Universiteit
Leiden
The Netherlands

Improving IGM temperature constraints using wavelet analysis on high-redshift quasars

Wolfson, M.; Hennawi, J.F.; Davies, F.B.; Oñorbe, J.; Hiss, H.; Lukić, Z.

Citation

Wolfson, M., Hennawi, J. F., Davies, F. B., Oñorbe, J., Hiss, H., & Lukić, Z. (2021). Improving IGM temperature constraints using wavelet analysis on high-redshift quasars. *Monthly Notices Of The Royal Astronomical Society*, 508(4), 5493-5513.
doi:10.1093/mnras/stab2920

Version: Submitted Manuscript (under Review)
License: [Leiden University Non-exclusive license](#)
Downloaded from: <https://hdl.handle.net/1887/3273920>

Note: To cite this publication please use the final published version (if applicable).

Improving IGM temperature constraints using wavelet analysis on high-redshift quasars

Molly Wolfson^{1*}, Joseph F. Hennawi^{1,2}, Frederick B. Davies^{1,3,4}, Jose Oñorbe⁵, Hector Hiss⁴, and Zarija Lukic³

¹*Department of Physics, University of California, Santa Barbara, CA 93106, USA*

²*Leiden Observatory, Leiden University, Niels Bohrweg 2, 2333 CA Leiden, Netherlands*

³*Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720, USA*

⁴*Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany*

⁵*Facultad de Físicas, Universidad de Sevilla, Avda. Reina Mercedes s/n, Campus de Reina Mercedes, 41012 Sevilla, Spain*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

The thermal state of the intergalactic medium (IGM) contains vital information about the epoch of reionization, one of the most transformative yet poorly understood periods in the young universe. This thermal state is encoded in the small-scale structure of Lyman- α (Ly α) absorption in quasar spectra. The 1D flux power spectrum measures the average small-scale structure along quasar sightlines. At high redshifts, where the opacity is large, averaging mixes high signal-to-noise ratio transmission spikes with noisy absorption troughs. Wavelet amplitudes are an alternate statistic that maintains spatial information while quantifying fluctuations at the same spatial frequencies as the power spectrum, giving them the potential to more sensitively measure the small-scale structure. Previous Ly α forest studies using wavelet amplitude probability density functions (PDFs) used limited spatial frequencies and neglected strong correlations between PDF bins and across wavelets scales, resulting in sub-optimal and unreliable parameter inference. Here we present a novel method for performing statistical inference using wavelet amplitude PDFs that spans the full range of spatial frequencies probed by the power spectrum and that fully accounts for these correlations. We applied this procedure to realistic mock data drawn from a simple thermal model parameterized by the temperature at mean density, T_0 , and find that wavelets deliver 1σ constraints on T_0 that are on average 7% more sensitive at $z = 5$ (12% at $z = 6$) than those from the power spectrum. We consider the possibility of combining wavelet PDFs with the power, but find that this does not lead to improved sensitivity.

Key words:

intergalactic medium – methods: statistical – quasars: absorption lines – dark ages, reionization, first stars

1 INTRODUCTION

The epoch of reionization, when the first luminous sources reionized the neutral hydrogen in the intergalactic medium (IGM), is one of the most dramatic periods of evolution in the young universe. During this time, ionization fronts impulsively heated reionized gas in the IGM to $\sim 10^4$ K. The exact amount of heat injected into the IGM depends on the properties of the luminous sources as well as the timing and duration of reionization (McQuinn 2012; Davies et al. 2016; D’Aloisio et al. 2019). After reionization, the IGM expands and cools through the adiabatic expansion of the universe and inverse Compton scattering off CMB photons. The combination of photoionization heating, Compton cooling, and cooling due to the expansion of the universe result in a tight power-law temperature-density relation for most of the IGM gas:

$$T = T_0 \Delta^{\gamma-1} \quad (1)$$

for overdensity $\Delta = \rho/\bar{\rho}$, the mean density of the Universe $\bar{\rho}$, temperature at mean density T_0 , and an expected slope γ (Hui & Gnedin 1997; Puchwein et al. 2015; McQuinn & Upton Sanderbeck 2016). However, the low-density IGM has long cooling times, so the thermal memory of reionization can persist for hundreds of Myr such that the thermal state of the IGM just after reionization ends contains important information on the state of the universe during reionization (Miralda-Escudé & Rees 1994; Hui & Gnedin 1997; Haehnelt & Steinmetz 1998; Theuns et al. 2002b; Hui & Haiman 2003; Lidz & Malloy 2014; Oñorbe et al. 2017; Oñorbe et al. 2017). Describing the thermal state of the IGM (T_0 and γ) just after reionization, $z \sim 5-6$, is key to understand the evolution of the universe during reionization.

The premier probe of the IGM is Ly α absorption along sightlines to bright quasars at high redshift, known as the Ly α forest (Gunn & Peterson 1965; Lynds 1971). The properties of these absorption features are sensitive to the thermal state of the IGM from two effects: Doppler broadening due to thermal motions and Jeans (pressure) smoothing of the underlying baryon distribution. The rate at which pressure forces erase gravitational fluctuations is set by the local sound speed, and at IGM densities the pressure scale sound cross-

* E-mail: mawolfson@ucsb.edu

ing time is approximately the Hubble time. Therefore, the pressure smoothing scale provides an integrated record of the thermal history of the IGM (Gnedin & Hui 1998; Kulkarni et al. 2015; Nasir et al. 2016; Oñorbe et al. 2017; Oñorbe et al. 2017; Rorai et al. 2017). Both of these effects reduce the small-scale structure of the Ly α forest.

Several statistics have been used to measure the thermal state of the IGM, including the flux probability density (Becker et al. 2007; Bolton et al. 2008; Viel et al. 2009; Calura et al. 2012; Lee et al. 2015), the curvature (Becker et al. 2011; Boera et al. 2014; Gaikwad et al. 2021), the Doppler parameter distribution (Schaye et al. 1999, 2000; Ricotti et al. 2000; Bryan & Machacek 2000; McDonald et al. 2001; Rudie et al. 2012; Bolton et al. 2010, 2012, 2014; Rorai et al. 2018; Gaikwad et al. 2021), and the joint distribution of the Doppler parameters with the Hydrogen Column Density (Hiss et al. 2018). One of the most commonly used statistics used to measure the structure of the Ly α forest is the 1D flux power spectrum ($P_F(k)$) (Theuns et al. 2000; Zaldarriaga et al. 2001; Yèche et al. 2017; Walther et al. 2017; Boera et al. 2019; Gaikwad et al. 2021). The reduction in small-scale structure in the Ly α forest leads to a cut-off in the power at high k values. However, with measurements of higher-redshift quasars, closer to reionization, the optical depth and its scatter for Ly α photons increase (Fan et al. 2006; Becker et al. 2015), leading to more absorption and Gunn-Peterson troughs in the Ly α forest. Calculating the 1D flux power spectrum at these high redshifts thus mixes high signal-to-noise ratio transmission spikes with noisy absorption troughs, potentially leading to a loss of information.

Wavelet analysis provides an alternative statistical method to measure the structure of the Ly α forest over a range of characteristic scales (Lidz et al. 2010; Garzilli et al. 2012; Gaikwad et al. 2021) (though see also Theuns & Zaroubi (2000); Theuns et al. (2002b); Zaldarriaga (2002); Meiksin (2000)). Wavelets are localized in both frequency and real space, which allows them to encode Fourier information while remaining in configuration space. Therefore, wavelet analysis has the benefit of keeping the absorption troughs distinct from the transmission spikes because it produces a full decomposition of wavelet amplitudes along the spectrum. The ultimate statistic used in wavelet analysis is the full wavelet amplitude probability density function (PDF). The PDF potentially contains more information than the average, which is effectively encoded in the power spectrum. However, these wavelet amplitude PDFs are complicated owing to the large correlations between bins in one wavelet amplitude PDF as well as between different wavelet amplitude PDFs.

Our work builds off and improves upon the previous implementation of wavelet analysis done by Lidz et al. (2010) and Gaikwad et al. (2021). The work done in Lidz et al. (2010) used one of the two characteristic wavelet scales explored to constrain the thermal state of the IGM. Each wavelet scale picks out a frequency in the flux so, to compare to the constraints on the thermal state of the IGM from $P_F(k)$, the number of smoothing scales used in wavelet analysis should be comparable to the number of band powers in $P_F(k)$. Only using one scale will reduce the constraining power of the wavelet amplitude PDFs because it is missing information in other Fourier modes. Lidz et al. (2010) also ignored correlations between the bins in the wavelet amplitude PDFs, potentially significantly affecting the resulting error bars. Gaikwad et al. (2021) used eight wavelet scales in their analysis and included the correlations between the bins within each wavelet amplitude PDF. Their method still ignored the correlations between the bins for wavelet amplitude PDFs of different scales, again potentially affecting the resulting error bars. They also combined their $P_F(k)$ measurements with their wavelet PDF measurements that were calculated from the same data set (along with the Doppler parameter distribution and the curvature statistic), ig-

norning correlations between all these statistics, to get a more precise measurement.

Our work quantifies the precision of parameter inference using wavelet amplitude PDFs and $P_F(k)$. We show that measuring T_0 from our simple thermal model from the wavelet amplitude PDFs results in a 7% reduction of the 1σ errors when compared to the measurement from $P_F(k)$ on the same mock data set. This confirms the potential that wavelet amplitude PDFs have to improve upon existing constraints on the thermal state of the IGM. Our wavelet analysis method uses more scales than previous works and spans the full range of scales probed by $P_F(k)$. For the first time, we calculate and present the full correlation matrices between the bins of the wavelet amplitude PDFs as well as the cross-correlations between $P_F(k)$ and the wavelet amplitude PDFs. We also combined the wavelet amplitude PDFs with $P_F(k)$ while taking the cross-correlations into account and found that this did not further improve the measurement. Finally, we characterized the effects of ignoring cross-correlations for the wavelet amplitude PDFs and the combination of the two statistics.

In addition to the thermal state of the IGM, the small-scale structure of the Ly α forest is also sensitive to departures from cold dark matter (CDM), including models of warm dark matter (WDM). For WDM, the linear power spectrum is exponentially suppressed when compared to CDM on scales smaller than the free-streaming length of the WDM particle (Narayanan et al. 2000). The mass of the WDM particle, m_{WDM} , can then be constrained by requiring the initial conditions to have sufficient small-scale power to reproduce the properties of the Ly α forest (Viel et al. 2013; Iršič et al. 2017; Garzilli et al. 2017). Wavelet analysis thus also has the potential to improve constraints on the mass of a WDM particle from the small-scale structure in the Ly α forest.

The structure of this paper is as follows. We describe our procedure for generating simulated Ly α forest sightlines in Section 2. We then introduce and explore the properties of our wavelet analysis in Section 3. Our method for statistical inference is laid out in Section 4. Our results comparing the measurements from the wavelet analysis and power spectrum is in Section 5. We summarize in Section 6.

2 SIMULATING LYMAN-ALPHA FOREST SPECTRA

2.1 Hydro Simulations

For this work we use one simulation run that uses the Nyx code. Nyx is a cosmological hydrodynamical simulation code designed for simulating the Ly α forest. For more details on the numerical methods, scaling, and the heating and cooling rates see Almgren et al. (2013) and Lukić et al. (2015). We use a standard Λ CDM cosmological model consistent with the constraints from Planck Collaboration et al. (2020): $\Omega_b = 0.04964$, $\Omega_m = 0.3192$, $\Omega_\Lambda = 0.6808$, $h = 0.67038$, $\sigma_8 = 0.826$, and $n_s = 0.9655$. The simulation we used has a box size of length, $L_{\text{box}} = 20 \text{ Mpc h}^{-1}$ and 1024^3 resolution elements. To simulate reionization, we use the *flash* model from Oñorbe et al. (2019) which reionizes at $z_{\text{reion}} = 7.75$, and uses $\Delta T = 2 \times 10^4$ to parameterize the instantaneous heat injection from reionization. In this framework every cell in the simulation will be ionized at $z = 7.75$ and heated to ΔT , unless the cell was previously ionized by a different process (i.e. collisional ionization). We consider two snapshots from this simulation at $z = 5$ and $z = 6$. We output 10,000 skewers of the Ly α forest from each snapshot to use in our analysis, which is equivalent of a total pathlength of 200 Gpc h^{-1} . The pixel scale of the simulation snapshot is $\Delta v = 2.7 \text{ km s}^{-1}$ at $z = 5$ and is $\Delta v = 2.9 \text{ km s}^{-1}$ at $z = 6$. Since there is a larger dataset available at

$z = 5$, we focus our work at this redshift. Figures shown in the main text will be at $z = 5$ unless otherwise specified while figures for $z = 6$ are available in Appendix C.

2.2 Thermal Models

In the *flash* reionization model, the majority of the IGM follows the tight temperature-density relation of equation (1) after reionization, see [Oñorbe et al. \(2019\)](#) for more details. In order to create simulation Ly α absorption sightlines with different values of T_0 , we adopt a semi-numerical approach to ‘paint’ on the temperature. We do this to each simulation cell using the density output from the simulation and setting the temperature according to equation (1) with our desired T_0 . This is done for all densities with no cutoff. This is a simplistic model that does not take into account the full evolution of the thermal state of the IGM. However, the purpose of this paper is to present our statistical method and demonstrate its accuracy and precision on simulated data so a simple temperature model for the IGM thermal state is sufficient to achieve these aims. We use $\gamma = 1.35$, which was calculated by fitting the initial simulation snapshot to a power law. Our thermal grid consists of 81 values of T_0 from $\log(T_0) = 3.4$ to $\log(T_0) = 4.4$ with $\Delta \log(T_0) = 0.0125$. The Ly α opacity, $\tau_{\text{Ly}\alpha}$ is related to the temperature via $\tau_{\text{Ly}\alpha} = n_{\text{HI}} \sigma_{\text{Ly}\alpha} \propto T^{-0.7}/\Gamma_{\text{HI}}$, see [Rauch \(1998\)](#). Because UV background photoionization, Γ_{HI} , is sourced by complex galaxy physics, it is not uniquely determined by the simulation. We therefore follow standard practice and adjust each model to have the same mean flux by rescaling τ such that $\langle e^{-\tau} \rangle = \langle F \rangle = 0.16$ at $z = 5$, which is within 1σ of the measurement presented in [Boera et al. \(2019\)](#). At $z = 6$ we use $\langle F \rangle = 0.011$ which is also consistent with recent measurements ([Becker et al. 2015](#); [D’Aloisio et al. 2018](#)).

2.3 Forward Modeling Real Observations

To mimic realistic observational data from echelle spectrographs, (e.g. from Keck/HIRES, VLT/UVES, and Magellan/MIKE) we forward model a resolution of $R = 30,000$ and a signal to noise ratio per pixel (SNR) of the unabsorbed continuum of 10 (35) at $z = 5$ ($z = 6$). The resolution smooths the flux by a Gaussian filter with $\text{FWHM} = 10 \text{ km s}^{-1}$ which means our simulations have ~ 4 pixels per FWHM of this resolution filter. For simplicity, we add flux-independent noise in the following way. We generate one 10,000 skewer $\times$ 1024 length realization of random noise all drawn from a Gaussian with $\sigma_N = 1/\text{SNR}$ and add this noise realization to every temperature model. An example skewer of our initial and forward-modeled data is shown in the top panels of Figures 2 and 4 respectively. Using the same noise realizations over the different models ensures that different noise realizations will not adversely affect the inference on the T_0 for mock data.

We assume a fiducial data set size of 8 quasar spectra at both $z = 5$ and $z = 6$ that probe a redshift interval of $\Delta z = 0.2$ per quasar for a total pathlength of $\Delta z = 1.6$ (equivalent to 29 skewers). In the discussion going forward, each mock data set consists of a random selection of 29 skewers without replacement.

3 WAVELET ANALYSIS

3.1 Formalism

Wavelets are localized in frequency and configuration space which allows wavelet amplitudes provide a breakdown of Fourier informa-

Figure 1. A complex Morlet Wavelet filter in real space with $s_n = 51.09 \text{ km s}^{-1}$. The solid line shows the real part of the wavelet while the dashed line shows the imaginary part. The width of the oscillations are set by the smoothing scale.

tion at all locations along a quasar sightline. Following [Lidz et al. \(2010\)](#), we calculate wavelet amplitudes from a “complex Morlet wavelet”, which is shown in Figure 1 and has the functional form:

$$\Psi_n(x) = A \exp(ik_0 x) \exp\left[-\frac{x^2}{2s_n^2}\right]. \quad (2)$$

The normalization, A , is set by requiring that $|\Psi_n(k)| = 1$. With this normalization, the Fourier transform of a complex Morlet wavelet is

$$\Psi_n(k) = \pi^{-1/4} \sqrt{\frac{s_n}{\Delta u}} \exp\left[-\frac{(k - k_0)^2 s_n^2}{2}\right]. \quad (3)$$

This is a Gaussian in configuration space centered on k_0 with width $\sigma_k = \sqrt{2}/s_n$. We also require that $k_0 s_n = 6$ to ensure these filters have a close to zero mean.

To begin the analysis on our simulated spectra, we first compute the flux contrast of the Ly α forest, δ_F :

$$\delta_F = \frac{F - \bar{F}}{\bar{F}}. \quad (4)$$

Then we convolve this flux contrast field with a wavelet filter of smoothing scale s_n resulting in a filtered spectrum, a_n :

$$a_n(x) = \int dx' \Psi_n(x - x') \delta_F(x') \quad (5)$$

The filtered spectrum is a complex number, the modulus of which is called the “wavelet amplitude” $A_n(x) = |a_n(x)|^2$. We define the power spectrum as

$$\langle \delta_F(k) \delta_F(k') \rangle = 2\pi P_F(k) \delta_D(k - k') \quad (6)$$

where δ_D is the Dirac Delta function. With this definition of the power, the average wavelet amplitude is

$$\langle A_n(x) \rangle = \int_{-\infty}^{\infty} dk' 2\pi [\Psi_n(k')]^2 P_F(k'). \quad (7)$$

In words, this means that the average wavelet amplitude is the power spectrum averaged over a Gaussian centered on wave-number $k_0 = 6/s_n$ with standard deviation $\sqrt{2}/s_n$. Therefore, this average wavelet amplitude is effectively a band-power.

Two wavelet amplitude spectra for an ideal simulated skewer at $z = 5$ are shown in the bottom two panels of Figure 2. For illustrative purposes in this section, we will mainly show wavelet amplitudes for $s_n = 51.09 \text{ km s}^{-1}$ though Figure 2 also shows $s_n = 77.44 \text{ km s}^{-1}$ for

a comparison. Ultimately in our analysis at $z = 5$, we will use fifteen logarithmically spaced values of $2200 \text{ km s}^{-1} > s_n > 5 \text{ km s}^{-1}$ as described in Section 4.2. For $z = 6$ we still use fifteen logarithmically spaced values of s_n with slightly shifted values due to the redshift dependence of the simulation resolution and box size.

The purpose of Figure 2 is to show the relationship between the flux and wavelet amplitudes for different smoothing scales. The top panel shows the flux for the three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). The middle panel shows the wavelet amplitude spectra for $s_n = 77.44 \text{ km s}^{-1}$ and the bottom show the wavelet amplitude spectra for $s_n = 51.09 \text{ km s}^{-1}$, both with the same values of T_0 as the top panel. The smoothing scale sets the size of the features in the flux that are picked out, when the smoothing scale and the feature size in the flux align the resulting wavelet amplitude is greater. The middle panel has a greater value of s_n than the bottom panel, so it is going to pick out wider features in the flux. Consider the peak in the flux at $\sim 550 \text{ km s}^{-1}$, which is smoother (and smaller) for $\log(T_0) = 4.1625$ (orange) flux than for $\log(T_0) = 3.4$ (blue). The corresponding wavelet amplitudes in the middle panel are greatest for $\log(T_0) = 3.4$ (blue) while the bottom panel are greatest for $\log(T_0) = 4.1625$ (orange), showing that the smaller feature in the flux agreed better with the smaller smoothing scale, as expected. The flux at $\log(T_0) = 4.4$ (green) is even smoother than the flux at $\log(T_0) = 4.1625$ (orange) but it does not have greater wavelet amplitude values in the bottom panel, this is because this peak corresponds to an ever smaller smoothing scale.

Figure 2 shows that the largest values of wavelet amplitudes correspond to peaks in the flux that are roughly the same width as the oscillations set by the wavelet smoothing s_n scale (an example of these oscillations can be seen in Figure 1). There can be offsets between features in flux and the corresponding features in the wavelet amplitude spectra, since the wavelets pick out features with the specific width set by s_n and, at larger scales, the wavelet will combine multiple features in the flux spectrum. This figure also demonstrates how wavelet analysis presents Fourier information in configuration space since the different wavelet amplitudes values convey frequency information along the quasar sightline.

In order to compare the spatial correlations between different values of s_n , consider Figure 3. The top panel of this figure shows a color plot of the wavelet amplitudes for different values of s_n along one line of sight; this is known as a "periodogram". The bottom panel of the plot is the flux used to calculate the wavelet amplitudes, which is the same as in Figure 2 for $\log(T_0) = 4.1625$. The large trough in the flux at $\sim 500 \text{ km s}^{-1}$ is seen in the wavelet amplitudes for scales up to $s_n \sim 40 \text{ km s}^{-1}$. The other troughs in the flux, such as the one near 750 km s^{-1} , are also seen in the wavelet amplitudes across multiple smoothing scales, most prominently at the smaller values of s_n . The overall decline in the average wavelet amplitude value for smaller values of s_n follows from the cutoff in the power spectrum, as is expected from equation (7).

As discussed in Section 2.3, we forward modeled our simulation skewers to mimic real data by including the effects of the resolution and noise. We illustrate the change in the flux as well as the wavelet amplitudes for one example skewer with $\log(T_0) = 4.1625$ in Figure 4. Note that the "clean" flux and wavelet amplitudes in this figure matches the model in Figure 2 from the same temperature model. From the Figure, we see that adding noise to the flux is able to shift features, add additional features, and change the amplitude of features in the wavelet amplitude spectrum. These effects are more prominent on smaller scales, as seen in the large differences between the models in the bottom panel of Figure 4, since the noise power

becomes comparable or greater than the flux power at these scales. Overall, this panel has greater values and more high-valued wavelet amplitudes for the forward modeled skewers than those without noise.

3.2 Wavelet Probability Density Function

As illustrated by Figures 2, 3, and 4 wavelet analysis converts the flux of the Ly α forest into wavelet amplitude spectra parameterized by s_n . The average value of the wavelet amplitude spectra corresponds to $P_F(k)$ via equation (7). The statistic we measure in our analysis is the wavelet amplitude probability density function (PDF), since this contains information on the full distribution of the wavelet amplitude values, rather than only the average.

PDFs for $s_n = 51.09 \text{ km s}^{-1}$ are shown in Figure 5 for three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). The top panel shows the PDFs calculated from the ideal simulation with clean flux. The bottom panel shows the PDFs after forward-modeling the simulation output with resolution and noise to mimic real data, as discussed in Section 2.3. In both the top and the bottom panel, the black dotted line shows the same PDF for pure noise draws with SNR = 10 and our pixel resolution.

In the top panel, the ideal PDFs are skewed to the left, with lower IGM temperatures corresponding to a higher mean value, as is expected from $P_F(k)$ and equation (7). The main effect of forward-modeling is the shift of the PDF from small values to larger values as was also seen in the bottom panel of Figure 4. This causes the suppression of wavelet amplitude values below $\sim 10^{-3}$. Initially the $\log(T_0) = 4.4$ (green) PDF had the largest tail below 10^{-3} , so the shift from small values to large values causes this model to change most dramatically from the top to the bottom panel. The PDF both shifts to the right and greatly increases the value of the PDF at the peak. The PDFs on the bottom panel are much more similar to the PDF for pure noise than in the top panel, which shows how the noise PDF is able to dominate over the signal. As the smoothing scale decreases and the overall PDF values decrease with it, as is inferred from Figure 3 and equation (7), the PDFs will become more dominated by the noise contribution.

Figure 5 demonstrates the ability of the wavelet amplitude PDF to differentiate T_0 models both with and without forward modeling. This confirms that the wavelet amplitude PDFs are promising statistics to measure the thermal state of the IGM. In addition, it illustrates how the PDF quantifies the full distribution of wavelet amplitudes for multiple sightlines, rather than the values along one sightline or the average value which is encoded in the the power spectrum.

4 STATISTICAL METHODS

The goal of this paper is to calculate the statistical precision with which a realistic quasar data set can constrain the parameters governing the small-scale structure of the IGM, here limited to T_0 , using wavelet analysis, specifically wavelet amplitude PDFs. The precision from this method can then be directly compared to the canonical approach using $P_F(k)$. We will also consider the precision achieved when combining the wavelet amplitude PDFs and power spectrum as has recently been attempted in the literature (Gaikwad et al. 2021).

To calculate the statistical precision, we will use Bayes' Theorem:

$$P(T_0|\text{data}) = \frac{P(\text{data}|T_0)P(T_0)}{P(\text{data})}. \quad (8)$$

Here the "data" vector depends on the statistical method for which

Figure 2. The top panel shows the flux from one simulation skewer at $z = 5$ for the three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). The middle panel shows the wavelet amplitude spectra for $s_n = 77.44 \text{ km s}^{-1}$ and the bottom show the wavelet amplitude spectra for $s_n = 51.09 \text{ km s}^{-1}$, both with the same T_0 values as the top panel. This shows that the largest values of wavelet amplitudes correspond to peaks in the flux that are roughly the same width as the oscillations set by the wavelet smoothing s_n scale.

we are calculating the precision. For the power spectrum, the "data" are the band-powers comprising $P_F(k)$. For the wavelet analysis, we will use multiple values of s_n and thus have multiple wavelet amplitude PDFs we must consider. In this case, the "data" will be the wavelet amplitude PDFs concatenated one after the other from largest to smallest s_n (which corresponds to smallest to largest k). Finally, when combining the wavelet and power spectrum analysis, the "data" vector will be the concatenated PDFs vector from the wavelet case with $P_F(k)$ added onto the end of it.

We assume a flat prior, $P(T_0)$, over the range of T_0 values we have simulation data for and will normalize the posterior, $P(T_0|\text{data})$, to unity so we don't need to explicitly calculate $P(\text{data})$. In order to calculate the likelihood, $P(\text{data}|T_0) = \mathcal{L}$, we assume a multivariate Gaussian distribution. This likelihood has the form:

$$\mathcal{L} = \frac{1}{\sqrt{\det(\Sigma)(2\pi)^n}} \exp\left(-\frac{1}{2}(\text{data} - \text{model})^T \Sigma^{-1}(\text{data} - \text{model})\right) \quad (9)$$

where $\Sigma = \Sigma(T_0)$ is the model dependent covariance matrix, n is the number of points in the data vector. Both the data and model vectors depend on the statistic we are using and will be discussed in their respective sections. The choice of a multivariate Gaussian distribution for the likelihood has been used in previous wavelet studies (Lidz et al. 2010; Gaikwad et al. 2021) as well as for studies

using the Ly α forest flux PDF (Lidz et al. 2006; Eilers et al. 2017). The base assumption is that each band power for $P_F(k)$ or each bin of the wavelet amplitude PDFs are Gaussian distributed. We show that this assumption is valid for our data in Appendix A.

In our analysis, we estimate the covariance matrix from mock draws of the data by

$$\Sigma(T_0) = \frac{1}{N_{\text{mocks}}} \sum_{i=1}^{N_{\text{mocks}}} (\text{mock}_i - \text{model})(\text{mock}_i - \text{model})^T \quad (10)$$

where N_{mocks} is the number of forward-modeled mock draws used. This method estimates a model dependent covariance, not the covariance of the data itself, since we are using many draws in our calculation. For the power spectrum calculation we use $N_{\text{mocks}} = 5,000$. We increase the number of mocks to $N_{\text{mocks}} = 1,000,000$ for the wavelet amplitude PDFs and the combination of the power spectrum and the wavelet amplitude PDFs, since these matrices are much larger with more values close to zero. Note that mocks are a random combination of 29 skewers without replacement. In theory, there are $(10,000!)/(29! \times 9,971!) \approx 10^{85}$ unique sets of 29 skewers from 10,000 skewers. This means that mock data sets will be correlated since they will contain skewers that are also in other mock data sets. However, we do not approach the the total possible number of combinations for these calculations and expect this effect to be negligible.

To visualize the covariance matrix for each method, we define

Figure 3. The bottom panel shows the flux from one simulation skewer with $\log(T_0) = 4.1625$. The top panel shows a periodogram of the wavelet amplitudes along the skewer for different smoothing scales. The different values of s_n will set different widths of oscillations that they pick out from the spectrum. This plot compares the location of the peaks and troughs in the wavelet amplitudes for different smoothing scales. It shows correlations between troughs at different values of s_n , for example the trough in the flux at -500 km s^{-1} can clearly be seen in the wavelet amplitudes for scales up to $s_n \sim 40 \text{ km s}^{-1}$. Note that the minimum wavelet amplitude shown on the plot is fixed at 10^{-8} for visual purposes.

the correlation matrix, C . The correlation matrix is the covariance matrix with the diagonal normalized to 1. This is done to the i th, j th element by

$$C_{ij} = \frac{\Sigma_{ij}}{\sqrt{\Sigma_{ii}\Sigma_{jj}}}. \quad (11)$$

4.1 Power Spectrum Likelihood

The resolution modifies $P_F(k)$ in a known way that can be corrected via the Fourier transform of the Gaussian resolution filter. There is also a white noise contribution to $P_F(k)$ due to the spectral noise which can be subtracted off. Therefore, the well known estimator for the true power is:

$$P_{\text{true}}(k) = \left\langle \frac{P_{\text{raw}}(k) - P_{\text{noise}}(k)}{W_R^2(k, \sigma_R, \Delta\nu)} \right\rangle \quad (12)$$

where $W_R(k, \sigma_R, \Delta\nu)$ is the Window function

$$W_R(k, \sigma_R, \Delta\nu) = \exp\left(-\frac{1}{2}(k\sigma_R)^2\right) \frac{\sin(k\Delta\nu/2)}{(k\Delta\nu/2)}. \quad (13)$$

We have Gaussian white noise with $\text{SNR} = 10$ added to the flux contrast, adding an extra factor of $\bar{F}$. The noise power is flat and has

a value of

$$P_{\text{noise}}(k) = \Delta\nu \left(\frac{1}{\text{SNR} \cdot \bar{F}} \right)^2 \quad (14)$$

where the factor of $\Delta\nu$ is our velocity pixel grid spacing which is $\Delta\nu = 2.7 \text{ km s}^{-1}$ at $z = 5$.

For $R = 30,000$, $\exp\left(-\frac{1}{2}(k\sigma_R)^2\right) < 0.24$ when $k \geq 0.4$. This implies that $W_R^2(k, \sigma_R, \Delta\nu) < .06$ when $k \geq 0.4$, so correcting these band-powers by this window function means dividing a noisy quantity, $P_F(k)$, by a very small number. When correcting by the window function, these band-powers blow up and the model covariance matrices we calculate via equation (10) are singular and ill-posed for inversion. We therefore choose to not correct by $W_R(k)$ in the "model" and "mock" data to ensure well-behaved covariance matrices. However, for visualization purposes we always show the resolution corrected power of equation (12) in the figures. The "model" is the power calculated from 10,000 flux skewers forward modeled with the resolution but not the noise, since there is no need to add additional noise when computing the mean. 10,000 skewers is equivalent to a total path-length of $200 \text{ Gpc } h^{-1}$. We calculate the "mock" data by computing the average $P_F(k)$ for 29 fully forward modeled skewers and then subtracting off the noise power, equation (14). This data set size is equivalent to an 8 quasar data set as discussed in Section 2.3.

To choose the k values for our mock, we used 15 logarithmic band-powers spanning from $2\pi/l_{\text{skewer}} = 0.0023 \text{ s km}^{-1}$ to $\pi/\Delta\nu = 1.2 \text{ s km}^{-1}$ at $z = 5$. The centers of the band-powers are listed in the first column of Table 1. We chose 15 band-powers in order to fully sample the shape of the power spectra while ensuring the low k (large scales) band-powers were populated by the discrete Fourier transform of the data. Figure 6 shows the power spectrum measurement, $P_F(k)$, for one mock data set with $\log(T_0) = 4.1625$ (black points) at $z = 6$. An equivalent figure for $z = 6$ can be found in Appendix C1. The 1σ error bars are calculated from the square root of the diagonal of the covariance matrix. Also shown are three models of the power spectra for three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). This mock data set visually seems to best agree with the model for $\log(T_0) = 4.1625$ (orange) for $k > 0.5$, which is the true T_0 of the model.

The model correlation matrix (see equation (11)) for the power spectrum at $\log(T_0) = 4.1625$ is shown in Figure 7. There are positive correlations (red) between the band-powers where $4 \times 10^{-2} \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$. The correlations between band-powers with $4 \times 10^{-2} \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$ and $k < 4 \times 10^{-2} \text{ s km}^{-1}$ the correlations are negative (blue). This behavior arises from the underlying spatial correlations of the Ly α forest and is consistent with what has been seen for real data (Walther et al. 2017; Boera et al. 2019). At values of $k > 0.2 \text{ s km}^{-1}$ (the smallest scales) noise dominates over the power spectrum, since the Ly α forest power spectrum exhibits a thermal cut-off at high- k , whereas the noise power spectrum is flat. Random Gaussian noise is uncorrelated, making it hard to recover the signal from the Ly α forest and results in the very weak correlations shown in the correlation matrix in the regions where $k > 0.2 \text{ s km}^{-1}$. We looked into the correlation matrix for $\text{SNR} = 50$ and $\text{SNR} = 100$ and found the values of the correlation matrix in the column above $k = 0.178 \text{ s km}^{-1}$ were stronger. This agrees with our interpretation of the the weak correlations in Figure 7 since with higher SNR the noise power is smaller and will not dominate until higher k .

Figure 4. The top panel shows the flux for the $\log(T_0) = 4.1625$ model from the simulation (orange line) as well as forward modeled with resolution and noise (black histogram). The middle panel shows the wavelet amplitude spectra for $s_n = 77.44 \text{ km s}^{-1}$ and the bottom show the wavelet amplitude spectra for $s_n = 51.09 \text{ km s}^{-1}$, both with the same T_0 value as the top panel. The simulation skewer is the same as that shown in Figure 2. This shows the effect noise has on the flux and the resulting wavelet amplitudes for one skewer.

Table 1. The first column contains the band-powers for the power spectrum. Next are the corresponding smoothing scales (s_n) used to calculate the wavelet amplitudes. The last two columns contain the minimum and maximum values used for the PDF estimation for each smoothing scale. These were chosen as the 1.5th and 98.5th percentiles of the data for the whole thermal grid at these smoothing scales.

$k \text{ (s km}^{-1}\text{)}$	$s_n = 6/k \text{ (km s}^{-1}\text{)}$	$\min \log(A_n)$	$\max \log(A_n)$
0.00278	2157.35	-0.671	1.683
0.00422	1423.32	-0.499	1.826
0.00639	939.04	-0.789	1.643
0.00968	619.54	-0.876	1.574
0.0147	408.74	-1.026	1.474
0.0222	269.67	-1.241	1.344
0.0337	177.91	-1.508	1.173
0.0511	117.38	-1.863	0.931
0.0775	77.44	-2.256	0.564
0.117	51.09	-2.651	0.039
0.178	33.71	-2.943	-0.430
0.270	22.24	-3.032	-0.585
0.409	14.67	-3.073	-0.627
0.620	9.68	-3.129	-0.688
0.939	6.39	-3.280	-0.830

4.2 Wavelet Amplitude PDF Likelihood

In previous work, Lidz et al. (2010) measured the thermal state of the IGM with wavelets by assuming a Gaussian likelihood and ignoring correlations between PDF bins as well as between PDFs from different smoothing scales. Gaikwad et al. (2021) measured the thermal state with wavelets assuming a Gaussian likelihood including correlations between bins of the same PDF but not between the PDFs for different smoothing scales. Here, we improve upon these previous work and present the likelihood calculation taking into consideration all correlations, both between PDF bins and between PDFs of different smoothing scales.

For wavelet amplitude PDFs, there is no analytic way to correct for the window function and subtract off the full noise PDF. Instead, we choose to use 10,000 forward-modeled skewers (with resolution and noise) to calculate the "model" wavelet amplitude PDFs, which is equivalent to calculating the wavelet amplitude PDFs for a total pathlength of 200 Gpc h^{-1} . The "mock" data is calculated from the same forward-modeled skewers as the "model", though "mock" data is the average of 29 skewers (equivalent to an 8 quasar data set, see Section 2.3).

As was mentioned in Section 3.1, we use fifteen values of s_n to get fifteen wavelet amplitude PDFs. These fifteen s_n correspond to the centers of the power spectrum band-powers, k , that were discussed in Section 4.1 and are listed in the second column of Table 1. Qualita-

Figure 5. This figure shows the PDFs for $s_n = 51.09 \text{ km s}^{-1}$ for three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). The top panel shows the PDFs calculated from 10,000 ideal simulation skewers (equivalent to a pathlength of 200 Gpc h^{-1}). The bottom panel shows the PDFs after forward-modeling with resolution and noise to mimic real data. In both the top and the bottom panel, the black dotted line shows the same PDF for pure noise draws with $\text{SNR} = 10$ and our pixel resolution. The difference in the mean values of these PDFs in each panel is expected from $P_F(k)$ and equation (7). The main effect of forward-modeling is the suppression of all wavelet amplitude values below $\sim 10^{-3}$, where the data is beginning to be dominated by the noise.

Figure 6. The power spectrum measurement, $P_F(k)$, for one mock data set with $\log(T_0) = 4.1625$ (black points). The 1σ error bars are calculated from the square root of the diagonal of the covariance matrix. Also shown are model values of the power spectra for three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green).

Figure 7. The correlation matrix for the power spectrum at $\log(T_0) = 4.1625$. The positive (red) and negative (blue) correlations on scales $k < 0.2 \text{ s km}^{-1}$ arise from underlying spatial correlations of the Ly α forest. The very weak correlations seen in the regions where $k > 0.2 \text{ s km}^{-1}$ are due to uncorrelated random Gaussian noise which dominates the signal on small-scales (high k).

tively, this should ensure that the wavelet amplitude PDFs contain at least as much information as the power spectrum due to equation (7), allowing us to compare the resulting precision on an equal footing. To estimate the wavelet amplitude PDFs for each smoothing scale, we calculate histograms. This introduces three histogram parameters into our analysis: the maximum wavelet amplitude considered, the minimum wavelet amplitude considered, and the number of bins in the histogram.

When selecting the minimum and maximum wavelet amplitude considered for our PDF estimation, we want to ensure that all bins will be populated for the whole thermal grid so that the covariance matrix is well posed for inverting. We also want to make sure the maximum and minimum values span a large enough range to capture the most significant differences in the shape of the PDF. For these reasons, we chose the maximum and minimum values for the PDFs by calculating the 1.5th and 98.5th percentile of the "model" wavelet amplitudes calculated for every thermal model in our grid. The maximum and minimum values of the wavelet amplitudes considered for each smoothing scale s_n are listed in Table 1. We also need to select a number of bins that will sufficiently sample the shape of the PDF without making the data vector too long and the covariance matrix ill-suited for inversion. We found that 10 bins was a reasonable choice to achieve these aims.

Figure 8 shows the PDFs from one mock data set for each s_n with $\log(T_0) = 4.1265$ (black points) and $z = 5$. An equivalent figure for $z = 6$ can be found in Appendix C1. The 1σ error bars are calculated from the square root of the diagonal of the covariance matrix. Each panel also shows the "model" values of the PDFs for three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). This figure qualitatively illustrates the ability of the wavelet PDF to differentiate between different T_0 values, which we formally quantify with Bayesian inference as discussed in Section 4. The "model" PDFs for $s_n < 22.24 \text{ km s}^{-1}$ all overlap because the noise dominates the signal on these scales and all three PDFs are equivalent to the pure noise PDF.

In order to understand the correlations present between the bins of a single wavelet PDF, we first calculate the model covariance

Figure 8. The black points show the PDFs from one mock data set for each s_n with $\log(T_0) = 4.1265$. The 1σ error bars are calculated from the square root of the diagonal of the covariance matrix. Each panel also shows the "model" values of the PDFs from the stated smoothing scale for three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). This figure qualitatively illustrates the ability of the wavelet PDF to differentiate between different T_0 values, which we formally quantify with Bayesian inference as discussed in Section 4.

matrix for $s_n = 51.09 \text{ km s}^{-1}$ and $\log(T_0) = 4.1265$ and then plot the correlation matrix in Figure 9. There are positive correlations (red) between the bins at small wavelet amplitudes $A_n < 5 \times 10^{-2}$ with the other small values. For larger values, there are negative correlations (blue) between the larger wavelet amplitudes $A_n > 0.1$ and all other wavelet amplitude values. These effects are due to the shape of the PDF as well as the constraint that the PDF must integrate to 1. Increasing the counts for any wavelet amplitude value will cause the counts in the peak of the PDF (around $A_n \sim 0.1$ as seen in Figure 8 for $s_n = 51.09 \text{ km s}^{-1}$) to decrease due to the integral constraint on the PDF. Meanwhile, the shape of the PDF means that when one bin

along the tail ($A_n < 4 \times 10^{-2}$ as seen in Figure 8 $s_n = 51.09 \text{ km s}^{-1}$) increases in counts, the other tail bins will increase as well.

Ultimately, we will combine fifteen wavelet amplitude PDFs, each with a different value of s_n , in our measurement. Our measurement will include the correlations between the different PDFs, unlike the measurements from both Lidz et al. (2010) and Gaikwad et al. (2021) which ignore these correlations. We include these correlations by using non-zero off diagonal terms in each covariance matrix, $\Sigma(T_0)$, when computing the likelihood in equation (9). The correlations between different wavelet amplitude PDFs have never been considered in the previous literature on the Ly α forest. Our data vector

Figure 9. The correlation for the wavelet amplitude PDF for $\log(T_0) = 4.1625$ for $s_n = 51.09 \text{ km s}^{-1}$. There are positive correlations (red) between the bins at small wavelet amplitudes $A_n < 5 \times 10^{-2}$ with the other small values. For larger values, there are negative correlations (blue) between the wavelet amplitudes $A_n > 0.1$ and all other wavelet amplitude values. These effects are due to the shape of the PDF as well as the constraint that the PDF must integrate to 1.

is a concatenation of each wavelet amplitude PDF starting with the largest value of s_n going down to the smallest value, making it $n_{\text{bins}} \times n_{s_n} = 10 \times 15 = 150$ points long. We expect that these correlations between different s_n values will be non-negligible due to the spatial correlations shown in the periodogram (Figure 3) as well as in the power spectrum (Figure 7).

In this case, the correlation matrix has dimension 150×150 and is shown in Figure 10. For visual purposes, the axes are labeled by the smoothing scale used to calculate the wavelet amplitude PDFs, but the correlations shown are between the wavelet amplitude bins (such as the labels in Figure 9). Each 10×10 block along the diagonal is the correlation matrix for a single s_n value. These diagonal blocks all appear very similar to the example shown for $s_n = 51.09 \text{ km s}^{-1}$ in Figure 9, as expected from the similar shaped PDFs.

The wavelet amplitude PDFs for large smoothing scales, $2157 \text{ km s}^{-1} \geq s_n \geq 33.7 \text{ km s}^{-1}$, have significant correlations off the diagonal. The correlations between the PDFs with $177.9 \text{ km s}^{-1} > s_n > 33.71 \text{ km s}^{-1}$ have the same pattern as the diagonal blocks modified by a small positive number (appearing mostly red). The correlations between the PDFs for $s_n > 408.7 \text{ km s}^{-1}$ and $177.9 \text{ km s}^{-1} > s_n > 33.71 \text{ km s}^{-1}$ have the same pattern as the diagonal blocks modified by a small negative number (appearing mostly blue). These modifications follow the same pattern as that in the correlation from the power spectrum shown in Figure 7 where there are positive correlations (red) between $4 \times 10^{-2} \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$ and negative correlations between $4 \times 10^{-2} \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$ and $k < 4 \times 10^{-2} \text{ s km}^{-1}$. This pattern arises from the underlying spatial correlations of the Ly α forest as was discussed for the power spectrum. For $s_n \leq 22.2 \text{ km s}^{-1}$, the wavelet amplitudes begin to be dominated by noise, so the correlations between the PDFs of different smoothing scales become very small. This again mimics the behavior seen for $k > 0.2 \text{ s km}^{-1}$ in the power spectrum correlation matrix shown in Figure 7.

Many of these off diagonal covariance elements are very small, and it is challenging to measure small correlations from finite noisy data

sets. For the full set of fifteen wavelet amplitude PDFs, the covariance matrix is 100 times larger than the power spectrum covariance matrix, making the calculation even more time consuming and difficult. This noise in the wavelet amplitude PDFs covariance matrix becomes quite noticeable in the posterior measurement on T_0 . We reduce this noise by smoothing the covariance matrices over multiple thermal grid with one spline per matrix element. A detailed discussion of this smoothing can be found in Appendix B.

4.3 Joint Wavelet-Power Likelihood

Gaikwad et al. (2021) combined the wavelet amplitude PDFs with the power spectrum as well as the Doppler parameter distribution and curvature statistics to improve upon each of the individual measurements of the thermal state of the IGM. They did this by ignoring the correlations between PDFs for different smoothing scales as well as between the PDFs and the power spectrum despite the fact that these statistics were all measured from the same data set, and are thus surely correlated. This application has motivated us to combine the power spectrum and wavelet amplitude PDFs while paying careful attention to correlations to see if this improves the precision of our mock measurement. We expect there to be non-negligible correlations between the wavelet amplitude PDFs and the power spectrum from equation (7), since this says the mean wavelet amplitude, i.e. the first moment of the wavelet PDF, contains the same information as a band power.

When combining the wavelet amplitude PDFs and the power spectrum, the data vector is the 150 element wavelet amplitude PDFs, i.e. 10 PDF bins $\times$ 15 smoothing scales discussed in Section 4.2, with the addition of the 15 band-powers of $P_F(k)$ added to the end. This makes the full data vector 165 points long and the correlation a complicated 165×165 matrix. To build intuition, we will first consider a subset of the full correlation matrix that consists of a single wavelet amplitude PDF with the power spectrum. The correlation matrix in this situation will be only 25×25 , i.e. 10 wavelet PDFs values and 15 band-powers.

The correlation matrix for the wavelet amplitude PDF from $s_n = 51.09 \text{ km s}^{-1}$ and the power spectrum is shown in Figure 11. The top panel has the full correlation matrix with the axes labeled by either the smoothing scale, $s_n = 51.09 \text{ km s}^{-1}$, which was used to calculate the A_n or "Power" representing the different values of k . The top right 15×15 diagonal block is identical to the correlation matrix for the power spectrum shown in Figure 7 and the bottom left 10×10 diagonal block is identical to the correlation matrix for one $s_n = 51.09 \text{ km s}^{-1}$ shown in Figure 9.

The off diagonal blocks show the correlations between the wavelet amplitude PDFs and the power spectrum. The bottom right rectangle of the correlation matrix is blown up in the bottom panel of the figure with the axes appropriately labeled by the wavelet amplitude A_n from the PDF and the k from the power spectrum. The strongest correlations (both positive and negative) between the power and wavelet amplitude PDF are found in the column at $k = 6/s_n = 0.12 \text{ s km}^{-1}$. This k value corresponds to the same scales probed by $s_n = 51.09 \text{ km s}^{-1}$. As the value of this k bin increases, we expect the wavelet amplitude PDF to shift to higher values so that the average wavelet amplitude in the PDF increases, as required by equation (7). This shift causes the larger values of A_n ($A_n > 0.2 \text{ km s}^{-1}$) to be more common, resulting in a positive correlation (red) with larger PDF bins, while the smaller values of A_n ($A_n < 0.1 \text{ km s}^{-1}$) are less common, resulting in a negative correlation (blue).

The behavior seen in the $k = 0.12 \text{ s km}^{-1}$ column of the bottom panel is replicated for the columns above $3 \times 10^{-2} \text{ s km}^{-1} \lesssim$

Figure 10. The correlation matrix for fifteen wavelet amplitude PDFs at $\log(T_0) = 4.1625$. The wavelet amplitude PDFs for large smoothing scales, $2157 \text{ km s}^{-1} \geq s_n \geq 33.7 \text{ km s}^{-1}$, have significant correlations off the diagonal. The correlations between the PDFs with $177.9 \text{ km s}^{-1} > s_n > 33.71 \text{ km s}^{-1}$ have the same pattern as the diagonal blocks modified by a small positive number (appearing mostly red). The correlations between the PDFs for $s_n > 408.7 \text{ km s}^{-1}$ and $177.9 \text{ km s}^{-1} > s_n > 33.71 \text{ km s}^{-1}$ have the same pattern as the diagonal blocks modified by a small negative number (appearing mostly blue). For $s_n \leq 22.2 \text{ km s}^{-1}$, the wavelet amplitudes begin to be dominated by noise, so the correlations between the PDFs for different values of s_n become very small. This pattern mimics that seen in the power spectrum correlation shown in Figure 7.

$k \lesssim 0.2 \text{ s km}^{-1}$ modified by a small positive number. The columns where $k < 4 \times 10^{-2} \text{ s km}^{-1}$ show the same behavior modified by a small negative number. These modifications mimic the pattern in the correlations for the power spectrum, as shown in the upper right quadrant of the top panel and Figure 7. In particular, it replicates the positive (red) correlations for $3 \times 10^{-2} \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$ with $k = 0.12 \text{ s km}^{-1}$ and the negative (blue) correlations for $k < 4 \times 10^{-2} \text{ s km}^{-1}$ with $k = 0.12 \text{ s km}^{-1}$.

As discussed in the beginning of this section, the total data vector for the combination of the wavelet amplitude PDFs and the power spectrum is 165 points long. The full 165×165 correlation matrix is shown in Figure 12. The axes are labeled by either the smoothing scale used to calculate A_n or "Power" representing the k bands. The bottom left 150×150 diagonal block is identical to the correlation matrix for the wavelet amplitude PDFs shown in Figure 10 while the top right 15×15 diagonal block is identical to the correlation matrix for the power spectrum shown in Figure 7. The right most column above "Power" shows similar behavior for each smoothing scale as

was discussed for the bottom panel of Figure 11. The column with the strongest correlations for each smoothing scale always corresponds to $k = 6/s_n$ and the behavior in other columns above "power" follow the strongest bin modified by the correlations between the power bins.

This data vector is larger than the data vector of the wavelet amplitude PDFs, which was discussed in Section 4.2. Similarly, the noise in this covariance matrix is non-negligible and so we smooth the covariance matrix over the thermal grid with a spline in order to calculate the posteriors. This is discussed in more detail in Appendix B.

5 RESULTS

5.1 T_0 Measurements

We can calculate the posterior probability of T_0 given a mock data set, $P(T_0|\text{data})$, from equation (8). In Figure 13 we compare the posterior

Figure 11. The top panel shows the correlation matrix for wavelet amplitude PDF with $s_n = 51.09 \text{ km s}^{-1}$ and the power spectrum for $\log(T_0) = 4.1625$. The dotted lines separate the part of the matrix that corresponds to the wavelet amplitude PDF (labeled by $s_n = 51.09 \text{ km s}^{-1}$) and the power spectrum (labeled "Power"). The bottom left diagonal block is identical to the correlation matrix shown in Figure 9 while the top right diagonal is identical to the correlation matrix shown in Figure 7. The bottom right rectangle of the correlation matrix is blown up in the bottom panel, which shows the correlations between the power spectrum and the wavelet amplitude PDFs with the appropriate labels of A_n and k . The strongest correlations (both positive and negative) between the power and wavelet amplitude PDF are found in the column at $k = 6/s_n = 0.12 \text{ s km}^{-1}$. This k value corresponds to the same scales probed by $s_n = 51.09 \text{ km s}^{-1}$. The behavior seen in the $k = 0.12 \text{ s km}^{-1}$ column of the bottom panel is replicated for the columns above $3 \times 10^{-2} \text{ s km}^{-1} \leq k \leq 0.2 \text{ s km}^{-1}$ modified by a small positive number. The columns where $k < 4 \times 10^{-2} \text{ s km}^{-1}$ show the same behavior modified by a small negative number. These modifications mimic the pattern in the correlations for the power spectrum, in particular the positive (red) correlations for $3 \times 10^{-2} \text{ s km}^{-1} \leq k \leq 0.2 \text{ s km}^{-1}$ with $k = 0.12 \text{ s km}^{-1}$ and the negative (blue) correlations for $k < 4 \times 10^{-2} \text{ s km}^{-1}$ with $k = 0.12 \text{ s km}^{-1}$.

distribution of T_0 from one mock data set at $z = 5$ from three different methods: the power spectrum (blue triangles), the wavelet amplitude PDFs (orange circles), and both the power spectrum and wavelet amplitude PDFs (green triangles). This mock data set is the same one shown in Figures 6 and 8. Visually, the wavelet amplitude PDFs provide a more precise posterior for T_0 than the power spectrum does while. The measurement of T_0 for these two methods are $T_0 = 14,900^{+1500}_{-1500} \text{ K}$ (power spectrum) and $T_0 = 14,100^{+1400}_{-1400} \text{ K}$ (wavelet amplitude PDFs). These errors are calculated by interpolating the cumulative distribution function (CDF) of the posterior onto the 15.9th and 84.1th percentiles, which correspond to the 1σ percentiles for a normal distribution. This region between these percentiles will be referred to as the 1σ region and the errors calculated from it as the equivalent 1σ errors throughout the end of this paper. The 1σ region

Table 2. This table shows the mean and variance of the 1σ values from 1,000 mock data sets at $\log(T_0) = 4.1265$ at both $z = 5$ and $z = 6$. The 1σ errors for the wavelet amplitudes PDFs are on average 7% smaller at $z = 5$ and 12% smaller at $z = 6$ than those for the power spectrum, though they do have a higher variance. The error calculated from combining the power spectrum and the wavelet analysis PDFs at $z = 5$ does not improve the errors on average over the wavelet amplitude PDFs alone.

z	Method	$\bar{\sigma}_+$	$\bar{\sigma}_-$
5	Power Spectrum	1490 ± 50	-1520 ± 50
	Wavelet Amplitude PDFs	1400 ± 200	-1400 ± 200
	Both	1400 ± 200	-1400 ± 200
6	Power Spectrum	3030 ± 190	-3140 ± 200
	Wavelet Amplitude PDFs	2700 ± 600	-2700 ± 600

is 7% smaller for the wavelet amplitude PDFs posterior than the power spectrum posterior, showing that the wavelet amplitude PDF is more sensitive to T_0 than the power for this mock data set. Combining the power spectrum and the wavelet PDFs has a negligible effect on the posterior distribution resulting in $T_0 = 14,100^{+1500}_{-1300} \text{ K}$. This 1σ region has the same width as the one for the wavelet amplitude PDFs alone and so the combination does not improve the measurement's precision.

In Figure 14 we compare the posterior distribution of T_0 from one mock data set at $z = 6$ from two different methods: the power spectrum (blue triangles) and the wavelet amplitude PDFs (orange circles). This mock data set is the same one shown in Figures C1 and C3. The measurement of T_0 for these two methods are $T_0 = 11,000^{+3000}_{-3000} \text{ K}$ (power spectrum) and $T_0 = 13,000^{+2000}_{-3000} \text{ K}$ (wavelet amplitude PDFs) where the errors are calculated in the same way as they were for $z = 5$. At this redshift, the wavelet amplitude PDF measurements resulted in a 20% improvement of the 1σ errors when compared to the results from power spectrum measurement from the same mock data, almost three times the improvement seen at $z = 5$, though here the errors only have one significant figure so improvement is a coarser measurement.

To further quantify the difference in the precision of these posteriors, we calculated the equivalent 1σ errors for 1,000 mock data sets at $\log(T_0) = 4.1265$ for both $z = 5$ and $z = 6$. These resulting mean and variance of these values are listed in Table 2. On average, the posteriors for the wavelet amplitude PDFs at $z = 5$ are 7% smaller than those from the power spectrum, again showing that the wavelet amplitude PDFs are more sensitive on average than the power spectrum. However, the variance on these power spectrum errors are 75% smaller, meaning the power spectrum posteriors are more consistently large while the wavelet amplitude PDFs vary more in size. The average errors on the posteriors from combining both the wavelet amplitude PDFs and the power spectrum show no improvement over the errors from the wavelet amplitude PDFs alone again showing this combination does not improve the measurement.

For $z = 6$, the posteriors for the wavelet amplitude PDFs are 12% smaller than those from the power spectrum while the variance on the power spectrum errors are 67% smaller. This means that the wavelet amplitude PDFs are again more sensitive on average but the errors vary more than the power spectrum. This agrees with the results at $z = 5$, though again we find that the wavelets lead to an even greater improvement on the average sensitivity by a factor of two. Physically, at this higher redshift, more of the spectra consists of absorption troughs, giving the wavelet amplitude PDFs even greater potential to improve on the power spectrum measurements since they maintain spatial information.

Figure 12. The correlation for all fifteen wavelet amplitude PDFs and the power spectrum combined for $\log(T_0) = 4.1625$. The axes are labeled by either the smoothing scale used to calculate A_n or "Power" representing the k bands. The bottom left 150×150 diagonal block is identical to the correlation matrix for the wavelet amplitude PDFs shown in Figure 10 while the top right 15×15 diagonal block is identical to the correlation matrix for the power spectrum shown in Figure 7. The right most column above "Power" shows similar behavior for each smoothing scale as was discussed for the bottom panel of Figure 11. The column with the strongest correlations for each smoothing scale always corresponds to $k = 6/s_n$ and the behavior in other columns above "power" follow the strongest bin modified by the correlations between the power bins.

5.2 Inference Testing

In order to test the fidelity of our statistical inference procedure and results, we perform an inference test. The goal of this test is to check that this calculated posterior behaves as a posterior probability should: if the true value of T_0 for the mock data falls into the equivalent of the 1σ region of the posterior $\sim 68\%$ of the time (and the 2σ region 95% of the time). We again calculate these equivalent 1σ and 2σ regions for our posteriors in the same way as discussed in Section 5.1. We integrate the posterior to get the CDF onto the 15.9th and 84.1th percentiles for 1σ and onto the 2.3rd and 97.7th percentile for 2σ . These percentiles correspond to the 1σ and 2σ percentiles for a normal distribution. From here, we count the number of times the true value of T_0 fell into these regions region. Ideally, the true value of T_0 should fall into the 1σ region 68.3% of the time and it should fall into the 2σ region 95.4% of the time.

We did this for 1,000 mock data sets at three different values of T_0 for $z = 5$ and one value of T_0 at $z = 6$. We chose to only look at one value of T_0 at the higher redshift because the posteriors are broader

and we are more likely to run into edge effects at the other T_0 values. The errors are calculated by $\sqrt{N}/10000$ where N is the number of times the true value fell into the desired region and 1,000 is the total number of mocks used. The results, shown in Table 3, are consistent with the expected values of 68.3% and 95.4% within the calculated errors and we pass this inference test.

5.3 Ignoring Correlations

We further investigated the posterior distributions from the wavelet amplitude PDFs and the combined wavelet amplitude PDFs and power spectrum measurements at $z = 5$ while ignoring certain correlations. To begin, we considered the posterior from the wavelet amplitude PDFs alone. We constructed three distinct covariance matrices from our initial full calculation, which is shown in Figure 10. The first covariance matrix considered is made up of the same values along the diagonal and zeros for all off-diagonal elements. This is similar to the covariance considered in Lidz et al. (2010) and is

Figure 13. The posterior on T_0 for one mock data at $z = 5$ set from three different methods: the power spectrum (blue triangles), the wavelet amplitude PDFs (orange circles), and the combination of the power spectrum and the wavelet amplitude PDFs (green triangles). The mock data set used to calculate these posteriors is shown for the power in Figure 6 and for the wavelet amplitude PDFs in Figure 8. The vertical dotted red line shows the true value of T_0 for the mock data set. Qualitatively, the posterior from the power spectrum is less precise than the posterior from the wavelet amplitude PDFs and the combination of the power spectrum and wavelet amplitude PDFs do not improve the precision of the posterior over the wavelet amplitude PDFs alone. The text in the corner is a quantitative measurement of T_0 , giving the median value with the equivalent 1σ errors for each method according to their colors in the same order as the legend.

Table 3. This table shows the results of our inference test for three values of T_0 and three statistical methods (power spectrum, wavelet amplitude PDFs, and the combination of the two) at $z = 5$. We have also included our results for one value of T_0 and two statistical methods (power spectrum and wavelet amplitude PDFs) at $z = 6$. We calculated the equivalent 1σ and 2σ regions from the CDF and then determined the frequency with which the true T_0 values fell into these regions. These results are presented for 1,000 mock data sets and are consistent with a true distribution function with our expected errors.

z	Method	$\log(T_0)$	% in 1σ	% in 2σ
5	Power Spectrum	3.9	70.0 ± 2.6	94.6 ± 3.1
		4.1265	68.9 ± 2.6	96.0 ± 3.1
		4.2875	68.7 ± 2.6	96.2 ± 3.1
	Wavelet Amplitude PDFs	3.9	63.5 ± 2.5	94.0 ± 3.1
		4.1265	67.6 ± 2.6	95.0 ± 3.1
		4.2875	69.3 ± 2.6	95.9 ± 3.1
	Both	3.9	63.7 ± 2.5	93.1 ± 3.1
		4.1265	67.1 ± 2.6	94.1 ± 3.1
		4.2875	68.5 ± 2.6	95.3 ± 3.1
6	Power Spectrum	4.1265	68.1 ± 2.6	95.6 ± 3.1
	Wavelet Amplitude PDFs	4.1265	65.7 ± 2.6	93.6 ± 3.1

referred to as the “no correlations” model in Figure 15 and Table 4. Next we construct a covariance matrix that includes the correlations between bins of the individual wavelet amplitude PDFs but ignores the correlations between different values of s_n . The resulting correlation matrix would have the same values for the fifteen 10×10 diagonal blocks in Figure 10 and zeros at all other locations. This is the similar to the covariance model considered in Gaikwad et al. (2021) and is referred to as the “PDF bin correlations” model in Figure 15 and Table 4. Finally we considered the full covariance matrix presented in this work in Figure 10, which we have labeled as “all correlations” in Figure 15 and Table 4. The resulting posteriors from

Figure 14. The posterior on T_0 for one mock data set at $z = 6$ from two different methods: the power spectrum (blue triangles) and the wavelet amplitude PDFs (orange circles). The mock data set used to calculate these posteriors is shown for the power in Figure C1 and for the wavelet amplitude PDFs in Figure C3. The vertical dotted red line shows the true value of T_0 for the mock data set. Qualitatively, the posterior from the power spectrum is less precise than the posterior from the wavelet amplitude PDFs. The text in the corner is a quantitative measurement of T_0 , giving the median value with the equivalent 1σ errors for each method according to their colors in the same order as the legend.

Figure 15. The posterior on T_0 for one mock data set for the wavelet amplitude PDFs using three different covariance matrices. The three matrices are described in more details in Section 5.3. They are: a diagonal-only covariance matrix which includes no correlations (blue triangles), a diagonal-block matrix that only contains correlations between PDF bins for the same wavelet scale (orange circles), and the full covariance matrix with all correlations (green triangles). The posterior from the diagonal matrix which has no correlations is much more narrow than the other two posteriors which are roughly the same width and height as each other.

these three models is shown in Figure 15. For the mock dataset shown in this figure (which is the same one shown throughout the rest of the paper) the “no correlations” model reduces the width of the posterior while the “PDF bin correlations” remains a similar width when compared to “all correlations”. The median value does not agree for any of these posteriors though the whole distribution of the posteriors have significant overlap. We also performed the same inference test on the posteriors for these models which will be discussed at the end of this section with results in Table 4.

Figure 16. The posterior on T_0 for one mock data set for the combination of wavelet amplitude PDFs and the power spectrum using three different covariance matrices. The three matrices are described in more details in Section 5.3. They are: a diagonal-block matrix with 16 distinct blocks for the wavelet PDF bin correlations and power correlations separately (blue triangles), a diagonal-block matrix with 2 distinct diagonal blocks that contains the full wavelet correlations and the power correlations separately (orange circles), and the full covariance matrix with all correlations including the cross-correlations between the wavelet and the power spectrum (green triangles). Adding additional correlations caused the posterior distribution to broaden each time.

Next, we again constructed three different covariance matrices for the combination of the wavelet amplitude PDFs and the power spectrum where the initial covariance matrix is shown in Figure 12. First we considered only the correlations between PDF bins for the wavelet amplitude and the full correlations for the power spectrum. The resulting covariance matrix has fifteen 10×10 diagonal blocks followed by one 15×15 diagonal block and zeros at all other locations. This is similar to how the combination of different wavelet scales and different statistics were done in Gaikwad et al. (2021). We refer to this model as the “PDF bin correlations” model in Figure 16 and Table 4. The subset of this matrix for the wavelet PDFs matches that used in Figure 15 with the same name. Next we consider the full wavelet correlations for the PDF bins as well as the different wavelet scales combined with the full power correlations but ignoring all cross correlations. The resulting matrix would have two diagonal blocks: one for the wavelet amplitude PDFs that has dimensions 150×150 and is shown in Figure 10, and one for the power spectrum that has dimensions 15×15 and is shown in Figure 7. We refer to this model as the “wavelet correlations” model in Figure 16 and Table 4. We also again considered the full covariance matrix presented in this work in Figure 12, which we have labeled as “All correlations” in Figure 16 and Table 4. The resulting posteriors from these three models is shown in Figure 16. For the mock dataset shown in this figure (which is the same one shown throughout the rest of the paper) the “PDF bin correlations” has the most narrow posterior while the “wavelet correlations” is more narrow than “all correlations” but broader than “PDF bin correlations”. Again, each posterior has a shifted median value compared to the others, though all the posteriors have significant overlap with each other.

We repeat the inference test described in section 5.2 for these models and have presented the results in Table 4. For the wavelet amplitude PDF models, only the “all correlations” model passed our inference test with the true value of T_0 falling in our 1σ region for 67.6% of mock posteriors and the true value of T_0 falling in our

2σ region for 95.0% of mock posteriors. In comparison, the “no correlations” model only had the true value of T_0 falling in our 1σ region for 35.8% of mock posteriors and the true value of T_0 falling in our 2σ region for 62.7% of mock posteriors. We can roughly estimate that, since $68/36 \sim 1.9$, we would need to widen the posterior for the “no correlations” model by a factor of 1.9 to pass the inference test. If we take the “all correlations” model as the true information contained in the wavelet amplitude PDFs then the “no correlations” posterior needs to be both shifted and broadened to a lesser extent to match this posterior. A similar calculation can be done for the “PDF bin correlations” model which is more similar to the “all correlations” model initially.

The results for the inference test on the different models of the combined wavelet amplitude PDF and power spectrum correlations are also shown in Table 4. Here, again, only the “all correlations” model passed our inference test with the true value of T_0 falling in our 1σ region for 67.1% of mock posteriors and the true value of T_0 falling in our 2σ region for 94.1% of mock posteriors. In comparison, the “PDF bin correlations” model only had the true value of T_0 falling in our 1σ region for 47.7% of mock posteriors and the true value of T_0 falling in our 2σ region for 79.4% of mock posteriors. We can roughly estimate that, since $68/48 \sim 1.4$, we would need to widen the posterior for the “PDF bin correlations” model by a factor of 1.4 to pass the inference test. If we take the “all correlations” model as the true information contained in the combination of the wavelet amplitude PDFs and the power spectrum then the “PDF bin correlations” posterior needs to be both shifted and broadened to a lesser extent to match this posterior. A similar calculation can be done for the “wavelet correlations” model which is more similar to the “all correlations” model to start.

5.4 Comparison to Previous Work

Lidz et al. (2010) made a measurement of the thermal state of the IGM with one wavelet amplitude PDF while ignoring the correlations between bins of the PDF. This measurement of T_0 is higher than those from most other statistics and also has larger error bars. This is shown in Walther et al. (2019), which used the same quasar data set as Lidz et al. (2010) when measuring the power spectrum but added additional data to roughly double the data set size. Figure 15 of Walther et al. (2019) shows the resulting thermal state constraints compared to Lidz et al. (2010) as well as other measurements. We investigated the effect of ignoring the correlations between PDF bins in Section 5.3 and found that this would result in underestimated errors. This would imply that the error bars from Lidz et al. (2010) would not reflect the true precision of the measurement. Additionally, Figure 15 shows that ignoring these correlations also shifts the peak of the posterior. However, we did not consider only using one smoothing scale as was done in Lidz et al. (2010), which we would expect to broaden the posterior of the measurement. Thus we can not precisely estimate how shifted or underestimated the measurement and errors from Lidz et al. (2010) would have been.

Gaikwad et al. (2021) made a measurement of the thermal state of the IGM using eight wavelet amplitude PDFs with $30 \text{ km s}^{-1} \leq s_n \leq 100 \text{ km s}^{-1}$, the power spectrum, and by combining these statistics along with the Doppler parameter distribution and curvature statistics. All of these statistics were calculated from the same data set and the correlations between PDF bins of different smoothing scale and the individual statistics were ignored. The main results are shown in Figure 14 of Gaikwad et al. (2021), where the different statistics have measurements that agree with each other but differ in error bar size and the joint constraints visually appear to have the smallest error bars

Table 4. This table shows the results of our inference test when ignoring correlations for either the wavelet amplitude PDFs or the combination of the wavelet amplitude PDFs and the power spectrum. The three models of correlations for each statistic is described in Section 5.3. The inference test was done for only one true value of T_0 , $\log(T_0) = 4.1265$. We calculated the equivalent 1σ and 2σ regions from the CDF and then determined the frequency with which the true T_0 values fell into these regions for 1,000 mock data sets. Only the models that considered all the correlations as presented in this paper labeled as “all correlations” for each statistic passed the inference test. The other two models for both statistics did not recover the true value of T_0 the expected number of times.

Method	Correlations	% in 1σ	% in 2σ
Wavelet Amplitude PDFs	No correlations	35.8 ± 1.9	62.7 ± 2.5
	PDF bin correlations	55.2 ± 2.3	87.0 ± 2.9
	All Correlations	67.6 ± 2.6	95.0 ± 3.1
Both PDFs and Power	PDF bin correlations	47.7 ± 2.2	79.4 ± 2.8
	Wavelet correlations	55.3 ± 2.4	86.5 ± 2.9
	All Correlations	67.1 ± 2.6	94.1 ± 3.1

most frequently. Our analysis has shown that combining the wavelet amplitude PDFs and the power spectrum from the same data set does not improve the measurement of T_0 on average. However, we did not further investigate the additional statistics that Gaikwad et al. (2021) considered nor did we consider the situation where the smoothing scales used to calculate the wavelet amplitudes did not span the range of k values considered in the power measurement. For these reasons, our work here varies considerably from the work done by Gaikwad et al. (2021). It is likely that combining the power spectrum and wavelet amplitude PDFs when there isn’t a full correspondence between s_n and k would lead to an improvement of the measurement from either statistic alone. However, as long as some k and s_n values overlap we expect there to be non-negligible correlations that were been ignored in Gaikwad et al. (2021). The two additional statistics are also likely to improve the combined measurement beyond the combination of only power spectrum and wavelet amplitude PDFs. Adding these statistics to the work presented here would require calculating the correlations between the different statistics as we did between the power spectrum and wavelet amplitude PDFs. Calculating these statistics and exploring the relevant correlations is beyond the scope of this work. For these reasons we can not precisely estimate the correct size of the error bars from their combined measurement.

We did consider the effect of ignoring the correlations between PDFs from different smoothing scales on the posterior for the wavelet amplitude PDFs alone in Section 5.3 and Figure 15. We found that ignoring these correlations caused the posterior to shift and underestimate the errors (the orange and green lines in Figure 15). Using our inference test in Table 4, for the wavelet amplitude PDFs with “PDF bin correlations” only the true value of T_0 fell in the 1σ region 55%. This would imply the need to grow the 1σ region by a factor of $68/55 \sim 1.24$ or a 24% increase. It is therefore likely that the errors on the measurement from only the wavelet amplitudes in Gaikwad et al. (2021) are underestimated.

Overall, Figures 9, 10, 11, and 12 show that the off-diagonal terms in the covariance matrix are non-negligible and should be included in future analysis using wavelet amplitudes in order to achieve accurate error estimates. Figures 15 and 16 and Section 5.3 additionally demonstrate and discuss the effects of ignoring these correlations on the posteriors.

6 CONCLUSION

We have expanded upon the wavelet analysis methods used by Lidz et al. (2010) and Gaikwad et al. (2021) to study the thermal state of the IGM. Our method combines fifteen wavelet amplitude PDFs with smoothing scales that span the full range of scales probed by $P_F(k)$ and, for the first time, provides a full accounting of the correlations between these PDFs. We also calculated $P_F(k)$ from the

same simulated data in order to compare the precision of measurements on T_0 from these statistics. In order to rigorously combine the wavelet amplitude PDFs and power spectrum, we calculated the cross-correlations between $P_F(k)$ and the wavelet amplitude PDFs. We presented examples of each of these correlation matrices in Figures 9, 10, 11, and 12. Figures 10 and 12 showed the non-negligible off-diagonal correlations between the different smoothing scales and the different statistics. With our method at $z = 5$, the posterior of T_0 using the wavelet amplitude PDFs is on average 7% more precise than the power spectrum measurement on the same data. This means getting the same precision measurement with the power spectrum requires $\sim 15\%$ more data. Combining the power spectrum and wavelet amplitude PDFs did not significantly improve in precision of the posterior on T_0 over that from the wavelet amplitude PDFs alone, indicating that they contain the same information. At $z = 6$ we found that the posterior of T_0 using the wavelet amplitude PDFs is on average 12% more precise which would require $\sim 15\%$ more data to achieve the same accuracy with the power spectrum. Additionally, we calculated posteriors on T_0 at $z = 5$ with covariance matrices that ignored the off-diagonal correlations between PDF bins, between smoothing scales, and between the different statistics in Figures 15 and 16. We were unable to pass an inference test with these posteriors (as reported in Table 4) which implies that the errors are underestimated in these cases. This further demonstrated the significance of the off-diagonal terms in the covariance matrices and that they must be computed for a robust statistical analysis.

Here we adopted a simple model of the thermal state of the IGM which depended on a single parameter T_0 . For the more common and general case of multiple model parameters, the wavelet amplitude PDFs have even greater potential to better constrain model parameters when compared to the 1D flux power spectrum. To reiterate, the wavelet amplitude PDF characterizes the full wavelet distribution while the 1D flux power spectrum contains information on the mean of the wavelet amplitude PDF, see Equation (7). If we are only varying one model parameter and this parameter shifts the wavelet amplitude PDFs, the mean of the wavelet PDF may effectively contain all the information on the differences of the model. This is true in our thermal model with T_0 , as can be seen in Figures 5 and 8. In more sophisticated models, like those of reionization (Oñorbe et al. 2019; Boera et al. 2019) and WDM (Viel et al. 2013; Iršič et al. 2017), we would want to vary multiple model parameters (such as γ , $\langle F \rangle$, and m_{WDM}). Multiple model parameters are likely to cause changes in the full distribution of wavelet amplitudes beyond shifts in the mean. Wavelet amplitude PDF are sensitive to these additional changes while the power spectrum is not, meaning the wavelet amplitude PDFs could better discriminate between models to an even greater extent than they do for only one model parameter when compared to the power spectrum. Investigation of models with multiple thermal

parameters is beyond the scope of this paper but is a more realistic and promising area to explore wavelet analysis.

Wavelets are an independent statistic that can be used to probe the small-scale structure of the IGM through the Ly α forest. They can be used as a check on alternative statistics such as the power spectrum, Doppler parameter distribution, and curvature statistics since, in principle, each statistic may be sensitive to different systematics. In addition, the wavelet amplitude PDFs are higher precision than the power spectrum. Wavelets have an added benefit of providing Fourier information in configuration space which may be useful in other areas, such as in quasar proximity zones (Khrykin et al. 2016) or when looking for temperature fluctuations in the IGM (Theuns & Zaroubi 2000; Theuns et al. 2002a; Zaldarriaga 2002; Fang & White 2004; Lai et al. 2006; McQuinn et al. 2011). We have not studied the effects of a late ending reionization with remaining temperature fluctuations or a varying UVB background on the shape of the wavelet PDFs but this is an interesting area to explore in the future (Davies & Furlanetto 2016; Becker et al. 2018; Kulkarni et al. 2019; Keating et al. 2020; Nasir & D'Aloisio 2020; Gaikwad et al. 2020).

Implementing wavelet analysis can be challenging due to the large size of data vectors and covariance matrices involved. Often the cross correlations are ignored (Lidz et al. 2010; Gaikwad et al. 2021) which can lead to inaccurate parameter constraints. It is also impossible to remove the effect of both the resolution and the noise on the full wavelet amplitude PDFs, meaning you have to forward-model these effects, unlike in power spectrum analysis where the noise can be subtracted and the resolution effects removed by a window function correction. When studying data, we would calculate the covariance matrix from bootstrapping the data itself as has been done when using other statistics of the Ly α forest (see Boera et al. (2019) for example). This would reduce the computational time required since we would not be need to compute the covariance matrix for each model.

An interesting subject for future work will be to build an emulator using wavelet PDFs analogous to Ly α forest power spectrum emulators (see, e.g. Walther et al. (2019)). One issue of concern for wavelets is the large number of functions that need to be emulated (15 wavelets versus 1 power spectrum), although it could be possible to simply emulate the wavelet likelihood which is a single function. The emulation field has shifted towards iterative sampling, informed by the posterior probability distribution for a given observational dataset (Rogers et al. 2019; Takhtaganov et al. 2021), making emulating the likelihood consistent with current methods.

In the future, our method of wavelet analysis can be applied to quasar data and more sophisticated simulations to obtain precision constraints on the thermal state of the IGM. One can expand our analysis to constrain the timing of reionization as well as models of dark matter. As mentioned above, our approach can also be adapted to analyses in proximity zones or searches for IGM temperature fluctuations exploiting the space-preserving properties of wavelets.

ACKNOWLEDGEMENTS

We acknowledge helpful conversations with the ENIGMA group at UC Santa Barbara. JFH acknowledges support from the National Science Foundation under Grant No. 1816006. We acknowledge PRACE for awarding us access to JUWELS at GCS@JSC, Germany.

DATA AVAILABILITY

The simulation data analyzed in this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Almgren A. S., Bell J. B., Lijewski M. J., Lukić Z., Van Andel E., 2013, *ApJ*, **765**, 39
- Becker G. D., Rauch M., Sargent W. L. W., 2007, *ApJ*, **662**, 72
- Becker G. D., Bolton J. S., Haehnelt M. G., Sargent W. L. W., 2011, *MNRAS*, **410**, 1096
- Becker G. D., Bolton J. S., Madau P., Pettini M., Ryan-Weber E. V., Venemans B. P., 2015, *MNRAS*, **447**, 3402
- Becker G. D., Davies F. B., Furlanetto S. R., Malkan M. A., Boera E., Douglass C., 2018, *ApJ*, **863**, 92
- Boera E., Murphy M. T., Becker G. D., Bolton J. S., 2014, *MNRAS*, **441**, 1916
- Boera E., Becker G. D., Bolton J. S., Nasir F., 2019, *ApJ*, **872**, 101
- Bolton J. S., Viel M., Kim T.-S., Haehnelt M. G., Carswell R. F., 2008, *MNRAS*, **386**, 1131
- Bolton J. S., Becker G. D., Wyithe J. S. B., Haehnelt M. G., Sargent W. L. W., 2010, *MNRAS*, **406**, 612
- Bolton J. S., Becker G. D., Raskutti S., Wyithe J. S. B., Haehnelt M. G., Sargent W. L. W., 2012, *MNRAS*, **419**, 2880
- Bolton J. S., Becker G. D., Haehnelt M. G., Viel M., 2014, *MNRAS*, **438**, 2499
- Bryan G. L., Machacek M. E., 2000, *ApJ*, **534**, 57
- Calura F., Tescari E., D'Odorico V., Viel M., Cristiani S., Kim T. S., Bolton J. S., 2012, *MNRAS*, **422**, 3019
- D'Aloisio A., McQuinn M., Davies F. B., Furlanetto S. R., 2018, *MNRAS*, **473**, 560
- D'Aloisio A., McQuinn M., Maupin O., Davies F. B., Trac H., Fuller S., Sanderbeck P. R. U., 2019, *ApJ*, **874**, 154
- Davies F. B., Furlanetto S. R., 2016, *MNRAS*, **460**, 1328
- Davies F. B., Furlanetto S. R., McQuinn M., 2016, *MNRAS*, **457**, 3006
- Eilers A.-C., Hennawi J. F., Lee K.-G., 2017, *ApJ*, **844**, 136
- Fan X., et al., 2006, *AJ*, **132**, 117
- Fang T., White M., 2004, *ApJ*, **606**, L9
- Gaikwad P., et al., 2020, *MNRAS*, **494**, 5091
- Gaikwad P., Srianand R., Haehnelt M. G., Choudhury T. R., 2021, *MNRAS*, **506**, 4389
- Garzilli A., Bolton J. S., Kim T. S., Leach S., Viel M., 2012, *MNRAS*, **424**, 1723
- Garzilli A., Boyarsky A., Ruchayskiy O., 2017, *Physics Letters B*, **773**, 258
- Gnedin N. Y., Hui L., 1998, *MNRAS*, **296**, 44
- Gunn J. E., Peterson B. A., 1965, *ApJ*, **142**, 1633
- Haehnelt M. G., Steinmetz M., 1998, *MNRAS*, **298**, L21
- Hiss H., Walther M., Hennawi J. F., Oñorbe J., O'Meara J. M., Rorai A., Lukić Z., 2018, *ApJ*, **865**, 42
- Hui L., Gnedin N. Y., 1997, *MNRAS*, **292**, 27
- Hui L., Haiman Z., 2003, *ApJ*, **596**, 9
- Iršič V., et al., 2017, *Phys. Rev. D*, **96**, 023522
- Keating L. C., Weinberger L. H., Kulkarni G., Haehnelt M. G., Chardin J., Aubert D., 2020, *MNRAS*, **491**, 1736
- Khrykin I. S., Hennawi J. F., McQuinn M., Worsek G., 2016, *ApJ*, **824**, 133
- Kulkarni G., Hennawi J. F., Oñorbe J., Rorai A., Springel V., 2015, *ApJ*, **812**, 30
- Kulkarni G., Keating L. C., Haehnelt M. G., Bosman S. E. I., Puchwein E., Chardin J., Aubert D., 2019, *MNRAS*, **485**, L24
- Lai K., Lidz A., Hernquist L., Zaldarriaga M., 2006, *ApJ*, **644**, 61
- Lee K.-G., et al., 2015, *ApJ*, **799**, 196
- Lidz A., Malloy M., 2014, *ApJ*, **788**, 175
- Lidz A., Heitmann K., Hui L., Habib S., Rauch M., Sargent W. L. W., 2006, *ApJ*, **638**, 27–44
- Lidz A., Faucher-Giguère C.-A., Dall'Aglio A., McQuinn M., Fechner C., Zaldarriaga M., Hernquist L., Dutta S., 2010, *ApJ*, **718**, 199

- Lukić Z., Stark C. W., Nugent P., White M., Meiksin A. A., Almgren A., 2015, *MNRAS*, **446**, 3697
- Lynds R., 1971, *ApJ*, **164**, L73
- McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., 2001, *ApJ*, **562**, 52
- McQuinn M., 2012, *MNRAS*, **426**, 1349
- McQuinn M., Upton Sanderbeck P. R., 2016, *MNRAS*, **456**, 47
- McQuinn M., Hernquist L., Lidz A., Zaldarriaga M., 2011, *MNRAS*, **415**, 977
- Meiksin A., 2000, *MNRAS*, **314**, 566
- Miralda-Escudé J., Rees M. J., 1994, *MNRAS*, **266**, 343
- Narayanan V. K., Spergel D. N., Davé R., Ma C.-P., 2000, *ApJ*, **543**, L103–L106
- Nasir F., D’Aloisio A., 2020, *MNRAS*, **494**, 3080
- Nasir F., Bolton J. S., Becker G. D., 2016, *MNRAS*, **463**, 2335
- Oñorbe J., Hennawi J. F., Lukić Z., 2017, *ApJ*, **837**, 106
- Oñorbe J., Hennawi J. F., Lukić Z., Walther M., 2017, *ApJ*, **847**, 63
- Oñorbe J., Davies F. B., Lukić Z., Hennawi J. F., Sorini D., 2019, *MNRAS*, **486**, 4075
- Planck Collaboration et al., 2020, *A&A*, **641**, A6
- Puchwein E., Bolton J. S., Haehnelt M. G., Madau P., Becker G. D., Haardt F., 2015, *MNRAS*, **450**, 4081
- Rauch M., 1998, *ARA&A*, **36**, 267–316
- Ricotti M., Gnedin N. Y., Shull J. M., 2000, *ApJ*, **534**, 41
- Rogers K. K., Peiris H. V., Pontzen A., Bird S., Verde L., Font-Ribera A., 2019, *J. Cosmology Astropart. Phys.*, **2019**, 031
- Rorai A., et al., 2017, *Science*, **356**, 418
- Rorai A., Carswell R. F., Haehnelt M. G., Becker G. D., Bolton J. S., Murphy M. T., 2018, *MNRAS*, **474**, 2871
- Rudie G. C., Steidel C. C., Pettini M., 2012, *ApJ*, **757**, L30
- Schaye J., Theuns T., Leonard A., Efstathiou G., 1999, *MNRAS*, **310**, 57
- Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, **318**, 817
- Takhtaganov T., Lukić Z., Müller J., Morozov D., 2021, *ApJ*, **906**, 74
- Theuns T., Zaroubi S., 2000, *MNRAS*, **317**, 989
- Theuns T., Schaye J., Haehnelt M. G., 2000, *MNRAS*, **315**, 600
- Theuns T., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell R. F., 2002a, *MNRAS*, **332**, 367
- Theuns T., Schaye J., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell B., 2002b, *ApJ*, **567**, L103
- Viel M., Bolton J. S., Haehnelt M. G., 2009, *MNRAS: Letters*, **399**, L39
- Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, *Phys. Rev. D*, **88**
- Walther M., Hennawi J. F., Hiss H., Oñorbe J., Lee K.-G., Rorai A., O’Meara J., 2017, *ApJ*, **852**, 22
- Walther M., Oñorbe J., Hennawi J. F., Lukić Z., 2019, *ApJ*, **872**, 13
- Yèche C., Palanque-Delabrouille N., Baur J., du Mas des Bourboux H., 2017, *J. Cosmology Astropart. Phys.*, **2017**, 047
- Zaldarriaga M., 2002, *ApJ*, **564**, 153
- Zaldarriaga M., Hui L., Tegmark M., 2001, *ApJ*, **557**, 519

APPENDIX A: LIKELIHOOD CHOICE

We chose to use a multivariate Gaussian distribution for the likelihood of our data. This assumption explicitly means that if we take multiple mock data sets and look at the distribution of a single point in our data vector (either a single bin from the wavelet amplitude histogram or a single k band from power spectra) it will be Gaussian. It also assumes that when looking at any two points from the data vector, we expect the resulting distribution to be a two-dimensional Gaussian (thus looking at many points will be a multi-dimensional Gaussian).

We can visually check the assumption that any two points from the data vector will result in a two-dimensional Gaussian distribution over many mocks. We will show this for one point in the power spectrum ($k = 0.12 \text{ s km}^{-1}$) and one point in the wavelet amplitude

Figure A1. The distribution of mock draws for the power spectrum at $k = 0.12 \text{ s km}^{-1}$ and the wavelet amplitude PDF at $A_n = 0.32$. The bottom left panel shows the 2D distribution of these bins where the red ellipse shows the 3σ region calculated from the covariance matrix for these two bins. The bottom right panel shows the distribution of values only for the power spectrum. The top left panel shows the distribution of values only for the wavelet amplitude PDF. All panels show good agreement with the assumption of a multi-variate Gaussian distribution.

PDF for $s_n = 51.09 \text{ km s}^{-1}$ ($A_n = 0.32$). The distribution of these values for 1,000 mock draws is shown in Figure A1 along with the distributions of the individual Histogram and $P_F(k)$ values. The red ellipse represents the 3σ contour from the 2×2 covariance matrix calculated for these two bins. The two dimensional distribution agrees very well with the ellipse by eye and the two histograms also visually appear Gaussian within the errors expected from counting statistics.

APPENDIX B: COVARIANCE MATRIX CALCULATION

As defined by equation (10), each model covariance matrix is calculated from mock draws of the data. This will inherently be a noisy calculation that will converge as $1/\sqrt{N}$ where N is the number of draws in the covariance matrix. As stated in the text, we used 1,000,000 mock draws when calculating the covariance matrix for the wavelet amplitude PDFs as well as the combination of the wavelet amplitude PDFs and the power spectrum. To check that 1,000,000 mock draws are sufficient to minimize the error from this calculation, we looked at the behavior of elements of the wavelet amplitude PDF covariance matrix in Figure B1. The values in the plot have been normalized such that at 10^6 draws they are 1. The four elements have been chosen such that there is one on the diagonal, one off the diagonal where there are strong correlations, and two off the diagonal where there are weak correlations. This plot shows that as we approach 10^6 draws the values vary significantly less than they do at lower values and thus the covariance elements are converging.

For both the wavelet amplitude PDFs and the combined power spectrum and wavelet amplitude PDFs covariance matrix, the data

Figure B1. This figure shows the values of four distinct elements in the covariance matrix for the wavelet amplitude PDFs as we increased the number of mock draws used to calculate the covariance matrix. For simplicity, these distinct values are labeled by their index, rather than the smoothing scale and wavelet amplitude values associated for each bin in the PDFs. These points were chosen such that we had one on the diagonal, one off the diagonal where there are strong correlations, and two off the diagonal where there are weak correlations. As we approach 10^6 draws the values converge, showing that 10^6 is sufficient for our covariance matrix calculation.

vector is long enough that there are many elements with very small cross correlations. These small values vary more (as seen in the (134, 23) and (134, 121) elements in Figure B1) such that they can still have non-negligible noise. This noise in the covariance leads to noise in the posterior measurement on T_0 (as discussed in Sections 4.2 and 4.3). At this point, for computational reasons, we decided not to increase the number of mock draws of data. Instead, we chose to smooth the covariance matrix across our thermal grid with a spline. We did this by fitting each individual element of the covariance matrix to spline with 20 equally spaced breakpoints. This results in 150×150 (165×165) splines for the wavelet amplitude PDFs (combined) covariance matrix. We chose 20 breakpoints to allow the spline to be flexible enough to find real patterns in the data while smoothing out noise.

We show one example of this spline in Figure B2. This is the spline for one of the elements of the covariance matrix shown in Figures 11 and 12. Specifically, this is for $k = 0.41 \text{ s km}^{-1}$ and $A_n = 0.027$. This figure shows how the spline can replicate patterns in the calculated covariance matrix values while still reducing noise. For example, with the noisy values near $\log(T_0) = 4.2$.

The covariance matrices used in equation (9) for both the wavelet analysis and the combined wavelet and power spectrum analysis use the values of the spline at every T_0 .

APPENDIX C: REDSHIFT 6

Here we have included the figures for a mock data set at $z = 6$. As a reminder, at $z = 6$ we considered $\langle F \rangle = 0.011$, $R = 30,000$, and $\text{SNR} = 35$. We restricted our study to a higher SNR because the lower observed flux makes the noise power more dominant, see equation (14) for details. The mock data set still considers 8 quasars (equivalent to 29 simulation skewers).

First, in Figure C1 we show a mock power spectrum at $\log(T_0) = 4.1625$ with three models analogous to Figure 6 at $z = 5$. The error bars blow up at the largest values of k due to removing the noise power which is dominant at these small scales as well as correct-

Figure B2. The covariance value for the same bin of the covariance matrix at different T_0 . The bin chosen here corresponds to the histogram bin of $A_n = 0.027$ for $s_n = 51.09 \text{ km s}^{-1}$ and the power spectrum band $k = 0.41 \text{ s km}^{-1}$ and $A_n = 0.027$. The solid line shows the spline fit to these points with 20 equally spaced breakpoints.

Figure C1. The power spectrum measurement, $P_F(k)$, for one mock data set with $\log(T_0) = 4.1625$ (black points) at $z = 6$. The 1σ error bars are calculated from the square root of the diagonal of the covariance matrix. Also shown are model values of the power spectra for three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green).

ing the window function. This Figure also shows three model power spectra at $z = 6$: $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). The corresponding correlation plot for $\log(T_0) = 4.1625$ is shown in Figure C2. The structure here looks quite different than the one at $z = 5$ shown in Figure 7. The correlations come from underlying correlations in the high- z Ly α forest with the overall low mean flux. Again, the off-diagonal elements for high- k (small scales) are very small due to the dominance of the noise power over the signal at these scales.

Next, we look at the figures relevant for the wavelet amplitude PDFs at $z = 6$. The mock wavelet amplitude PDFs at all scales considered are shown in Figure C3. Note that the scales considered here are slightly different than those at $z = 5$ (as listed in table 1 and Figure 8) because the size of the skewers and the nyquist frequency vary at these redshifts. At both redshifts we chose 15 logarithmically spaced values for s_n (and k). At this redshift and signal to noise ratio, we can see distinct bumps for the contribution of noise and wavelet amplitudes for $290.75 \text{ km s}^{-1} \geq s_n \geq 55.09 \text{ km s}^{-1}$. Again we have also shown three model wavelet amplitude PDFs: $\log(T_0) = 3.4$

Figure C2. The correlation matrix for the power spectrum at $\log(T_0) = 4.1625$ and $z = 6$. The very weak correlations seen in the regions where $k > 0.2 \text{ s km}^{-1}$ are due to uncorrelated random Gaussian noise which dominates the signal on small-scales (high k).

(blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). These wavelet amplitude PDFs all agree with each other for the smallest scales $s_n < 23.98 \text{ km s}^{-1}$ where noise dominates the PDF which is the same for all temperature models. The correlation for all wavelet amplitude PDFs at $\log(T_0) = 4.1625$ is shown in Figure C4. The 10×10 diagonal blocks for the largest scales ($s_n > 440.7 \text{ km s}^{-1}$) and the smallest scales ($36.34 \text{ km s}^{-1} > s_n$) look similar to that shown for $z = 5$ in Figure 10 since the PDFs also have the same shape. The 10×10 diagonal blocks for the mid-range values of s_n are different from the others due to the combined PDF shape as seen in Figure C3. Again the off-diagonal blocks show a similar pattern to the diagonal blocks modified by positive or negative numbers, mimicking the off-diagonal correlations from the power at $z = 6$ seen in Figure C2. The off-diagonal correlation values between the $s_n \leq 23.98 \text{ km s}^{-1}$ and all other values of s_n show no strong correlations because the uncorrelated noise dominates at these small scales.

At this redshift we chose not to investigate the combination of the wavelet amplitude PDFs and the power spectrum due to the computational time required and the results at $z = 5$ which showed no significant improvement from combining these statistics when the same scales are covered in both.

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

Figure C3. The black points show the PDFs from one mock data set for each s_n with $\log(T_0) = 4.1625$ and $z = 6$. The 1σ error bars are calculated from the square root of the diagonal of the covariance matrix. Each panel also shows the "model" values of the PDFs from the stated smoothing scale for three different values of T_0 : $\log(T_0) = 3.4$ (blue), $\log(T_0) = 4.1625$ (orange), and $\log(T_0) = 4.4$ (green). This redshift shows broad PDFs with arguable two bumps in the mid-range values of s_n where the flux and noise power levels are comparable.

Figure C4. The correlation matrix for fifteen wavelet amplitude PDFs at $\log(T_0) = 4.1625$ and $z = 6$. The wavelet amplitude PDFs for large smoothing scales, $2326 \text{ km s}^{-1} \geq s_n \geq 36.34 \text{ km s}^{-1}$, have significant correlations off the diagonal. The off diagonal blocks show a similar repeating shape to those on the diagonal modified by numbers. These numbers are positive close to the diagonal and are negative further from the diagonal. For $s_n \leq 29.98 \text{ km s}^{-1}$, the wavelet amplitudes begin to be dominated by noise, so the correlations between PDFs for different values of s_n become very small. This pattern mimics that seen in the power spectrum correlation shown in Figure C2. The pattern of the diagonal blocks are the most different for the mid range values of $290.7 \text{ km s}^{-1} \geq s_n \geq 83.5 \text{ km s}^{-1}$, which is where the PDF shapes are the most different from the typical shape, as seen in Figure C3.