2023, MNRAS, 521, 1489
Montero-Camacho P., Mao Y., 2021, MNRAS, 508, 1262
Murdoch H. S., Hunstead R. W., Pettini M., Blades J. C., 1986, ApJ, 309, 19
Murphy M. T., Kacprzak G. G., Savorgnan G. A. D., Carswell R. F., 2019, MNRAS, 482, 3458
Nasir F., Bolton J. S., Becker G. D., 2016, MNRAS, 463, 2335
O'Meara J. M. et al., 2015, AJ, 150, 111
Palanque-Delabrouille N. et al., 2013, A&A, 559, A85
Palanque-Delabrouille N. et al., 2016, A&A, 587, A41

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

Pâris I. et al., 2011, A&A, 530, A50
Pavičević M., Iršič V., Viel M., Bolton J. S., Haehnelt M. G., Martin-Alvarez S., Puchwein E., Ralegankar P., 2025, Phys. Rev. Lett., 135, 071001
Pieri M. M. et al.. 2016, in Reylé C., Richard J., Cambrésy L., Deleuil M., Pécontal E., Tresse L., eds, Vauglin I., SF2A-2016: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics. Société Française d'Astronomie et d'Astrophysique (SF2A), Paris, France, 259.
Pistis F. et al., 2025, A&A, 698, A292
Prochaska J. X., Madau P., O'Meara J. M., Fumagalli M., 2014, MNRAS, 438, 476
Puchwein E. et al., 2023, MNRAS, 519, 6162
Rauch M. et al., 1997, ApJ, 489, 7
Ravoux C. et al., 2023, MNRAS, 526, 5118
Rogers K. K., Poulin V., 2025, Phys. Rev. Res., 7, L012018
Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, MNRAS, 318, 817
Schaye J., Aguirre A., Kim T.-S., Theuns T., Rauch M., Sargent W. L. W., 2003, ApJ, 596, 768
Springel V., 2005, MNRAS, 364, 1105
Szakacs R. et al., 2023, MNRAS, 526, 3744
Theuns T., Schaye J., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell B., 2002, ApJ, 567, L103
Viel M., Haehnelt M. G., Springel V., 2004, MNRAS, 354, 684
Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, Phys. Rev. D, 88, 043502
Villasenor B., Robertson B., Madau P., Schneider E., 2023, Phys. Rev. D, 108, 023502
Wang B. et al., 2022, ApJS, 259, 28
Wu X., McQuinn M., Kannan R., D'Aloisio A., Bird S., Marinacci F., Davé R., Hernquist L., 2019, MNRAS, 490, 3177
Zhao C. et al. 2026, Science China Physics, Mechanics & Astronomy, 69, 119511.
Zheng Y., Montero-Camacho P., Cai Z., Mao Y., 2026, MNRAS, 545, staf2014

## APPENDIX A: POSTERIOR DISTRIBUTIONS

Figs A1 and A2 show the marginalized posterior distributions for all five parameters ($\tau_{\rm eff}$, $T_0$, $\gamma$, $u_0$, $A_{\rm p}$) varied in each $P_{{\rm Ly}\,\alpha}$ redshift bin. Fig. A1 corresponds to the best fit $P_{{\rm Ly}\,\alpha}$ displayed in Fig. 10, while Fig. A2 corresponds to the 1D posteriors displayed in Fig. 11.

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

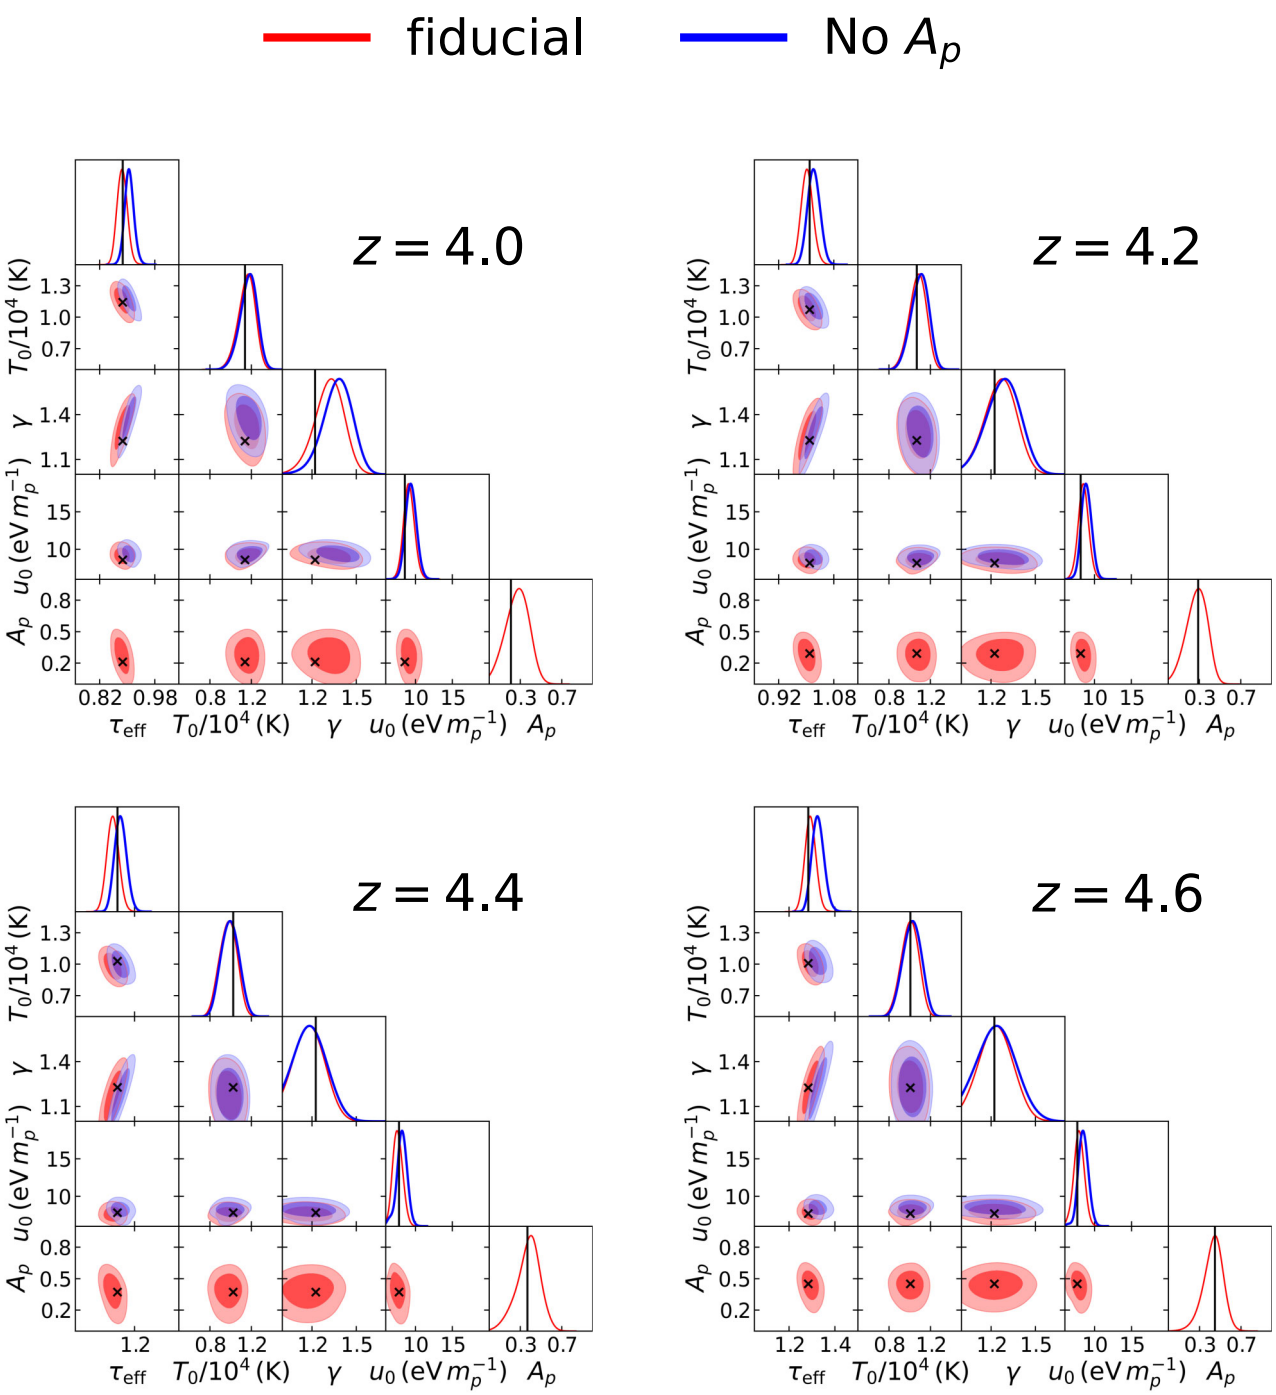



**Figure A1.** The 1D and 2D posterior distributions for $\tau_{\rm eff}$, $T_0$, $\gamma$, $u_0$, and $A_{\rm p}$ for our fiducial model (red contours) and no-$A_{\rm p}$ model (blue contours). The shaded regions show the 68 per cent and 95 per cent confidence intervals for the joint distributions. The four panels are for the redshift bins used in this study, namely $z = 4.0$, 4.2, 4.4, and 4.6. Expected parameter values are shown as black crosses in the contour panels and vertical black lines in the one-dimensional posterior panels. Note that only the fiducial posterior is shown in panels involving $A_{\rm p}$. The similarity of the posteriors in the two cases indicates the preference for patchy reionization is not driven by significant shifts in the other IGM parameters.

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

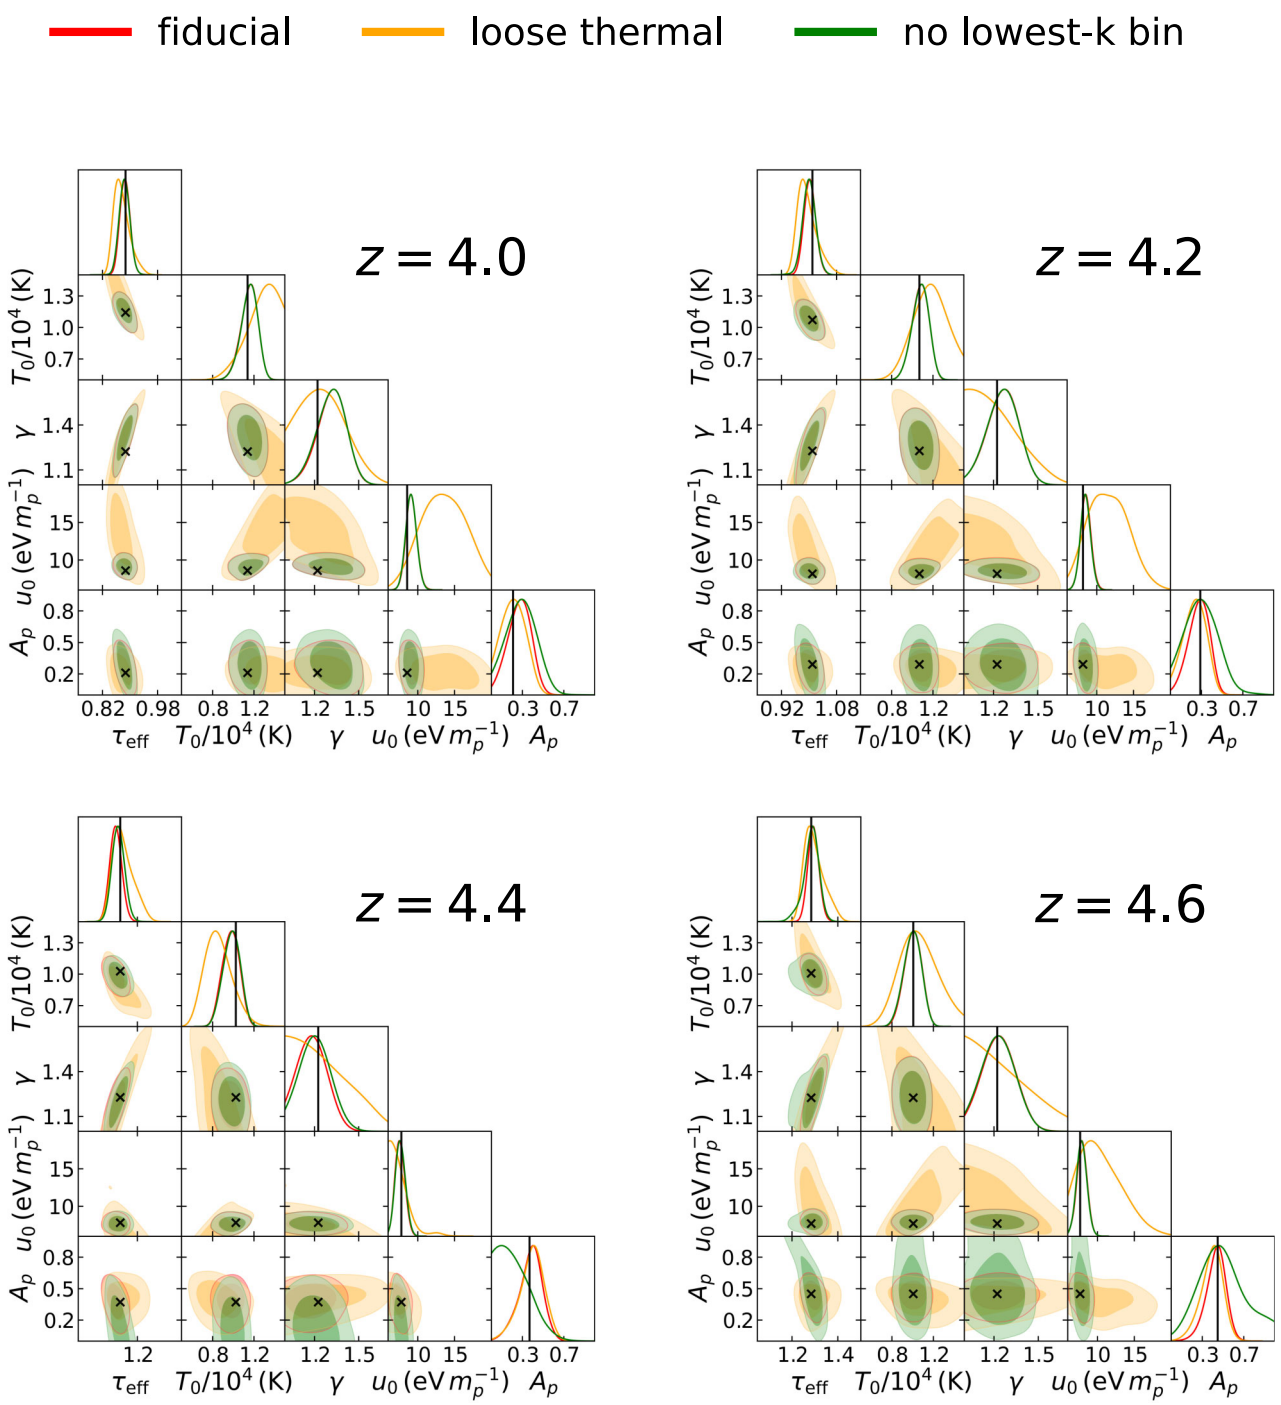



**Figure A2.** As for Fig. A1, but now showing our fiducial model (red contours), loose thermal prior model (orange contours), and no-lowest-$k$ model (green contours). The posterior for $A_{\rm p}$ does not change significantly with the thermal prior, although the constraint on $A_{\rm p}$ is significantly weakened when ignoring the largest-scale $k$ bin.

This paper has been typeset from a TEX/LATEX file prepared by the author.

Downloaded from academic.oup.com/mnras/article/553/3/stag2028/9004843 by guest on 26 November 2026

© The Author(s) 2026.
Published by Oxford University Press on behalf of Royal Astronomical Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.