





Galaxies Probing Galaxies in PRIMUS. II. The Coherence Scale of the Cool Circumgalactic Medium

Kate H. R. Rubin¹ , Aleksandar M. Diamond-Stanic², Alison L. Coil³ , Neil H. M. Crighton⁴, and Kyle R. Stewart⁵

¹San Diego State University, Department of Astronomy, San Diego, CA 92182, USA; krubin@sdsu.edu

²Bates College, Department of Physics & Astronomy, 44 Campus Ave., Carnegie Science Hall, Lewiston, ME 04240, USA

³Center for Astrophysics and Space Sciences, Department of Physics, University of California, San Diego, 9500 Gilman Dr., La Jolla, CA 92093, USA

⁴Center for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia

⁵Department of Mathematical Sciences, California Baptist University, 8432 Magnolia Ave., Riverside, CA 92504, USA

Received 2017 November 29; revised 2018 June 21; accepted 2018 July 21; published 2018 December 3

Abstract

The circumgalactic medium (CGM) close to $\sim L^*$ star-forming galaxies hosts strong Mg II $\lambda 2796$ absorption (with equivalent width $W_{2796} > 0.1 \text{ \AA}$) with a near-unity covering fraction. To characterize the spatial coherence of this absorption, we analyze the W_{2796} distribution in the CGM of 27 star-forming galaxies detected in deep spectroscopy of bright background (b/g) galaxies first presented in Rubin et al. The sample foreground (f/g) systems have redshifts $0.35 \lesssim z \lesssim 0.8$ and stellar masses $9.1 < \log M_*/M_\odot < 11.1$, and the b/g galaxies provide spatially extended probes with half-light radii $1.0 \text{ kpc} \lesssim R_{\text{eff}} \lesssim 7.9 \text{ kpc}$ at projected distances $R_\perp < 50 \text{ kpc}$. Our analysis also draws on literature W_{2796} values measured in b/g QSO spectroscopy probing the halos of f/g galaxies with a similar range in M_* at $z \sim 0.25$. By making the assumptions that (1) samples of like galaxies exhibit similar circumgalactic W_{2796} distributions and, (2) within a given halo, the quantity $\log W_{2796}$ has a Gaussian distribution with a dispersion that is constant with M_* and R_\perp , we use this QSO–galaxy pair sample to construct a model for the $\log W_{2796}$ distribution in the CGM of low-redshift galaxies. Adopting this model, we then demonstrate the dependence of the observed $\log W_{2796}$ distribution on the ratio of the surface area of the b/g probe to the projected absorber surface area ($x_A \equiv A_G/A_A$), finding that distributions that assume $x_A \geq 15$ are statistically inconsistent with that observed toward our b/g galaxy sample at a 95% confidence level. This limit, in combination with the b/g galaxy sizes, requires that the length scale over which W_{2796} does not vary (i.e., the “coherence scale” of Mg II absorption) is $\ell_A > 1.9 \text{ kpc}$. This novel constraint on the morphology of cool, photoionized structures in the inner CGM suggests that either these structures each extend over kiloparsec scales or the number and velocity dispersion of these structures are spatially correlated over the same scales.

Key words: galaxies: halos – quasars: absorption lines

1. Introduction

Within the last decade, quasi-stellar object (QSO) absorption-line experiments have revealed the gaseous material enveloping low-redshift galaxies to be a dominant component of their host halo’s baryonic mass. Galvanized by the unprecedented sensitivity of the Cosmic Origins Spectrograph on the *Hubble Space Telescope*, these studies report a mass of cool (temperature $T \sim 10^4 \text{ K}$), diffuse baryons of nearly $10^{11} M_\odot$ filling the regions extending to 160 kpc from isolated, $\sim L^*$ galaxies (Stocke et al. 2013; Peeples et al. 2014; Werk et al. 2014; Prochaska et al. 2017). Simultaneous observations of the highly ionized metal species O VI have been interpreted to indicate the presence of another, warmer gas phase at $T \gtrsim 10^5 \text{ K}$, estimated to contain a mass of more than $10^9 M_\odot$ (Tumlinson et al. 2011; Prochaska et al. 2017). Moreover, an enduring prediction of galaxy formation theory is the shock heating of gas as it falls onto dark matter halos (Rees & Ostriker 1977; White & Rees 1978; Kereš et al. 2005; Nelson et al. 2013), resulting in a ubiquitous “hot” phase ($T \sim 10^6 \text{ K}$) filling halos with masses $\gtrsim 10^{11} M_\odot$ (Birnboim & Dekel 2003; Kereš et al. 2009). These empirical and theoretical findings imply the omnipresence of a massive gas reservoir composed of material over a broad range of temperatures surrounding luminous galaxies in the nearby universe.

The predominance of this baryonic component in turn implies a crucial role in the regulation of galaxy growth. Indeed, hydrodynamical simulations of galaxy formation predict that this reservoir is fed by the accretion of material

from the intergalactic medium, the stripping of satellite galaxies as they merge with the central massive host, and large-scale outflows of gas driven from star-forming regions via feedback processes (e.g., Oppenheimer et al. 2010; Shen et al. 2013; Ford et al. 2014; Hopkins et al. 2014). At the same time, these simulations predict distinct spatial distributions and morphologies for each gas phase. The O VI–absorbing material, for instance, typically exhibits a relatively smooth morphology extending well beyond the virial radius of a halo of mass $M_h \sim 10^{12} M_\odot$ (e.g., Shen et al. 2013; Oppenheimer et al. 2016), while optically thick H I and gas traced by absorption in low-ionization metal transitions (e.g., Si II, C II) is distributed in narrow filaments or small clumps (Shen et al. 2013; Fumagalli et al. 2014; Faucher-Giguère et al. 2015). These morphologies may ultimately be linked to the physical origin of each phase and therefore may potentially corroborate interpretations based on other factors (e.g., metallicity or kinematics).

Moreover, a characterization of the detailed structure of circumgalactic material is crucial to our understanding of its hydrodynamics (Crighton et al. 2015). As the region through which gas accretes onto galaxies and the reservoir receiving galactic wind ejecta, the circumgalactic medium (CGM) cannot be understood as a static gaseous body (e.g., Werk et al. 2014; Fielding et al. 2017; Oppenheimer et al. 2018). The best estimates of the volume density of the photoionized phase indicate that it is too rarified to be in pressure equilibrium with a virialized hot gas halo (Werk et al. 2014). Furthermore, it is

predicted that such cool “clumps” will be susceptible to Rayleigh–Taylor and Kelvin–Helmholtz instabilities as they travel through the surrounding hot medium (Schaye et al. 2007; Heitsch & Putman 2009; Joung et al. 2012; Crighton et al. 2015; McCourt et al. 2015; Armillotta et al. 2017). In the absence of additional stabilizing mechanisms, hydrodynamical simulations predict that these clumps are almost completely disrupted on timescales similar to the cloud-crushing time,

$$t_{\text{crush}} \sim \left(\frac{\rho_{\text{clump}}}{\rho_{\text{hot}}} \right)^{1/2} \frac{R_{\text{clump}}}{v_{\text{clump}}} \quad (1)$$

(Jones et al. 1994; Heitsch & Putman 2009; Joung et al. 2012; Schneider & Robertson 2017), with $(\rho_{\text{clump}}/\rho_{\text{hot}})$ the ratio of the densities of the two phases, v_{clump} the relative velocity between the cool and hot media, and R_{clump} the size of the cool structure. Taken at face value, this relation indicates that the survival time of the cool phase has a strong dependence on both its size and kinematics.

In spite of their apparent importance, however, empirical constraints on the sizes and morphologies of the phases of the CGM are very few. Photoionization modeling of absorption systems detected along QSO sightlines with, e.g., the CLOUDY spectral synthesis code (Ferland et al. 1998) can, in principle, constrain the extent of the structure along the line of sight. However, such modeling requires that several ionic metal transitions be observed, and the results are sensitive to the user’s assumptions regarding the background (b/g) ionizing radiation field. Even in analyses of some of the richest absorption-line data sets with coverage of numerous ionic species, the uncertainty in such size estimates is typically at least an order of magnitude (Werk et al. 2014). Alternatively, multiple images of gravitationally lensed QSOs can have transverse separations ranging from less than a kiloparsec to >100 kpc scales, and if aligned behind foreground (f/g) absorption, they can provide unique comparisons of the velocity structure and strength of the system over these scales (e.g., Weymann & Foltz 1983; Smette et al. 1995; Monier et al. 1998; Rauch et al. 2002; Ellison et al. 2004; Chen et al. 2014). However, such special sightlines are very rare on the sky, particularly if one also demands the presence of a close transverse f/g galaxy whose redshift is known a priori (Chen et al. 2010a, 2014; Zahedy et al. 2016).

Spectroscopy of spatially extended b/g sources, such as bright b/g galaxies, can also constrain the sizes of f/g absorbers, especially when analyzed in tandem with complementary QSO absorption spectroscopy (e.g., probing similar f/g systems). Galaxies that are sufficiently bright to enable high-S/N spectroscopic coverage of f/g metal-line transitions in the near-UV are rare; however, dense galaxy redshift surveys can facilitate the selection of statistical samples of such exceptional objects. When identified close in projection to f/g galaxies with known redshifts, near-UV spectroscopy of the b/g sightlines may be used to target the coherence scale of the cool, photoionized phase of the CGM in metal-line absorption transitions (e.g., Mg II $\lambda\lambda 2796, 2803$). b/g–f/g galaxy pair spectroscopy has indeed been presented in several previous studies (e.g., Adelberger et al. 2005; Barger et al. 2008; Rubin et al. 2010; Steidel et al. 2010; Bordoloi et al. 2011; Cooke & O’Meara 2015; Diamond-Stanic et al. 2016; Lee et al. 2016; Lopez et al. 2018; Peroux et al. 2018). However, until recently,

no study has achieved the S/N necessary to assess absorption equivalent widths associated with the photoionized CGM in more than one or two individual f/g galaxy halos.

In Paper I of this series (Rubin et al. 2018; hereafter **GPG1**), we presented spectroscopy obtained with the Keck/Low-Resolution Imaging Spectrometer (LRIS) and the Very Large Telescope (VLT)/Focal Reducer/Low-dispersion Spectrograph 2 (FOR2) of 72 projected pairs of galaxies at $0.4 \lesssim z \lesssim 1.0$ drawn from the PRISM Multi-object Survey (PRISMUS; Coil et al. 2011; Cool et al. 2013). Each individual b/g galaxy spectrum is sufficiently deep to be sensitive to Mg II $\lambda 2796$ absorption with equivalent width $W_{2796} \gtrsim 0.5 \text{ \AA}$ and, as such, provides consequential constraints on CGM absorption associated with the corresponding f/g halo. The pairs probe impact parameters as large as $R_{\perp} < 150$ kpc, but over two-thirds of the sample has $R_{\perp} < 50$ kpc, thus probing regions known to exhibit the strongest Mg II absorption in complementary QSO sightline experiments (e.g., Chen et al. 2010a). The f/g galaxies in these pairs have stellar masses in the range $10^9 M_{\odot} < M_{*} < 10^{11.2} M_{\odot}$ and lie predominantly along the star-forming sequence.

We then used these data to examine the relation between W_{2796} and R_{\perp} , demonstrating a negative correlation between these quantities within $R_{\perp} < 50$ kpc. We explored the median relation between W_{2796} and the intrinsic properties of the f/g hosts, finding that greater W_{2796} arises around galaxies with higher star formation rates (SFRs) and/or M_{*} . Finally, we compared these measurements with a sample of W_{2796} values obtained from studies of projected QSO–galaxy pairs in the literature, finding that the median W_{2796} observed toward both b/g galaxies and b/g QSOs at a given impact parameter around f/g galaxies of similar M_{*} are statistically consistent.

In the present work, we turn our focus from the mean and median CGM absorption strengths to an examination of the dispersion in the W_{2796} distributions observed toward b/g galaxies and QSOs and to a detailed comparison of these distributions. As we derive, the dispersion in W_{2796} as a function of R_{\perp} and intrinsic host galaxy properties is dependent on the size of the b/g beam relative to that of the f/g absorber and may therefore constrain the latter quantity. We begin our comparison by making the assumption that the Mg II–absorbing CGM as probed by our b/g QSO sample is the universal, or “fiducial,” CGM—i.e., that we are observing the same median W_{2796} profile as a function of R_{\perp} and M_{*} toward both these QSOs and our PRIMUS b/g galaxies. In Section 2.1 below, we develop a simple model for this fiducial W_{2796} profile. Then, in Section 2.2, we explore the relationship between the measured dispersion in this profile and the size of a given b/g beam relative to the sizes of the f/g absorbers in this fiducial CGM. Readers interested in the constraint on absorber size implied by the level of dispersion in W_{2796} measured toward our b/g galaxy sample may wish to focus on Section 2.3. In Section 3, we discuss the implications of these results for the physical nature of the Mg II–absorbing CGM (Section 3.1), describe complementary constraints on its small-scale structure (Section 3.2), address the limitations of our analysis (Section 3.3), and discuss the lifetime and fate of this cool, photoionized material (Section 3.4). We offer concluding remarks in Section 4. We adopt a Λ CDM cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_M = 0.3$, and $\Omega_{\Lambda} = 0.7$.

2. The Coherence Scale of W_{2796}

Our goal is to perform a quantitative comparison of the dispersion in W_{2796} measurements obtained toward b/g QSOs and b/g galaxies to constrain the physical scale of the f/g absorption. Our approach rests on a key assumption: that W_{2796} profiles obtained by assembling large samples of projected QSO–galaxy (or galaxy–galaxy) pairs are representative of the “fiducial” CGM. That is, we assume that the dispersion in these W_{2796} measurements is driven by the spatial fluctuations in W_{2796} in this fiducial CGM, rather than global variations in CGM properties from one host galaxy to another (at a given M_* , SFR, R_\perp , etc.). This assumption has not yet been justified empirically; moreover, the recent findings of Lopez et al. (2018) suggest that this may overestimate the intrinsic W_{2796} dispersion slightly. Larger samples of galaxies with extended or multiple b/g sightlines (e.g., gravitationally lensed QSOs or galaxies; Chen et al. 2014; Zahedy et al. 2016; Lopez et al. 2018) are needed to validate this picture.

Given this starting point, it follows that (as mentioned in Section 1) the Mg II–absorbing CGM probed by our PRIMUS b/g galaxy sample is the same “fiducial” CGM probed by existing projected QSO–galaxy pair samples at a similar epoch. Analysis presented in GPG1 demonstrated that the W_{2796} observed toward b/g galaxies is larger around f/g hosts with higher SFR and/or M_* at a given R_\perp , in qualitative agreement with the results of projected QSO–galaxy pair studies. In addition, we demonstrated that the median W_{2796} observed toward our b/g galaxy sample is consistent with that observed along b/g QSOs probing f/g halos over the same range in stellar mass ($9.1 < \log M_*/M_\odot < 10.7$). Although these findings do not test the validity of our key assumption, they are at least compatible with the concept of a fiducial CGM. Moreover, this dependency of W_{2796} on intrinsic galaxy properties (as well as on R_\perp ; Chen et al. 2010a; Nielsen et al. 2013; GPG1) implies that when comparing the dispersion in W_{2796} among various projected pair samples, we must account for the differing R_\perp distributions and f/g galaxy properties of each data set.

2.1. A Fiducial Model for the Cool CGM

To facilitate this accounting, we start by developing a model for the relationship between $\log W_{2796}$, R_\perp , and M_* of the f/g host. This parameterization was first explored in Chen et al. (2010b), who found that the inclusion of M_* as an independent model variable significantly reduced the intrinsic scatter in the relation between $\log W_{2796}$ and $\log R_\perp$ among their sample of 71 W_{2796} measurements obtained from b/g QSO spectroscopy probing “isolated” f/g galaxy halos at $z \sim 0.25$ (over $9 \text{ kpc} \lesssim R_\perp \lesssim 170 \text{ kpc}$). The demonstration of a positive relation between W_{2796} and M_* among the samples discussed in GPG1 suggests that this type of model may have a similar effect in the present context. Given that GPG1 also presented evidence for larger W_{2796} around f/g hosts with larger SFR, a version of the model including a dependence on SFR (rather than M_*) might similarly reduce the intrinsic scatter in the $\log W_{2796}$ – R_\perp relation. However, because GPG1 did not identify a significantly stronger relationship between W_{2796} and one of these intrinsic host properties relative to the other, for simplicity, we choose to focus here on the potential dependence on M_* only. Larger data sets sampling the CGM of many more f/g hosts are needed to isolate the relationships between these two correlated quantities and W_{2796} . Recent studies

have also suggested that W_{2796} may depend on the azimuthal angle of the b/g sightline relative to the f/g galaxy (Bordoloi et al. 2011; Bouché et al. 2012; Kacprzak et al. 2012), implying an additional reduction of the true intrinsic scatter in the W_{2796} distribution. We test for evidence of this dependence in our QSO–galaxy pair sample in the Appendix, again concluding that a larger comparison data set is required before such a dependence can be productively incorporated into this fiducial CGM model.

Our model thus simply includes a linear dependence on both R_\perp and $\log M_*$ as follows:

$$\log \bar{W}_{2796} = b + m_1 R_\perp + m_2 (\log M_*/M_\odot - 10.3), \quad (2)$$

with $\log \bar{W}_{2796}$ representing the predicted absorption strength and the arbitrary offset of 10.3 chosen to be close to the median M_* of the relevant data sets (described in more detail below). We adopt a likelihood function similar to that used in Chen et al. (2010a),

$$\begin{aligned} \mathcal{L}(\bar{W}) = & \left(\prod_{i=1}^n \frac{1}{\sqrt{2\pi s_i^2}} \exp \left\{ -\frac{1}{2} \left[\frac{W_i - \bar{W}}{s_i} \right]^2 \right\} \right) \\ & \times \left(\prod_{i=1}^m \int_{-\infty}^{W'_i} \frac{dW'}{\sqrt{2\pi s_i^2}} \exp \left\{ -\frac{1}{2} \left[\frac{W'_i - \bar{W}}{s_i} \right]^2 \right\} \right), \end{aligned}$$

here with each value of \bar{W} equal to the value of $\log \bar{W}_{2796}$ given by the model at each f/g galaxy’s $R_{\perp,i}$ and $\log M_{*,i}$. As described in GPG1, the first product includes direct $\log W_{2796}$ measurements (W_i), and the second includes upper limits (W'_i). The Gaussian variance $s_i^2 = \sigma_i^2 + \sigma_C^2$, with σ_i representing measurement uncertainty and σ_C representing the intrinsic scatter in the relation.

We constrain the four free parameters of this model (m_1 , m_2 , b , and σ_C) by fitting a subset of the QSO–galaxy pair samples available in the literature chosen to define the “fiducial” CGM at this epoch. These include the same data sets adopted in GPG1, namely, those of Chen et al. (2010a, 2010b) and Werk et al. (2013). Both of these studies build their samples using f/g galaxies whose redshifts are known a priori. We cull these samples to include only QSO sightlines that pass within $R_\perp < 50 \text{ kpc}$ of the corresponding f/g host for consistency with our primary PRIMUS pair sample. In addition, because some evidence suggests that the W_{2796} profile of star-forming host halos differs from that of halos hosting early-type galaxies (e.g., Bordoloi et al. 2011), and because the f/g systems in our PRIMUS pair sample are mostly star-forming, we exclude any pairs with galaxies lying below the star-forming sequence as defined by Berti et al. (2017; see Equation (1) in GPG1). The final fiducial QSO–galaxy pair sample is shown in Figure 1 and, in total, includes 39 measurements from Chen et al. (2010b) and 11 measurements from Werk et al. (2013).

Following the methodology of GPG1, we use the Markov chain Monte Carlo technique to sample the posterior probability density function (PPDF) for our model given these data (using the software package `emcee`; Foreman-Mackey et al. 2013). We use uniform probability density priors over the ranges $-5.0 < m_1 < 5.0$, $-5.0 < m_2 < 5.0$, $-10.0 < b < 10.0$, and $-10.0 < \ln \sigma_C < 10.0$. As in GPG1, we find that chains generated by 100 “walkers” each taking 5000 steps thoroughly sample the PPDF. We again adopt the median and




Figure 1. The $\log W_{2796}$ vs. a linear combination of R_{\perp} and $\log M_*$ ($R_{\perp}^{\text{corr}} = R_{\perp} + (m_2/m_1)(\log M_*/M_{\odot} - 10.3)$) for the QSO–galaxy pair sample we use to define the parameters of our fiducial model for the Mg II–absorbing CGM. This includes pairs with f/g galaxies having $R_{\perp} < 50$ kpc and that lie on the star-forming sequence as defined by Bertl et al. (2017). The median redshift of these systems is $z = 0.25$. The best-fit values and ± 34 th percentile probability intervals for relevant model parameters are listed above the top x -axis. The quantity R_{\perp}^{corr} is calculated assuming m_1 and m_2 are equal to their best-fit values. The f/g galaxies having low, intermediate, and high values of SFR are indicated with magenta circles, orange stars, and green squares, respectively. The black line shows the best-fit linear relation, and the dark and light gray contours indicate the inner $\pm 34\%$ and $\pm 47.5\%$ of the locus of curves obtained from random draws from the PPDF of the model.

± 34 th percentiles of the marginalized PPDFs as the best-fit value and uncertainty for each parameter.

Figure 1 shows the outcome of this procedure. The best-fit parameter values with uncertainties are listed at the top of the plot. The x -axis shows a linear combination of the two independent variables, $R_{\perp}^{\text{corr}} = R_{\perp} + (m_2/m_1)(\log M_*/M_{\odot} - 10.3)$, with m_1 and m_2 set at their best-fit values. The black line shows the best-fit relation. The dark and light gray contours show the $\pm 1\sigma$ and 2σ uncertainty ranges in $\log \bar{W}_{2796}$, estimated by drawing 1000 random sets of parameters from the PPDF, calculating \bar{W} for each set, and then determining the inner ± 34 th and ± 47.5 th percentile values of \bar{W} as a function of R_{\perp}^{corr} .

The fitted slope m_2 is suggestive of a marginally statistically significant dependence on $\log M_*$, with $m_2 = 0.151^{+0.128}_{-0.126}$ larger than zero at a level of $\sim 1.2\sigma$. The best-fit value and uncertainty interval for σ_C , on the other hand, are indicative of a high level of intrinsic scatter around the best-fit linear relation. The data points are color-coded by SFR as indicated in the legend. The distributions of $\log W_{2796}$ values with respect to the best-fit relation are broadly consistent among these subsamples, suggesting that if there is an additional dependence of W_{2796} on SFR, this sample will not usefully constrain it.

We use this model and these best-fit parameters to define the form of and intrinsic scatter in $\log \bar{W}_{2796}$ as a function of R_{\perp} and $\log M_*$; i.e., our “fiducial” Mg II–absorbing CGM model. Our constraints on all of these parameters (including the intrinsic scatter, σ_C) will be leveraged in the analysis to follow. Here we remind the reader that we have assumed a Gaussian form for the variance in $\log \bar{W}_{2796}$, implying that the probability distribution of




Figure 2. Distribution of $\log W_{2796}$ values in our QSO–galaxy comparison sample in two bins of the quantity R_{\perp}^{corr} , with $R_{\perp}^{\text{corr}} = R_{\perp} + (m_2/m_1)(\log M_*/M_{\odot} - 10.3)$ and $m_2/m_1 = -10.8$. Two of these constraints on $\log W_{2796}$ are upper limits and are included in the histograms at the values of these limits. The blue and red curves are fits of Gaussian functions to the $\log W_{2796}$ distributions at $20 \text{ kpc} < R_{\perp}^{\text{corr}} < 30 \text{ kpc}$ and $30 \text{ kpc} < R_{\perp}^{\text{corr}} < 40 \text{ kpc}$, respectively. The Gaussians are shown only to demonstrate the form of the scatter in $\log W_{2796}$ that we assume. Their fitted parameters are not used in our analysis.

a measurement $\log W_{2796,i}$ at particular values of $R_{\perp,i}$ and $\log M_{*,i}$ is also a Gaussian centered at $\log \bar{W}_{2796,i}$ (Hogg et al. 2010). This assumption of Gaussianity in $\log W_{2796}$ (and hence lognormality in W_{2796}) has not yet been justified; however, we persist in this assumption for the following reasons. First, we lack the measurements needed to empirically constrain the form of the W_{2796} distribution, and we therefore consider the choice between normal and lognormal distributions arbitrary. Second, while negative values of W_{2796} may be measured in instances of noisy spectroscopy, the true W_{2796} due to diffuse Mg II ions f/g to a bright b/g source will always be nonnegative. A lognormal distribution is consistent with this constraint. A test of this assumption may, in principle, be performed as demonstrated in Figure 2; here we select subsamples of the QSO–galaxy comparison data set in two narrow ranges in the quantity $R_{\perp}^{\text{corr}} = R_{\perp} + (m_2/m_1)(\log M_*/M_{\odot} - 10.3)$ and show the distribution of $\log W_{2796}$ in each. We cannot use these subsamples to quantitatively disfavor a lognormal relative to a Gaussian model (or vice versa) for the W_{2796} distribution due to the small number of measurements. We encourage future studies with larger samples to address this issue.

2.2. Observing the Fiducial CGM toward Larger b/g Beams

The PPDF for the parameters of the fiducial model developed in the previous subsection may now be used to generate new samples of W_{2796} measurements that would be observed along b/g QSO sightlines given any arbitrary set of f/g galaxy M_* and impact parameter values. Our ultimate goal is to constrain the sizes of Mg II absorbers by comparing this model to our measurements of W_{2796} toward b/g galaxies (i.e., sources that emit UV continuum light over significantly larger areas than QSOs). We begin this comparison by considering how the distribution of W_{2796} values would change if such a




Figure 3. The $\log W_{2796}$ vs. $R_{\perp}^{\text{corr}} = R_{\perp} + (m_2/m_1)(\log M_*/M_{\odot} - 10.3)$ for our simulated observations of the fiducial Mg II-absorbing CGM. Here R_{\perp}^{corr} is calculated as described in the Figure 1 caption. The two-dimensional histogram indicates the number of simulated measurements in each bin normalized by the maximum number per bin within each column. The gray scale varies linearly with point density, with the darkest shading indicating the bins with the highest density. The cyan line shows the best-fit fiducial model as determined in Section 2.1, and the large points show the subset of the QSO–galaxy pair sample used to constrain this model. The points are colored as described in the legend and the Figure 1 caption.

fiducial CGM, once generated by this model, could be reobserved using b/g light sources of increasing size.

Perhaps the simplest configuration to consider is that of an extended b/g source consisting of numerous point-like (i.e., QSO-like) sources of equal intensity covering a projected area A_G . Picture the light from these point sources passing through CGM material and being absorbed by f/g Mg II ions. The strength of this absorption need not be the same along all of these sightlines; here we assume that the absorption observed toward each sightline j yields an equivalent width W_{2796}^j . In this case, the equivalent width observed along a spectrum that integrates the light from all of these point sources, $W_{2796}(A_G)$, is equal to the average of all W_{2796}^j , $W_{2796}(A_G) = \frac{1}{N} \sum_{j=1}^N W_{2796}^j$. This holds for any number of Mg II-absorbing structures with any column density and velocity distributions along these lines of sight, so long as the structures give rise to the total specified equivalent widths (W_{2796}^1, W_{2796}^2 , etc.).⁶

To make use of these mock extended b/g sources, we now refer to our fiducial model to generate a large sample of W_{2796} measurements (the distribution of which, by definition, is consistent with the W_{2796} distribution of our QSO–galaxy pair sample). We assume that the true parameters defining the relationship between R_{\perp} , $\log M_*$, and $\log W_{2796}$ are equal to the best-fit values of m_1 , m_2 , b , and σ_C as determined in Section 2.1. We then evenly sample the range in R_{\perp}^{corr} occupied by the QSO–galaxy pair data set. For a given R_{\perp}^{corr} value, we calculate the corresponding $\log \bar{W}_{2796}$ implied by the fiducial model and then generate 5000 random draws from a standard normal distribution centered at $\log \bar{W}_{2796}$ with a standard deviation equal to σ_C . We repeat this process at every R_{\perp}^{corr} and show the resulting distribution of $\log W_{2796}$ measurements in

⁶ If the point sources instead emit continua with varying intensities, the final observed $W_{2796}(A_G)$ will be a continuum flux-weighted average of the equivalent widths along each sightline.

Figure 3. Here the two-dimensional histogram indicates the number of realized $\log W_{2796}$ values in each bin normalized by the maximum number per bin in the corresponding histogram column. The gray scale varies linearly with density, with the darkest shading indicating the most frequently sampled bins. We note that formally, the variance in this simulated data set (σ_C^2) is slightly lower than that in the observed data set ($s_i^2 = \sigma_i^2 + \sigma_C^2$), as a proper realization of the observations would include an estimate of the typical measurement uncertainty (σ_W) and draw from a Gaussian with a dispersion $\sqrt{\sigma_W^2 + \sigma_C^2}$. However, the median value of σ_i in this sample is ~ 0.025 , over an order of magnitude lower than σ_C , such that its inclusion in the variance would yield a negligible difference in the final realized sample. Figure 3 shows the QSO–galaxy pair sample with colored points as in Figure 1 for comparison to our simulated data set. The best-fit fiducial relation is shown in cyan.

We may now use this simulated data set to predict the distribution of a new set of observations made with our simple extended b/g beam. First, we must consider a new quantity: the projected area of each Mg II absorber, A_A . Here we use “absorber” or “structure” to describe a system that gives rise to the same W_{2796} across its entire projected area. By definition, then, the edges of every absorber just touch those of the neighboring structures (so that there is no “empty space” or any overlap between absorbers). We remind the reader that the Mg II absorbers we are considering are strongly saturated, such that W_{2796} is likely more closely correlated with the velocity dispersion and number of absorbing clouds along the sightline than with the Mg II column density. We also note that although many of the absorbers in our sample have a 1:1 Mg II doublet ratio without exhibiting line-black troughs (see Table 4 and Figure 18 in GPG1), we are not able to usefully constrain the gas covering fraction via absorption-line analysis due to the low spectral resolution of the data.

The standard way of describing the $W_{2796}(A_G)$ that would be observed for one of these absorbers having an absorption strength W_{2796}^k (that is, assuming all neighboring absorbers have $W_{2796}^k = 0$ Å) is to write

$$W_{2796}(A_G) = C_f \int (1 - e^{-\tau_{\lambda}}) d\lambda, \quad (3)$$

where τ_{λ} is the optical depth in the Mg II 2796 transition, and where the integral above is the equivalent width that would be observed if the b/g source were fully covered (i.e., W_{2796}^k). In the case that the velocity structure and column density are coherent over A_A and $A_A \leq A_G$, the covering fraction $C_f^k = A_A/A_G$, and $W_{2796}(A_G) = C_f^k W_{2796}^k$. Moreover, an alternative scenario in which the column density and velocity structure of the absorption vary significantly within A_A (in such a way that they produce the same W_{2796} at every location within A_A) is also permitted in this framework and would yield equivalent results.

For our fiducial model, we assume that all absorbers have the same area (A_A), and we model the transverse size of each absorber using a square geometry. This allows us to populate a continuous two-dimensional cross section through the CGM. While there exist few empirical constraints on the shape and size of even individual Mg II absorbers, much less on these quantities as a function of W_{2796} , we expect that these two assumptions are significant oversimplifications of the physical

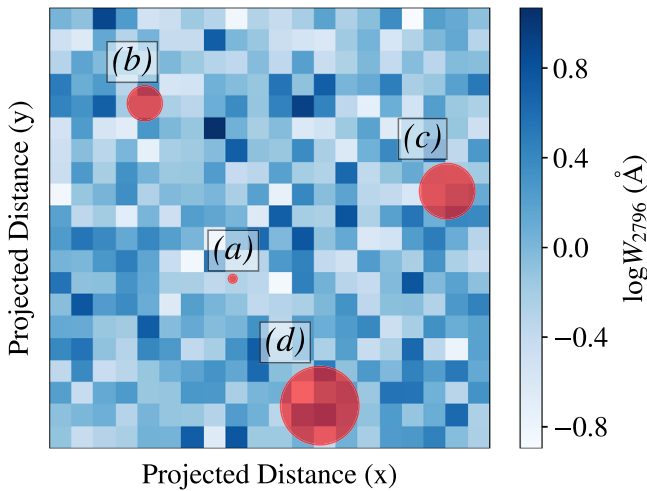


Figure 4. Realization of the $\log W_{2796}$ distribution within a small region of our model CGM. Each pixel in the map represents a different absorber and is colored by its W_{2796} , as indicated in the color bar. This particular realization is generated from a Gaussian $\log W_{2796}$ distribution centered at $\log \bar{W}_{2796}/\text{\AA} = 0$ and is thus meant to be representative of the CGM at $R_{\perp}^{\text{corr}} \sim 15$ kpc. The red circles represent extended b/g sources of various sizes, and their labels indicate the corresponding two-dimensional $\log W_{2796}$ distributions in the four panels of Figure 5. Note that the choice of physical size of each absorber (or pixel) in this toy model is unimportant; it is the ratio of the projected area of the b/g beam to the area covered by each absorber (x_A) that affects the observed $\log W_{2796}$ distribution.

layout of cool CGM material. We adopt them here to enable our demonstration of the qualitative effect of using extended b/g beams on the observed W_{2796} distribution. We will discuss requirements for increasing the realism of this model in Section 3.3.

We now consider the spatial distribution of absorbers within a small region of a halo (say a 10×10 kpc² projected “patch” at $R_{\perp}^{\text{corr}} \sim 15$ kpc). We may populate this region with absorbers by generating random draws from the standard normal distribution of the fiducial CGM centered at the $\log \bar{W}_{2796}$ appropriate for this value of R_{\perp}^{corr} . We place these absorbers at random locations within the region. An example of such a realization is shown in Figure 4. This 20×20 pixel box is populated with absorbers with a Gaussian $\log W_{2796}$ distribution having a mean $\log \bar{W}_{2796}/\text{\AA} = 0$ and a standard deviation σ_C . We then measure the absorption observed toward circular beams placed at random locations behind this region. For the purposes of this illustration, each absorber is 0.5×0.5 kpc² in area, and we are assuming that the gradient in $\log \bar{W}_{2796}$ over the range in R_{\perp}^{corr} covered by this patch is negligible. In the case that the size of the b/g beam is much smaller than the projected area of the f/g structures ($A_G \ll A_A$, or the ratio $x_A \equiv \frac{A_G}{A_A} \ll 1$), as demonstrated with the beam marked (a) in the figure, the observed absorption $W_{2796}(A_G)$ will almost always equal W_{2796}^* , the equivalent width of the particular structure probed. There may be some instances in which the b/g beam is placed behind the edges of two absorbers, such that $W_{2796}(A_G)$ is a weighted average of $W_{2796}^{k=1}$ and $W_{2796}^{k=2}$. However, this will be rare, given our condition for the relative sizes of the beams versus the absorbers. In this case, the observed distribution of $\log W_{2796}$ will be similar to that observed toward QSOs themselves, i.e., consistent with our realization of the fiducial CGM generated above.

This is demonstrated in Figure 5(a). To generate the two-dimensional histogram shown, we again evenly sample the

range in R_{\perp}^{corr} covered by the QSO–galaxy pair sample and simulate a CGM patch at each R_{\perp}^{corr} as described above with the appropriate $\log \bar{W}_{2796}$ value. To ensure a thorough sampling of the Gaussian, each patch includes 50×50 absorbers (or pixels). We adopt a beam radius of $r_G = 0.2$ pixels, such that $A_G = \pi r_G^2 = 0.13A_A$ in this case. We then place 2500 of these beams in random locations “behind” every patch, recording the $W_{2796}(A_G)$ observed toward each beam. The two-dimensional histogram in Figure 5(a) shows the distribution of these $W_{2796}(A_G)$ values for every patch, normalized as in Figure 3. Note that we generate a new patch for each step along the x -axis of size $\Delta R_{\perp}^{\text{corr}} \approx 1$ kpc, rather than adopting the same $\log W_{2796}$ distribution over a $\Delta R_{\perp}^{\text{corr}}$ interval of several kpc, as in the illustration shown in Figure 4. Our PRIMUS galaxy pair W_{2796} measurements, including only pairs having $R_{\perp} < 50$ kpc, those in which the b/g galaxy is not host to a bright active galactic nucleus (AGN), and absorption measurements that are unaffected by blending, are shown with large colored points (GPG1). The best-fit relation is again shown with a cyan line. As expected, the histogram in Figure 5(a) is just slightly more narrowly distributed about this best-fit model than the histogram shown in Figure 3.

As the size of the b/g beam approaches the size of the f/g absorbers, the frequency of sightlines probing more than one absorber increases. If $A_G > A_A$, every b/g sightline will intercept more than one absorber. To demonstrate this, we repeat the above exercise adopting b/g beams such that $x_A = 2, 5$, and 10 (as indicated by the beams (b), (c), and (d) in Figure 4) and show the resulting $\log W_{2796}$ distributions in Figure 5 panels (b)–(d), respectively. In effect, for each beam, we are performing the sum $W_{2796}(A_G) = \sum_{k=1}^N C_f^k W_{2796}^k$, where the covering fraction is adjusted from its standard value of $C_f^k = A_A/A_G$ for absorbers with edges that overlap those of the beam. The scatter in these distributions is reduced as the b/g beam increases in size; in addition, because large beams are measuring the arithmetic mean of W_{2796} values drawn from a lognormal distribution, the $\log W_{2796}$ that occurs with the highest frequency at each R_{\perp}^{corr} lies slightly above the original fiducial best-fit relation.

We may therefore constrain the ratio x_A by comparing these simulated two-dimensional $\log W_{2796}$ distributions with the dispersion in $\log W_{2796}$ measured in our PRIMUS galaxy pair data set. For consistency with the QSO–galaxy pair subsample used to construct our fiducial model, we must limit our comparison to include only the 27 galaxy pairs from GPG1 with f/g galaxies that are star-forming (shown with cyan and blue symbols in Figure 5; we also show measurements for pairs with passive f/g galaxies in red for completeness). Examining the figure by eye, it is clear that the galaxy pair sample $\log W_{2796}$ measurements exhibit a relatively high level of dispersion about the best-fit fiducial model, apparently comparable to that of the two-dimensional histogram in panel (a).

To quantify this dispersion, we calculate the offset between the best-fit fiducial relation and each $\log W_{2796}$ measurement ($\Delta \log W_{2796} = \log W_{2796} - \log \bar{W}_{2796}$) and show the distribution of these values with the orange histograms in Figure 6. Only measurements of the CGM around star-forming hosts are included here; in addition, in cases for which W_{2796} is an upper limit, the offset of this limit is shown only if $\Delta \log W_{2796}/\text{\AA} < 0.1$. We calculate the same offsets for the simulated CGM data sets shown in Figure 5, with the green, turquoise, cyan, and blue histograms showing the dispersion in data sets predicted for b/g beams with $x_A \ll 1$ and $x_A = 2, 5$, and




Figure 5. Fiducial Mg II-absorbing CGM model as observed toward extended b/g beams. Each column of each gray-scale histogram is calculated by generating a projected “patch” of randomly distributed Mg II absorbers with area A_A selected from a Gaussian $\log W_{2796}$ distribution having a mean $\log \bar{W}_{2796}$ appropriate for the R_{\perp}^{out} of that column (as demonstrated in Figure 4). We then “observe” each patch toward numerous randomly placed b/g beams with areas (A_G) chosen as indicated with the ratio $x_A \equiv \frac{A_G}{A_A}$ in the lower left corner of each panel (and the red circles in Figure 4). The histograms show the number of simulated $\log W_{2796}$ measurements per bin, normalized as described in the Figure 3 caption. The cyan line shows the best-fit fiducial linear model. The colored points show $\log W_{2796}$ measured along PRIMUS b/g galaxy sightlines that lack bright AGNs and with $R_{\perp} < 50$ kpc. Pairs with star-forming f/g galaxies are indicated with cyan/blue stars, and pairs in which the f/g galaxy is quiescent are indicated with red pentagons. Pairs with star-forming and quiescent f/g hosts that yield upper limits on $\log W_{2796}$ are shown with open blue stars and open red pentagons, respectively. The dispersion in the $\log W_{2796}$ distribution is predicted to decrease as the ratio x_A increases. The $\log W_{2796}$ distribution observed toward PRIMUS b/g galaxies appears consistent with that expected for a small value of this ratio ($x_A \ll 1$), suggesting that the strength of Mg II absorption varies on a scale larger than these b/g beams.

10, respectively. Here we simulate additional data sets that assume b/g beams having $x_A = 12, 15, 20,$ and 30 and show the resulting distributions of $\log W_{2796}$ offsets with purple, magenta, red, and brown histograms. Each histogram is generated from 2500 simulated W_{2796} (A_G) measurements and has been normalized such that the peak value is equal to the peak of the histogram showing the observed $\Delta \log W_{2796}$ distribution.

We demonstrate that these simulated data sets are approximately normally distributed by performing a nonlinear least-squares fit of a Gaussian to each, and we show the best-fit curves with thick colored lines. We print the mean (μ) and standard deviation (σ) of each sample above the corresponding plot panel.

To test for consistency between the observed $\Delta \log W_{2796}$ distribution and each of these simulated data sets, we implement a survival analysis using the Python package

lifelines (Davidson-Pilon et al. 2018).⁷ For each value of x_A , we first perform a log-rank test comparing the survival distributions of a Gaussian function having the corresponding values of μ and σ and the observed set of $\Delta \log W_{2796}$ values. Here we include all measurements constraining the CGM of star-forming f/g galaxies, including all upper limits on $\log W_{2796}$ (regardless of their value).⁸ The resulting P -values are printed above the appropriate panels in Figure 6.

Then, to account for the magnitude of the uncertainty in each observed value of $\Delta \log W_{2796}$, we consider both the

⁷ Available at <https://github.com/CamDavidsonPilon/lifelines>.

⁸ We also multiply both simulated and observed data sets by -1 , as the lifelines implementation of the log-rank test handles right-censored data only.

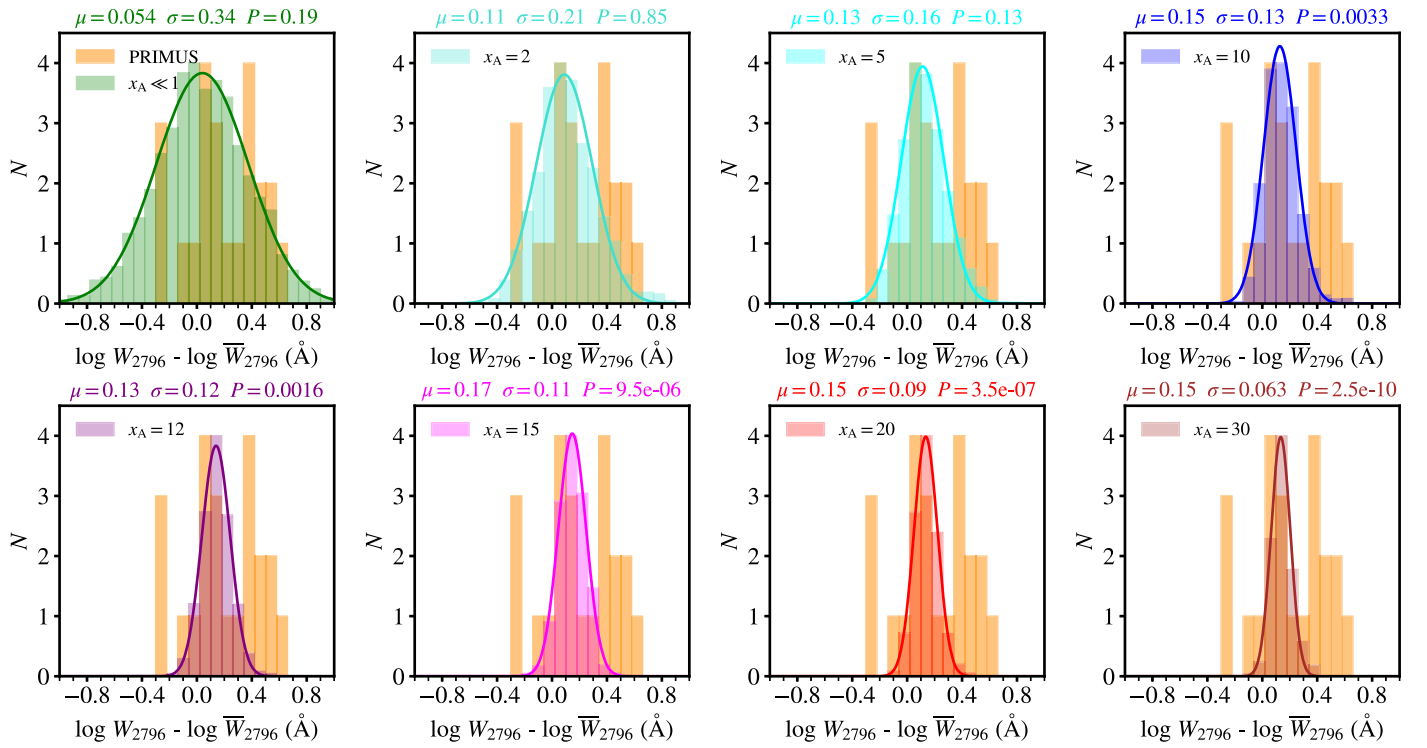


Figure 6. Distribution of the offsets ($\Delta \log W_{2796}$) between the best-fit fiducial Mg II-absorbing CGM relation and our $\log W_{2796}$ measurements toward PRIMUS b/g galaxies probing the halos of star-forming f/g galaxies (orange histograms). In cases for which $\log W_{2796}$ is an upper limit, the offset is included in the plotted distribution only if it is < 0.1 . The green, turquoise, cyan, and blue histograms (top panels) show the distribution of $\Delta \log W_{2796}$ offsets for each of the simulated data sets shown in Figure 5, i.e., assuming a fiducial Mg II CGM observed with b/g beams having $x_A \ll 1$ and $x_A = 2, 5,$ and $10,$ respectively. The purple, magenta, red, and brown histograms (bottom panels) show simulated distributions of $\Delta \log W_{2796}$ offsets assuming $x_A = 12, 15, 20,$ and 30 (not shown in Figure 5). The colored curves show Gaussian fits to each distribution. The mean (μ) and standard deviation (σ) of each simulated data set are printed above the corresponding panel, along with the P -value obtained from a log-rank test of the consistency between the observations and these simulated distributions.

uncertainty in the best-fit fiducial relation and the measurement uncertainty in $\log W_{2796}$ (σ_i) at each value of $R_{\perp,i}^{\text{corr}}$. First, we draw 1000 random sets of the parameters $m_1, m_2,$ and b from the PPDF of our linear Mg II-absorbing CGM model (as in Section 2.1). We adopt the parameters in each of these sets as fiducial parameters and recalculate $\log \bar{W}_{2796}$ accordingly. For each set, we also generate a new realization of the observed $\log W_{2796}$ values, perturbing each detection by a random draw from a Gaussian distribution having a standard deviation equal to σ_i . Finally, we perform the log-rank test for each of these 1000 realizations and record the resulting P -values.

We show the detailed distributions of these P -values in Figure 7. We summarize the distributions in box-and-whisker form and compare them to our null hypothesis rejection criterion ($P = 0.05$) in Figure 8. For data sets simulated with $x_A < 5$, values of $P > 0.05$ are quite common (see green and turquoise histograms), such that we fail to reject the null hypothesis in these cases (i.e., we find no evidence that the observed and simulated survival distributions differ). For values of $x_A \geq 5$ (for which P -values > 0.05 are relatively infrequent), we print the median and upper 95th- and 99th-percentile P -values in the corresponding panels of Figure 7.

In detail, these results indicate inconsistency between the observed and simulated $\Delta \log W_{2796}$ distributions with increasing statistical significance as x_A increases from 10 to 30. Referring first to the direct comparison between the observed set of $\Delta \log W_{2796}$ values (adopting the best-fit fiducial CGM) and the simulated distributions, we reject the null hypothesis wherever $x_A \geq 10$. Considering the effects of the combined uncertainties in the fiducial CGM model and our $\log W_{2796}$

measurements, we find that the null hypothesis is ruled out in the majority of realizations for $x_A = 5$ (with $P(\text{med}) = 0.005$); however, the log-rank test yields $P > 0.05$ for 29% of these realizations. This is reflected in Figure 8 in the overlap of the colored box at $x_A = 5$ showing the upper and lower quartile values of the P -value distribution with the horizontal line showing $P = 0.05$. When $x_A = 10$, the null hypothesis is ruled out for 90% of realizations, while for $x_A = 15, P < 0.05$ for all but 22 of the 1000 Monte Carlo realizations of the observed data set (such that the whisker at $x_A = 15$ in Figure 8 ends below $P = 0.05$). Thus, at the most conservative level, we conclude that within the framework of this CGM model, we must reject the null hypothesis for ratios $x_A \geq 15$ with 95% confidence.

Moreover, given the frequency with which the simulated $\Delta \log W_{2796}$ distribution for $x_A = 10$ is demonstrated to be inconsistent with that observed, we also consider models having $10 \leq x_A < 15$ to be at least marginally inconsistent with our PRIMUS pairs data set.

2.3. A Limit on the Coherence Scale of W_{2796} in the Inner CGM

We may now use these constraints on the ratio of the projected areas of our b/g galaxy beams to that of Mg II absorbers (x_A) to place a limit on the projected area over which W_{2796} does not fluctuate. We will refer to the square root of this area as the “coherence scale” or “coherence length” of Mg II absorbers, ℓ_A .

In GPG1, we performed a detailed analysis assessing the spatial extent of the rest-frame UV continuum emission arising from our b/g galaxy sample (see Section 6). We found that the vast majority of the galaxies for which *HST* imaging is

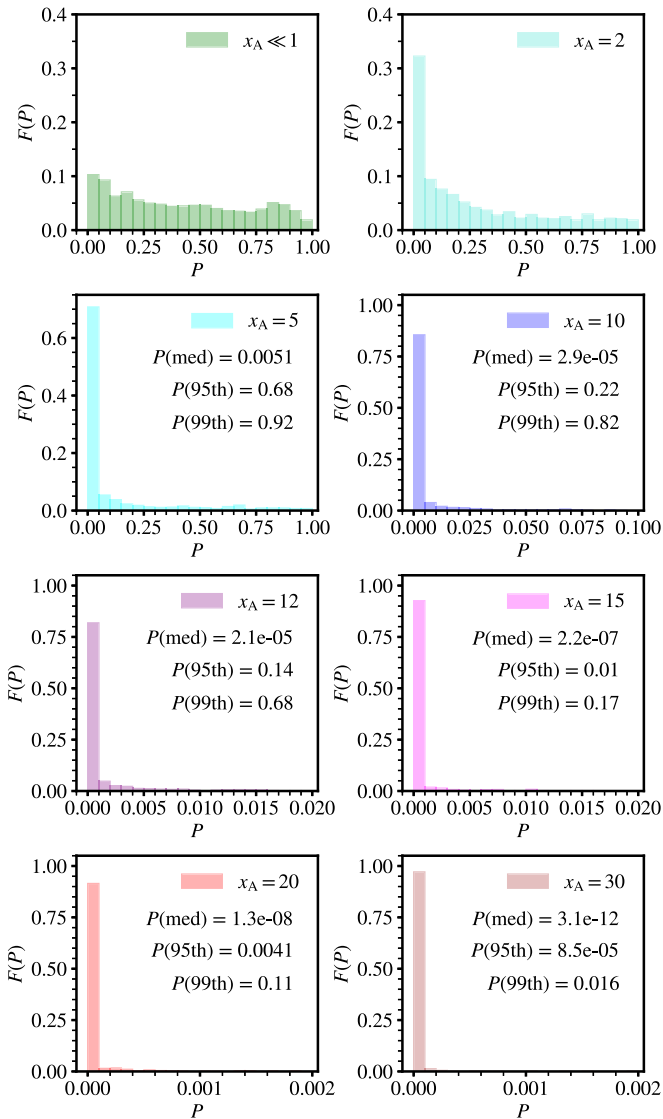


Figure 7. Normalized frequency distributions of P -values resulting from log-rank tests comparing the survival distributions of the observed $\Delta \log W_{2796}$ data set and those of the simulated data sets with $x_A \ll 1$ (green histogram) and $x_A = 2$ (turquoise histogram), 5 (cyan histogram), 10 (blue histogram), 12 (purple histogram), 15 (magenta histogram), 20 (red histogram), and 30 (brown histogram). The median and upper 95th- and 99th-percentile P -values are printed in the corresponding panels for $x_A \geq 5$. P -values ruling out the null hypothesis are common for these latter distributions, pointing to statistically significant inconsistencies between the observed and simulated $\Delta \log W_{2796}$ distributions.

available in the rest-frame optical (and that do not host bright AGNs) have effective radii $R_{\text{eff}}(z_{f/g}) > 2.0$ kpc at the redshift of the corresponding f/g galaxy ($z_{f/g}$). The median effective radius of this b/g galaxy sample is $R_{\text{eff}}(z_{f/g}) = 4.1$ kpc. We also demonstrated that in general, the half-light radii of bright galaxies at $z \sim 0.5$ –1 are very similar in both the rest-frame optical and rest-frame UV, justifying the assumption that the projected extent of our b/g galaxy sample as measured in the rest-frame optical reflects the sizes of the continua probing f/g Mg II absorption.

Adopting the size limit $R_{\text{eff}}(z_{f/g}) > 2.0$ kpc, such that $A_G = \pi (2.0^2) \text{ kpc}^2$, our most conservative limit on x_A ($x_A = A_G/A_A$ is not ≥ 15) requires $A_A > \frac{4.0\pi}{15} \text{ kpc}^2$, such that $\ell_A > 0.9$ kpc. A somewhat less conservative constraint may be calculated by adopting the median b/g galaxy size, yielding




Figure 8. Box-and-whisker plots representing each of the P -value frequency distributions shown in Figure 7. Each box extends to the upper and lower quartile values, with the median indicated by the central black line. The whiskers extend to the 5th- and 95th-percentile values. The box colors are the same as those used for the corresponding x_A values in Figures 6 and 7. The dashed blue horizontal line indicates $P = 0.05$. Over 95% of the realizations adopting $x_A = 15$ yield P -values below this limit, allowing us to reject the null hypothesis that the observed and simulated $\Delta \log W_{2796}$ distributions arise from the same parent population. Simulations with $10 \leq x_A \leq 12$ also yield P distributions dominated by values $\ll 0.05$, but in addition, they give rise to a nonnegligible fraction of realizations for which the null hypothesis cannot be rejected.

$A_G = \pi (4.1^2) \text{ kpc}^2$, $A_A > \frac{16.8\pi}{15} \text{ kpc}^2$, and $\ell_A > 1.9$ kpc. Given that the b/g galaxies in our sample that have been imaged by *HST* range in size from $R_{\text{eff}}(z_{f/g}) = 1.0$ kpc to as large as $R_{\text{eff}}(z_{f/g}) = 7.9$ kpc, we prefer this latter limit and will adopt it in the discussion that follows. If we instead were to adopt the less secure limit $x_A < 10$, the corresponding constraint on the coherence length (again using $R_{\text{eff}}(z_{f/g}) = 4.1$ kpc) would be $\ell_A > 2.3$ kpc. We expect that an expansion of this galaxy pair data set and a larger sample of QSO–galaxy pairs suitable for establishing the form of the “fiducial” CGM will eventually improve the robustness of this larger limit.

Regardless of potential future developments, the foregoing analysis has placed a unique limit on the coherence scale of the Mg II–absorbing CGM within $R_L < 50$ kpc of $\sim L^*$ galaxies in the low-redshift universe ($\ell_A > 1.9$ kpc). This limit is fully consistent with those derived in previous work analyzing single low-redshift galaxy pairs with impact parameters $R_L < 50$ kpc (which have yielded $\ell_A > 0.4$ kpc; Rubin et al. 2010; Diamond-Stanic et al. 2016; Peroux et al. 2018). However, as noted above, this does not necessarily imply that the material giving rise to absorption over this scale must occupy a single structure subtending this area. On the contrary, there is a high degree of degeneracy in the physical locations and velocities of “clouds” producing absorption of a given W_{2796} even along a single pencil-beam sightline. We will consider the implications of this constraint on the physical structure of Mg II absorption in the CGM in detail in Section 3.1.

3. Discussion

3.1. The Physical Nature of the Mg II–absorbing CGM

The power of QSO absorption-line spectroscopy goes far beyond the simple assessment of the incidence and distribution of the Mg II absorption strength (W_{2796}) in the halos of galaxies. High spectral resolution observations of QSO

sightlines have revealed the velocity and column density structure of these absorbers in exquisite detail. Here we describe these constraints and consider their implications for the interpretation of our limit on the coherence length of W_{2796} in the inner CGM.

3.1.1. Relating Cool Structures and Mg II Absorbers

The analysis of Keck/HIRES QSO spectroscopy probing ~ 23 Mg II absorbers at $z \sim 0.4\text{--}1.2$ by Churchill & Vogt (2001) and Churchill et al. (2003b) remains one of the most careful and germane in this context. These authors performed Voigt profile fitting of systems ranging in strength from $W_{2796} \sim 0.3$ to 1.5 \AA , finding that without exception, the absorbers consisted of multiple Voigt profile “components.” In general, these components are kinematically narrow, with Doppler parameters $b_D \sim 5 \text{ km s}^{-1}$ —a velocity width that naturally arises from thermal broadening in gas at temperatures $T \sim (3\text{--}4) \times 10^4 \text{ K}$ (Churchill et al. 2003b). These individual, narrow components have frequently been attributed to structures called absorbing “clouds” in the literature (Churchill et al. 1999; Churchill & Charlton 1999), and we will continue to use this term here.

In detail, Churchill et al. (2003b) found that the number of these clouds composing each Mg II absorber (N_{cl}) ranged from two to 18, with N_{cl} increasing approximately linearly with W_{2796} (albeit with significant scatter). A linear fit to their data set yielded a slope $= 0.058 \pm 0.004 \text{ \AA cloud}^{-1}$ and intercept $W_{2796} = 0.28 \text{ \AA}$ for $N_{\text{cl}} = 1$. A few systems lying well above this trend were found to be associated with heavily saturated Mg II components arising from damped Ly α systems (DLAs) or Lyman limit systems (LLSs), and the authors noted that the number of components in these cases could have been underestimated due to kinematic overlap or blending. Finally, the column density distribution function of all clouds in the absorber sample was measured to a limiting column density of $\log N(\text{Mg II}) \approx 11.6 \text{ cm}^{-2}$ and found to be consistent with a power law $f(N) \propto N^{-\delta}$ with $\delta = 1.59 \pm 0.05$. This implies that, for example, a cloud with column density $\log N(\text{Mg II}) \approx 12.5 \text{ cm}^{-2}$ occurs with more than an order of magnitude higher frequency than a cloud with $\log N(\text{Mg II}) \approx 13.5 \text{ cm}^{-2}$.

These results are suggestive of a scenario in which Mg II absorbers, regardless of their total strength, are composed of multiple kinematically cold clouds superimposed along the same line of sight. The resultant Mg II “system” will be stronger if the sightline pierces a larger number of these clouds and/or if there is a large kinematic dispersion from cloud to cloud. The Churchill et al. (2003b) sample does not include systems as strong as many of the absorbers detected in our PRIMUS b/g galaxy sightlines (with $W_{2796} > 2 \text{ \AA}$), and high-resolution spectroscopy of similarly strong absorbers tends to yield broad, line-black profiles (Mshar et al. 2007). While it is difficult to demonstrate that such profiles arise from multiple narrow components from analysis of Mg II alone, examination of weaker Fe II transitions associated with the same systems generally reveals a complex, multiple-cloud structure (Mshar et al. 2007).

3.1.2. Exploration of Degeneracies in Component Structure and Velocity Dispersion

The assumption that Mg II systems are indeed dominated by kinematically cold clouds has important implications for our




Figure 9. Distribution of W_{2796} measured from simulated Mg II line profiles. Each profile is composed of N_{cl} kinematically cold Mg II-absorbing “clouds.” The clouds in every profile have velocities drawn at random from a Gaussian distribution with FWHM $= \Delta v_{\text{FWHM}}$. For each value of N_{cl} and Δv_{FWHM} , we generate 100 line profile realizations and indicate the resulting range in W_{2796} with the green ($\Delta v_{\text{FWHM}} = 150 \text{ km s}^{-1}$), blue ($\Delta v_{\text{FWHM}} = 350 \text{ km s}^{-1}$), and red ($\Delta v_{\text{FWHM}} = 550 \text{ km s}^{-1}$) contours. The dashed line shows a linear fit to the number of clouds vs. W_{2796} for the sample of strong Mg II absorbers discussed in Churchill et al. (2003b).

interpretation of the coherence scale (ℓ_A) discussed in Section 2.3. We emphasize that ℓ_A refers to the scale over which W_{2796} does not fluctuate, rather than a length scale associated with the cold clouds themselves. These quantities must be considered independently due to the degeneracy in the relationship between W_{2796} , the number of clouds composing an absorber (N_{cl}), and the kinematic spread of these clouds. For instance, numerous saturated clouds all having a very similar radial velocity (but that are physically quite distant from each other) can, in principle, yield a much smaller W_{2796} than a few clouds with velocities separated by $> 50 \text{ km s}^{-1}$.

To better quantify the extent of this degeneracy, we simulate a suite of Mg II absorbers using code included in the Python package `linetools`.⁹ To simplify our experiment, we first assume that each cloud composing these absorbers has the same Doppler parameter and column density. We choose a Doppler parameter $b_D = 5.4 \text{ km s}^{-1}$, corresponding to the median value for the Churchill et al. (2003b) sample. We choose a column density of $\log N(\text{Mg II}) = 13.0 \text{ [cm}^{-2}\text{]}$, as this value is close to the center of the column density distribution of this sample (as shown in, e.g., Figure 4(a) of Churchill et al. 2003b). We then generate line profiles for absorbers having between $N_{\text{cl}} = 1$ and 20 cold clouds. For each system, the velocities of the clouds are drawn at random from a Gaussian with a full width at half maximum $\Delta v_{\text{FWHM}} = 150, 350, \text{ or } 550 \text{ km s}^{-1}$. These values sample the expected range in the velocity dispersions of dark matter halos hosting galaxies with $9.1 < \log M_*/M_\odot < 11.1$ at $z = 0.5$ (assuming halo masses in the range $11.3 \lesssim \log M_h/M_\odot \lesssim 13.3$; Maller & Bullock 2004; Moster et al. 2013). We generate 100 realizations of each system (i.e., at a given N_{cl} and Δv_{FWHM}) and measure the resulting W_{2796} .

The ranges in these W_{2796} values are indicated by the colored contours in Figure 9. The outlines of the contours are placed at

⁹ Available at <https://github.com/linetools>.

the minimum and maximum W_{2796} for every value of N_{cl} . For comparison, the linear fit to the $N_{\text{cl}} - W_{2796}$ relation for the Mg II absorber sample discussed in Churchill et al. (2003b) is indicated with a black dashed line. The intercept of this relation at $N_{\text{cl}} = 1$ is slightly higher than the W_{2796} of our $N_{\text{cl}} = 1$ line profiles, suggesting that the weakest absorbers in the Churchill et al. (2003b) sample have a Doppler parameter or column density slightly greater than the values we have chosen for each individual “cloud.” However, this fit otherwise corresponds quite closely to the lower envelope of W_{2796} values exhibited by profiles with $\Delta v_{\text{FWHM}} = 150 \text{ km s}^{-1}$. The overall higher W_{2796} distribution of our simulated sample is likely driven by the relatively low velocity spread of clouds giving rise to the absorbers in the Churchill et al. (2003b) sample. Most of these systems are produced by clouds with relative velocities $\Delta v < 100 \text{ km s}^{-1}$, and the system with the highest velocity spread in the sample has only $\Delta v \sim 400 \text{ km s}^{-1}$. Figure 9 demonstrates that systems with velocity dispersions $\lesssim 150 \text{ km s}^{-1}$ are unlikely to yield $W_{2796} > 2 \text{ \AA}$; indeed, the maximum W_{2796} among the Churchill et al. (2003b) absorbers is $\sim 1.5 \text{ \AA}$.

Figure 9 further demonstrates that there are numerous physical scenarios that may give rise to absorbers with a given W_{2796} . For instance, an absorber having W_{2796} close to the median value for the PRIMUS pair sightlines shown in Figure 5, $W_{2796} \approx 0.7 \text{ \AA}$, can arise from a system of $N_{\text{cl}} \sim 5$ –6 clouds with a large velocity spread or $N_{\text{cl}} \sim 9$ –10 clouds with a low velocity dispersion such that the individual cloud line profiles overlap in velocity space. The range of N_{cl} values that can plausibly yield a given absorber strength broadens as W_{2796} increases. Equivalent widths as large as the maximum of the observed distribution ($W_{2796} \approx 2.6 \text{ \AA}$) are only exhibited by the simulated systems having $N_{\text{cl}} \sim 19$ –20 and a “maximal” velocity spread. However, our choice to limit $N_{\text{cl}} \leq 20$ is not driven by any physical constraint, and such strong systems could easily result from absorbers with $N_{\text{cl}} > 20$ and less extreme kinematics.

3.1.3. Interpretation of the W_{2796} Coherence Scale

In light of these considerations, we now return to the interpretation of our constraint on the coherence scale of W_{2796} . Our analysis requires that $\ell_A > 1.9 \text{ kpc}$, or $A_A > 3.5 \text{ kpc}^2$. Thus, if we consider a contiguous “patch” of the CGM with a projected area A_A , the corresponding three-dimensional column of the CGM must be populated with absorbing structures that give rise to the same W_{2796} across this entire area. These structures, however, may have a wide variety of configurations.

We present cartoons illustrating three such configurations in Figure 10. The topmost cartoon diagrams the layout of the “inner” CGM. The location and approximate scale of the host galaxy is indicated with the red symbol. The gray region represents a slice through the CGM with impact parameter increasing to the right. We adopt the assumption that the absorption is in any case produced by numerous kinematically narrow clouds as suggested by Churchill et al. (2003b), and we represent these structures in green. We place these clouds along the vertical axis according to their velocity relative to the f/g host galaxy. The projected extent of the beam toward a typical b/g galaxy is indicated in blue.

The simplest scenario that can yield the required coherence invokes clouds that themselves have projected areas $> 3.5 \text{ kpc}^2$, as illustrated in panel (1) of Figure 10. Numerous clouds

observed along a given beam yield the strongest absorbers, while beams absorbed by only a few clouds yield weak absorption. Based on the computations discussed in Section 3.1.2, the weakest absorption observed toward our b/g galaxy sample is sufficiently strong that it likely arises from ~ 3 –5 such clouds, while the strongest absorbers would be composed of $\gtrsim 20$ of these extended clouds. (We have reduced these numbers for the purposes of clarity in the figure.)

Alternatively, each of these kinematically cold clouds could instead subtend a much smaller area, e.g., several square parsecs. The weakest absorbers in our b/g galaxy sample would in this case arise from many thousands of clouds (as would be required to cover an area $> 1 \text{ kpc}^2$). However, they would be distributed such that the sightline through each area element of the patch (dA_A) intercepts only ~ 3 –5 clouds. In regions of the CGM with the strongest observed absorption, many thousands of clouds would again be required, and here each element dA_A would be shadowed by $\gtrsim 20$ clouds. Note that the clouds along the sightline through any given subpatch dA_A may be completely physically distinct from those clouds within a neighboring subpatch. They may have different physical locations and velocities within the halo, and the number of clouds from subpatch to subpatch may also vary as long as the cloud velocity dispersion yields a consistent W_{2796} throughout the patch.

This second scenario is illustrated in panel (2) of Figure 10. Here each cloud subtends a smaller projected area than in panel (1) (they have diameters approximately one-third the width of the b/g galaxy beam). The clouds exhibit significant variations in their relative velocities from sightline to sightline. However, the numbers of clouds along different sightlines are spatially correlated, with a “correlation length” at least as large as $\ell_A > 1.9 \text{ kpc}$. In other words, regions of the halo having numerous small clouds along a given sightline tend to have projected areas $> 3.5 \text{ kpc}^2$. Low- W_{2796} regions arising from the absorption of only a few of these small clouds also tend to extend over areas larger than the clouds themselves. As noted above, the pencil-beam sightlines within a given region of coherence need not all pierce the same number of clouds (as is shown in the figure). We have illustrated the scenario this way for simplicity’s sake; however, our observations require only that the combination of Δv and N_{cl} yield similar W_{2796} within a region A_A .

Panel (3) of Figure 10 illustrates a final scenario, one that is ruled out by our observations. Here the individual Mg II-absorbing structures have small sizes, as in panel (2). However, the number of clouds varies significantly from one sightline to the next, such that W_{2796} varies on projected spatial scales much smaller than the b/g galaxy beam. Similar variations could also arise from significant changes in Δv from sightline to sightline. Such large variations in either quantity are unlikely given the results of the analysis described above.

In summary, on a phenomenological level, our measurements point to either (1) Mg II-absorbing structures that extend over projected areas $> 3.5 \text{ kpc}^2$ or (2) a spatial correlation in the numbers of these clouds and their velocity dispersion with a correlation length $> 1.9 \text{ kpc}$. These two types of configurations may plausibly arise from a variety of physical scenarios, e.g., extended cool streams of inflowing gas or numerous cold clouds embedded in a rapidly expanding and kinematically complex outflow. We also note here that echelle-resolution spectroscopy of extended b/g sources as presented in

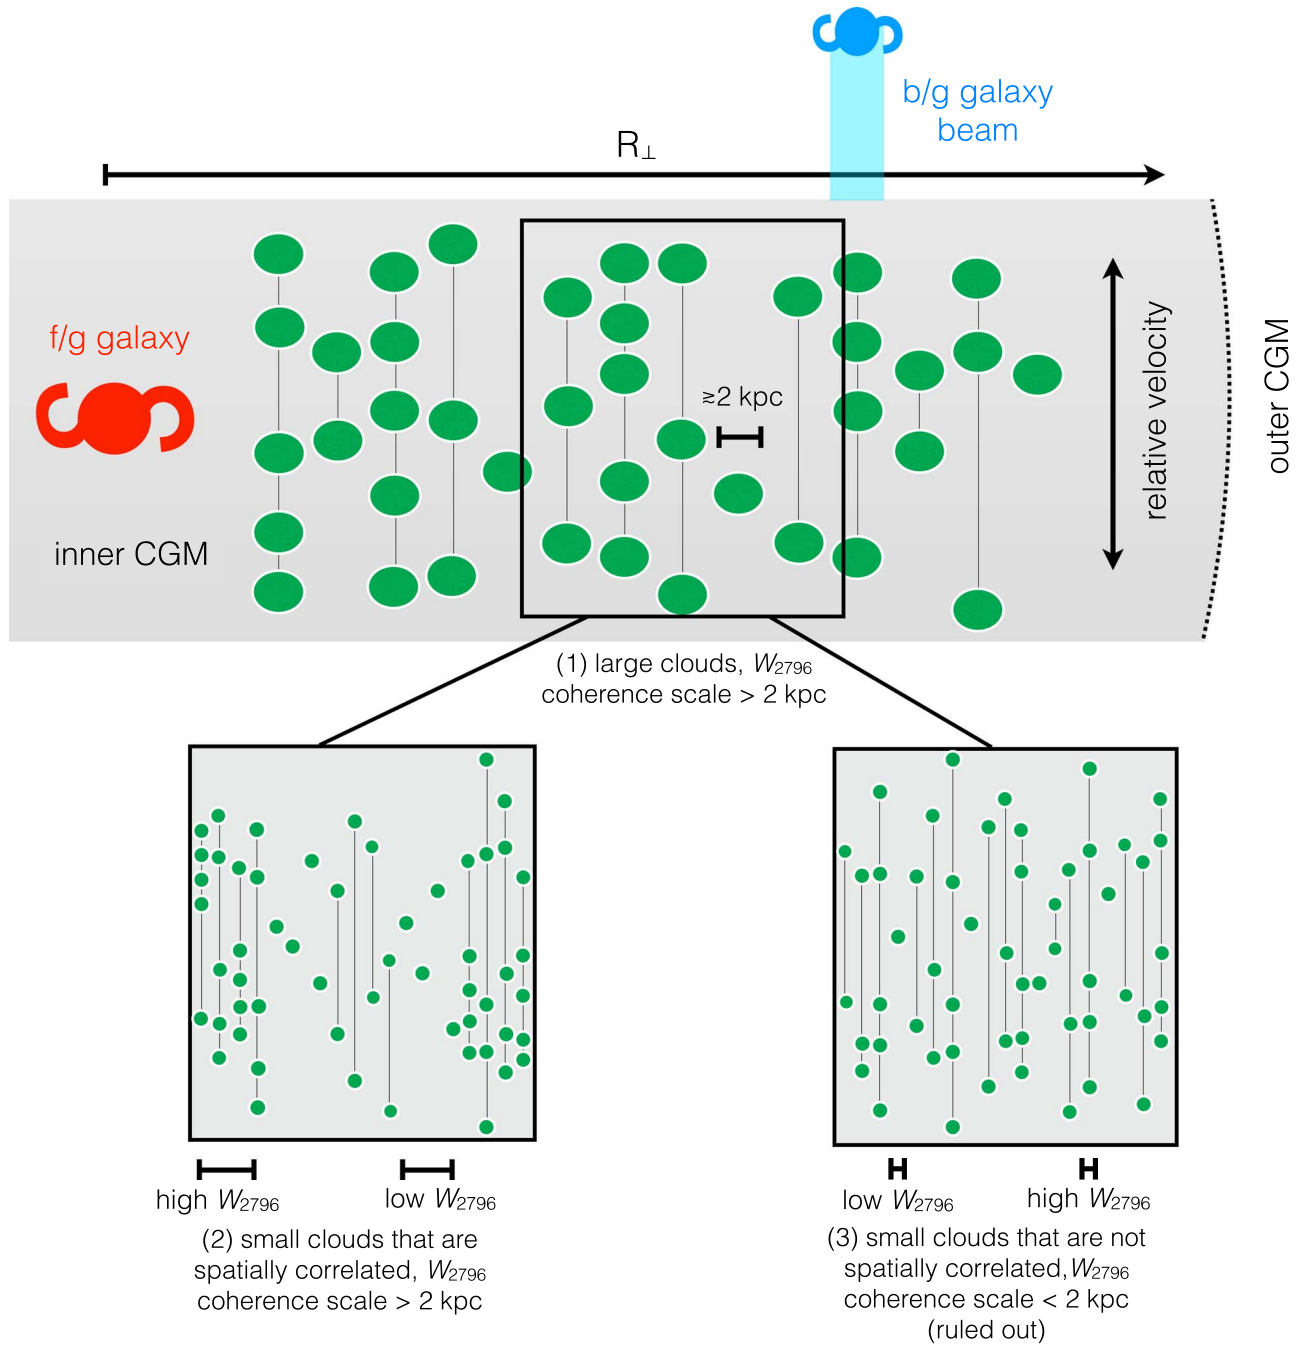


Figure 10. Several scenarios for the physical distribution of Mg II-absorbing structures in the inner CGM ($R_{\perp} < 50$ kpc). Panel (1) shows a slice through this inner halo region (gray) centered around a host galaxy indicated in red. Each green circle or oval represents a distinct structure or cloud, which we assume gives rise to saturated (or nearly saturated) Mg II absorption (having $\log N(\text{Mg II}) > 13.0$ [cm^{-2}]) with a narrow velocity dispersion ($b_D \sim 5\text{--}10$ km s^{-1}). These structures are placed along the vertical axis according to their mean line-of-sight velocity relative to the host galaxy. Sightlines that pass through several clouds (as indicated by the thin vertical black lines) give rise to large W_{2796} , while sightlines passing through only a single cloud yield $W_{2796} \sim 0.2$ Å. The rough size of one of our b/g galaxy beams is indicated by the blue symbol at the upper right. In panel (1), each individual cloud fully covers the beam of the b/g galaxy, such that the velocity structure within the beam is coherent. In panel (2), the physical scale over which W_{2796} changes is similar to that of the b/g galaxy beam; however, the high- W_{2796} regions contain numerous smaller clouds with incoherent velocities. The W_{2796} distributions observed toward the PRIMUS b/g galaxy sightlines would be similar for these two scenarios. Our analysis rules out the scenario shown in panel (3), in which the W_{2796} varies on scales smaller than that subtended by our b/g galaxy beams.

Diamond-Stanic et al. (2016) has the potential to differentiate between these scenarios, as it is sensitive to both velocity structure and variations in covering fraction across the beam. Indeed, these authors report a covering fraction of unity for both strong and weak Mg II velocity components at $R_{\perp} = 27$ kpc over a beam ≈ 0.4 kpc in radius, lending support to scenario (1) above. We will further consider the physical implications of our findings and these scenarios in Section 3.4.

3.2. Complementary Constraints on the Coherence of Mg II Absorption

Our data set is not the first to offer constraints on the sizes of cool, photoionized structures in the CGM. Spectroscopy of gravitationally lensed QSO sightlines yielded the first such measurements beginning in the 1980s (e.g., Young et al. 1981; Weymann & Foltz 1983; Foltz et al. 1984). Photoionization

modeling provides yet another assessment of absorber size, albeit indirect. Finally, spectroscopy of UV-bright stars or AGNs probing cool gas in the halo of our own Galaxy constrains its absorption coherence over a broad range of length scales (e.g., Smoker et al. 2015). Here we briefly review the findings of these experiments and consider how they complement the present experimental approach.

3.2.1. The Inner CGM of the Milky Way

A key testing ground for the coherence of circumgalactic material is the Milky Way itself, in which the extent of cool gas clouds can be viewed directly in 21 cm emission. The largest recent neutral hydrogen surveys have been conducted with single-dish radio telescopes with spatial resolutions of $4'–36'$ (e.g., Putman et al. 2002; Kalberla et al. 2005; Peek et al. 2011), corresponding to sizes of $\sim 25–210$ pc at a distance of 20 kpc. This imaging has revealed a great variety of structures, with those at the highest velocities (i.e., high-velocity clouds, or HVCs) tending to be located beyond ~ 5 kpc from the disk (Wakker 2001). The most massive of these structures are organized into extended cloud complexes or streams (e.g., the Magellanic Stream; Complex C) with physical dimensions of several to several tens of kpc (e.g., Thom et al. 2008; D’Onghia & Fox 2016). There is, in addition, a substantial population of “compact,” isolated HVCs with somewhat smaller sizes ($\sim 5'–15'$, with a median physical size of ~ 10 pc assuming a distance of 10 kpc; Putman et al. 2002; Saul et al. 2012). Interferometric imaging of a handful of neutral hydrogen cloud complexes in the halo has revealed that they are composed of numerous dense clumps only $\sim 30–50$ pc across (Richter et al. 2005; Ben Bekhti et al. 2009), suggesting that such fine substructure may be a common feature of the HVC population as a whole.

However, the ionized gas associated with these clouds is known to be significantly more extended than the neutral component (Lehner & Howk 2011). Furthermore, the HI velocities of HVCs typically vary by < 5 km s $^{-1}$ over much of their projected surface (Wakker et al. 2008), suggesting that if the associated low-ionization absorption were probed along multiple close sightlines, it would exhibit minimal velocity shear. In principle, analysis of the coherence of low-ion absorption toward halo stars or AGNs (e.g., André et al. 2004; van Loon et al. 2009, 2013; Nasouli-Shoar et al. 2010; Smoker et al. 2015) has the greatest potential of any of the techniques discussed here to reveal the detailed density structure of circumgalactic clouds, and such efforts are ongoing (J. K. Werk et al. 2018, in preparation).

3.2.2. Gravitationally Lensed QSOs

Gravitationally lensed images of a single QSO can provide powerful constraints on the coherence of the intervening gas. Depending on the properties of the lensing system and its configuration relative to the source, the QSO images may appear separated over a wide range of physical scales extending to several tens of kpc. The assembly of a large sample of such systems can therefore offer much stronger leverage on the sizes of f/g absorbers than the beams of unlensed b/g galaxies. Lensed QSOs are also often much brighter than the brightest b/g galaxies that may be selected for studies such as this; hence, they permit high-S/N, high spectral resolution observations probing the detailed kinematic coherence of the material.

Early spectroscopic studies of gravitationally lensed QSOs focused primarily on the size scales of Ly α and C IV absorbers at $z \sim 1–2$ (e.g., Young et al. 1981; Weymann & Foltz 1983; Foltz et al. 1984; Smette et al. 1992). Using spectroscopy of QSO images with projected separations ranging from $R_{\perp} \sim 1 h_{50}^{-1}$ to $35 h_{50}^{-1}$ kpc, these works emphasize the strong similarity between the equivalent widths of these species observed in adjacent sightlines. In one of the first works to focus on constraining the coherence of low-ion metal absorption using this technique, Smette et al. (1995) reported the identification of five Mg II absorbers along doubly lensed QSO sightlines separated by $\sim 10–25 h_{50}^{-1}$ kpc. None of the systems having $W_{2796} < 0.4 \text{ \AA}$ were detected in more than one sightline. The two stronger systems (with $W_{2796} = 0.5–1.0 \text{ \AA}$) yielded absorption that varied by $\sim 30\%–40\%$ from one sightline to the other, suggestive of a scenario in which both sightlines probe the inner CGM of an associated intervening halo (at $R_{\perp} < 50$ kpc).

More sophisticated analyses, made possible with the advent of *HST* and high-resolution spectrographs on 10 m class telescopes, have gradually bolstered evidence for a picture in which weak Mg II absorbers are composed of clouds extending over less than a kiloparsec, whereas stronger absorbers ($W_{2796} \gtrsim 1 \text{ \AA}$) are coherent over larger scales (Monier et al. 1998; Rauch et al. 1999; Petitjean et al. 2000; Rauch et al. 2002; Churchill et al. 2003a; Ellison et al. 2004). For example, in their study of *HST*/FOS spectroscopy of the four sightlines to the Cloverleaf QSO at $z \sim 2.54$ (probing scales of $\sim 1–6 h_{100}^{-1}$ kpc), Monier et al. (1998) presented a detailed comparison of the equivalent widths across every sightline for the three intervening systems with detected low-ionization metal absorption. They found that most of the species with equivalent widths $\gtrsim 0.5 \text{ \AA}$ exhibited absorption along all four sightlines, and that in the strongest of these systems, the low-ion equivalent widths are very similar from sightline to sightline.

Following this study, Rauch et al. (1999, 2001, 2002) presented a series of works leveraging high-S/N spectroscopy obtained with Keck/HIRES to examine the spatial structure of gas clouds on subkiloparsec scales. These authors introduced the term “coherence length” to describe the distance over which there are significant changes in the physical parameters of a cloud (e.g., column density or projected velocity; Rauch et al. 2001). We note that this usage differs from our intended meaning in the foregoing text: we have invoked this phrase to describe a lack of variation in equivalent width, rather than a significant variation in cloud physical parameters.¹⁰ Rauch et al. (1999) discussed a single system at $z \sim 3.54$ observed along sightlines separated by only $26 h_{50}^{-1}$ pc. Variations in the column densities of C II and Si II between the two sightlines by factors of $\sim 2–10$ and velocity offsets of ~ 10 km s $^{-1}$ were uncovered, suggesting that the cool structures are composed of numerous tiny “cloudlets.” Rauch et al. (2002) then presented similar observations of three additional systems giving rise to Mg II and/or Fe II absorption at $z \sim 0.5–1$ along three sightlines separated by $\sim 0.2–0.7 h_{50}^{-1}$ kpc. The strongest system ($W_{2796} = 1.2 \text{ \AA}$ at $z = 0.566$) had a qualitatively similar component structure across all of these sightlines, although the precise velocity centroids of the components were observed to shift slightly, and most of the components exhibited column density variations of up to 1.5 dex. The components comprising

¹⁰ As discussed in Section 3.1.3, coherent W_{2796} does not imply coherent velocity structure.

the weaker systems, on the other hand, were frequently not detected along one or two of the sightlines.

The more recent work of Chen et al. (2014) is one of very few studies to use lensed QSO sightlines to probe within $R_{\perp} < 50$ kpc of bright f/g galaxies whose redshifts were identified a priori. These authors targeted a quadruply lensed QSO at $z \sim 1.7$ lying close behind two $\gtrsim L^*$ galaxies at $z = 0.42$ and 0.78 . Magellan/MagE and MIKE spectroscopy of the four sightlines, each separated by ~ 5 – 10 kpc at the redshifts of the f/g halos, revealed ubiquitously strong ($W_{2796} \gtrsim 0.5 \text{ \AA}$) absorption. In particular, the $z = 0.42$ halo exhibited notably coherent velocity profiles with a velocity shear of only $\Delta v \sim 20 \text{ km s}^{-1}$ across the sightlines. The $z = 0.78$ halo, on the other hand, gave rise to qualitatively distinct Mg II profiles with line widths differing by up to $\sim 180 \text{ km s}^{-1}$ and velocity centroids offset by $\sim 90 \text{ km s}^{-1}$. The authors considered several physical origins for the absorption, finally concluding that the properties of infalling gas streams are most consistent with the observed line profiles (and are more so than the putative properties of an extended rotating disk or biconical galactic winds). It is suggested that the overall coherence of the profiles would arise from streams $\gtrsim 10$ kpc in width, and that turbulent motion associated with the streams must contribute to the velocity centroid offsets between sightlines. Zahedy et al. (2016) also made use of lensed QSO sightlines to probe the inner CGM, targeting the halos of three $\gtrsim L^*$ elliptical galaxies at $R_{\perp} < 15$ kpc. In the most spectacular of these systems, the two sightlines of a double lens, separated by ~ 8 kpc, uncovered remarkably similar (and complex) velocity profiles with an overall velocity offset of $\sim 350 \text{ km s}^{-1}$. The authors liken this high level of coherence to that observed in the Chen et al. (2014) study, while also noting the difficulties in differentiating between prospective physical origins for the gas.

In summary, spectroscopy of gravitationally lensed QSO sightlines has begun to reveal the detailed spatial structure of low-ionization absorption in a wide range of environments. Structures giving rise to relatively weak absorption ($W_{2796} \lesssim 0.4 \text{ \AA}$) tend to exhibit substantial variations in equivalent width and/or column density and velocity on scales less than ~ 1 – 2 kpc (e.g., Monier et al. 1998; Petitjean et al. 2000; Rauch et al. 2002; Ellison et al. 2004; Rogerson & Hall 2012). We note that such systems are relatively rare in the inner CGM of $\sim L^*$ galaxies at $z < 1$, occurring along only four of the 22 QSO sightlines passing within $R_{\perp} < 30$ kpc of the galaxy samples discussed in Chen et al. (2010a) and Werk et al. (2013; see also Figure 11 of GPG1). Even at $30 \text{ kpc} < R_{\perp} < 50 \text{ kpc}$, $W_{2796} < 0.4 \text{ \AA}$ absorption occurs in only 14 of the 32 sightlines in these samples. Stronger systems, in contrast, tend to yield strong absorption and similar velocity structure in lensed QSO sightlines separated by $\gtrsim 1$ – 5 kpc (Monier et al. 1998; Rauch et al. 2002; Chen et al. 2014). These results are at least qualitatively consistent with our finding that the equivalent width of Mg II absorbers observed along our b/g galaxy sightlines does not vary on length scales < 1.9 kpc. Larger samples of bright, lensed QSOs will be needed to perform a more quantitative test for the consistency of all of these constraints (from, e.g., the Hyper Suprime-Cam Survey, More et al. 2017; or VST-ATLAS, Schechter et al. 2017).

Moreover, these studies may also, in principle, be used to differentiate between scenarios (1) and (2) as described in

Figure 10 and Section 3.1.3. At present, the coherence of the CGM structures reported in Chen et al. (2014) is evocative of the scenario diagrammed in panel (1) of this figure and tends to disfavor scenario (2). We expect that expanded absorption-line studies using both lensed QSO and b/g galaxy spectroscopy will be crucial to ruling out and/or refining these models.

3.2.3. Photoionization Modeling

The sizes of cool circumgalactic gaseous structures are also manifest in the ionization state of the gas. The thickness of an absorbing cloud affects the extent to which it is penetrated by ionizing radiation, which in turn affects the cloud’s ionization fraction (e.g., Bergeron 1986; Donahue & Shull 1991; Ferland et al. 1998). Thus, under the assumption of a particular ionizing radiation field, cloud configuration, and metal abundance pattern, measurements of the column densities of rest-frame UV metal transitions spanning a range of ionization states can yield an estimate of the cloud size (e.g., Churchill & Charlton 1999; Rigby et al. 2002; Lan & Fukugita 2017).

Over the last several years, this technique has frequently been invoked in the context of $z \sim 2$ – 3 CGM studies, for which many of the relevant transitions are accessible in the optical. Crighton et al. (2013) analyzed absorption detected at $R_{\perp} = 58$ kpc from a bright star-forming galaxy at $z = 2.44$. Their high-S/N, high spectral resolution coverage of the Lyman series permitted very precise constraints on the neutral hydrogen column density in this system. The properties of one particularly strong velocity component of this absorber (having $N(\text{Mg II}) = 10^{13.2} \text{ cm}^{-2}$), together with their assumption of a Haardt & Madau (2012) ionizing background and solar abundance ratios, strongly suggested a cloud thickness $\lesssim 3$ kpc.

Following this work, Crighton et al. (2015) reported on the detection of a partial LLS at $R_{\perp} = 50$ kpc from an $\sim 0.2 L^*$ galaxy also at $z \sim 2.5$. The authors carefully treated the systematic errors associated with photoionization modeling, introducing a variable slope for the input radiation field. They also used the MCMC technique to calculate marginalized probability distributions for each model parameter. Their analysis constrained the cloud thicknesses to be ~ 100 – 500 pc, with typical uncertainties less than a factor of ~ 2.5 . Similarly small (or even smaller) cloud sizes were recovered for the strong absorption systems detected in the circumgalactic environments of $z \sim 2$ QSOs from analysis of the associated C II* $\lambda 1335$ transition (Prochaska & Hennawi 2009; Lau et al. 2016).

While such subkiloparsec or subparsec clouds are apparently typical of the $z \gtrsim 2$ CGM, photoionization modeling of the CGM at lower redshifts suggests that cool clouds exhibit a much wider range of sizes at the current epoch (e.g., Stocke et al. 2013). Werk et al. (2014) performed photoionization modeling of the CGM absorbers around a sample of $z \sim 0.2$, $\sim L^*$ galaxies observed as part of the COS-Halos survey (Tumlinson et al. 2011, 2013; Werk et al. 2012, 2013), leveraging a data set containing 44 close ($R_{\perp} < 160$ kpc) sightlines. Among this sample, 33 sightlines exhibited low- or intermediate-ionization absorption, permitting tight constraints on the ionization state and metallicity of this subset. The resulting values of $N_{\text{H}}/n_{\text{H}}$ (the ratio of the total hydrogen column density to the hydrogen volume density) range between 0.1 and 2000 kpc with uncertainties of several orders of magnitude. However, those absorbers with the strongest Mg II (i.e., $N(\text{Mg II}) > 10^{13.5} \text{ cm}^{-2}$) tend to also have particularly

well-constrained neutral hydrogen column densities from fitting of the damping wings of the Ly α profile, such that the corresponding cloud sizes could be determined to within ± 1 dex. Moreover, these systems, which we suggest are close analogs to the strong Mg II absorbers we observe in the inner CGM in the present sample, have overall larger sizes of ~ 7 –100 kpc (J. K. Werk 2018, private communication).

All together, we view this growing body of photoionization modeling analyses as broadly consistent with the general picture supported by the lensed QSO studies discussed in Section 3.2.2: weak, low-ionization absorbers, and/or those at high redshift ($z > 1$; see also Lan & Fukugita 2017), tend to yield cloud thicknesses less than a kiloparsec (or have sizes that are weakly constrained), whereas stronger systems at $z < 1$ residing close to a massive galaxy exhibit thicknesses over ~ 5 –100 kpc scales. These results must, of course, be considered with the caveat that they are subject to substantial systematic uncertainties (e.g., in the strength and shape of the ionizing radiation field, the cloud metal abundance pattern, or the cloud geometry). In particular, Stern et al. (2016) recently demonstrated that a cloud structure allowing for multiple gas densities (rather than a constant-density slab, as assumed by Werk et al. 2014) can more closely match the full range of absorption strengths measured in the COS-Halos data set and would yield much smaller Mg II absorber sizes (~ 50 pc; Stern et al. 2016). However, there is little a priori observational evidence favoring this more complex geometry over the simpler structures assumed in most previous studies. Again, larger samples of absorbers detected along lensed QSO sightlines will be useful for differentiating among these models.

3.3. The Significance of Our Modeling Assumptions

Our analysis has relied on a number of significant assumptions regarding the nature of the Mg II-absorbing CGM. Here we discuss some of the weaknesses of these assumptions and refinements to our modeling to be pursued in future work.

3.3.1. The Fiducial CGM and Its Intrinsic Dispersion

As described at the beginning of Section 2, a fundamental assumption of our approach is the existence of a “fiducial” CGM—i.e., that the Mg II absorption characteristics of all galaxies with a given set of intrinsic properties (e.g., M_* , SFR) are equivalent. Indeed, this assumption is implicit in any study using the assembly of single-sightline b/g probes to analyze CGM absorption as a function of f/g host galaxy parameters. It is currently difficult to justify this simplification given the paucity of systems for which absorption along more than one b/g sightline may be analyzed. This is of particular concern given that cosmological “zoom” simulations predict a large degree of variation in mass outflow rates from a given galaxy on timescales < 100 Myr (although this effect is less pronounced at $z < 1$; e.g., Muratov et al. 2015).

The growing samples of gravitationally lensed QSO sightlines passing close to bright f/g galaxies (e.g., Chen et al. 2014; Zahedy et al. 2016) will eventually constrain the dispersion in W_{2796} within individual halos and hence will provide an independent estimate of our cosmic scatter parameter σ_C . A yet more powerful experimental design is that of Lopez et al. (2018), who used the bright, extended arc of

a gravitationally lensed b/g galaxy to study Mg II absorption in a single f/g halo along 56 independent sightlines extending from $R_{\perp} = 15$ to 90 kpc. The arc, sampled at ~ 2 –4 kpc spatial resolution, revealed significant scatter in W_{2796} from sightline to sightline. The magnitude of this scatter was reportedly smaller than that inferred from QSO–galaxy pair samples (Chen et al. 2010a; Nielsen et al. 2013); however, the authors caution that the latter included f/g galaxies with a broad range of intrinsic properties (as well as absorption-selected systems). A detailed comparison between such multi-sightline probes and QSO–galaxy pair sightlines with similar f/g galaxy properties will likewise provide highly valuable constraints on the intrinsic scatter of W_{2796} . Finally, advancements in UV-sensitive integral field spectrographs (e.g., the Keck Cosmic Web Imager; Morrissey et al. 2018) will soon make the direct detection of Ly α emission from diffuse halo material routine, permitting both constraints on the clumping scale of this emission in individual halos and comparisons of emission properties among large galaxy samples. While direct detection of Mg II in emission may be more challenging, it will likely be possible for at least a handful of $z \sim 0.3$ –1 systems (e.g., Rubin et al. 2011; Martin et al. 2013).

A second assumption important to our analysis is that there is no evolution in this fiducial CGM between $z \sim 0.35$ –0.8 (the epoch probed by our PRIMUS pair sample) and the lower-redshift regime within which most of our QSO–galaxy comparison systems lie ($z \sim 0.1$ –0.3). Stated more generally, our assumption is that both our PRIMUS pairs and QSO–galaxy pairs probe equivalent circumgalactic media as a function of M_* and R_{\perp} . One systematic bias to consider is the decline in the average specific SFR (sSFR) for galaxies of a given M_* over this period (Noeske et al. 2007; Speagle et al. 2014), which could, in principle, drive enhanced W_{2796} values (or enhance variability in W_{2796} values) for higher-redshift systems. In addition, the masses of the dark matter halos hosting galaxies of a given M_* decrease with cosmic time (e.g., Moster et al. 2013), which may affect the observed velocity dispersion of cool clouds embedded in these structures. Finally, the virial radius of a dark matter halo of a particular mass is redshift-dependent ($R_{\text{vir}} \propto (1+z)^{-2/3}$; Maller & Bullock 2004), such that our choice to compare CGM properties as a function of R_{\perp} implies that the comparison is being carried out at inconsistent values of R_{\perp}/R_{vir} .

Regarding the first issue of sSFR evolution, GPG1 demonstrated that there is a +0.5 dex offset between the median sSFR of the star-forming f/g galaxies in the PRIMUS pair and the QSO–galaxy pair samples consistent with the expected decline in star formation activity from $z \sim 0.4$ to 0.2. However, this work also demonstrated consistency in the median W_{2796} measured around PRIMUS and QSO–galaxy subsamples of f/g galaxies selected to span the same ranges in M_* . Thus, at the sensitivity of our data set, we find no evidence for a systematic offset in W_{2796} values due to differences in star formation activity; however, we cannot rule out the possibility of differing dispersions in W_{2796} arising from this evolution. A sample of QSO–galaxy pairs with f/g host properties more similar to those of the PRIMUS f/g galaxy sample will be needed to control for this effect.

Regarding issues related to differences in the dark matter halo masses and virial radii between these two samples, abundance matching analyses predict that the halo mass for a

central galaxy with $\log M_*/M_\odot = 10$ decreases by approximately 0.2 dex from $\log M_h/M_\odot \approx 11.7$ to $\log M_h/M_\odot \approx 11.5$ between $z = 0.5$ and 0.0 (Moster et al. 2013). Accounting for the difference in the median redshifts of the two samples ($z = 0.44$ versus $z = 0.25$), the corresponding change in halo virial velocity is only $\sim 20 \text{ km s}^{-1}$ ($v_{\text{vir}} \approx 117 \text{ km s}^{-1}$ versus 95 km s^{-1} ; Maller & Bullock 2004), implying a difference in the FWHM of the Mg II velocity distribution observed along the line of sight of $< 40 \text{ km s}^{-1}$. Any enhancement in W_{2796} driven by the larger velocity dispersion of the PRIMUS host halos will therefore likely be $\lesssim 0.2 \text{ \AA}$ (Ellison 2006). The difference in the virial radii of two halos with $\log M_h/M_\odot \approx 11.7$ and $\log M_h/M_\odot \approx 11.5$ at these two epochs is also relatively insignificant: R_{vir} decreases from $\approx 157 \text{ kpc}$ at $z = 0.44$ to 148 kpc at $z = 0.25$ (i.e., by only $\sim 6\%$). Thus, we do not consider evolution in host dark matter halo properties a major source of systematic uncertainty in our modeling.

More generally, an intrinsic dependence of CGM properties on any host galaxy property not included in our fiducial model (e.g., host galaxy orientation and/or azimuthal angle; see the Appendix) will cause us to overestimate the intrinsic scatter in W_{2796} . Ultimately, as mentioned above, significantly larger QSO–galaxy and galaxy–galaxy pair samples will be required to reveal such detailed dependencies and effectively mitigate this issue. We expect, however, that as long as (1) these intrinsic dependencies are consistent between both pair samples and (2) the f/g galaxies in these samples exhibit similar ranges in the relevant properties, the results of our analysis are likely insensitive to such systematics.

3.3.2. Absorber and b/g Beam Morphology

Our modeling has also made some critical assumptions regarding the morphology of Mg II absorbers. First, we have assumed that every absorber has the same projected size regardless of its strength (W_{2796}). However, the discussion presented in Sections 3.2.2 and 3.2.3 motivates the exploration of models with stronger absorbers having larger projected areas. Given the small size of the PRIMUS pairs data set and its limited S/N, we do not attempt to incorporate such refinements in the present work. We also note that this type of analysis will be best leveraged if detailed constraints on the UV continuum morphology of the b/g galaxy sample are available. Imaging of these objects with *HST* would enable a more precise prediction of the observed absorber strength given a particular f/g spatial distribution of W_{2796} . Finally, future efforts must certainly move beyond our invocation of square absorbers to accommodate a more generic morphology as suggested by cosmological simulations (e.g., Faucher-Giguère et al. 2015; Nelson et al. 2016; Fielding et al. 2017; Oppenheimer et al. 2018).

3.4. Implications for the Thermodynamics of the CGM

Recent advancements in our empirical picture of the media surrounding galaxies have forced us to confront a fundamental deficiency in our understanding of its governing physics. On the one hand, numerous surveys have pointed to the ubiquity of cool ($T \sim 10^4 \text{ K}$), photoionized material extending to $R_{\text{L}} > 100 \text{ kpc}$ from luminous galaxies at low redshift (e.g., Chen et al. 2010a; Nielsen et al. 2013; Stocke et al. 2013; Werk et al. 2013). On the other hand, the frequent detection of O VI absorption along the same sightlines (e.g., Tripp et al. 2008;

Tumlinson et al. 2011; Savage et al. 2014); the detection in X-ray imaging of extended, hot ($T \sim 10^6 \text{ K}$) halo material around nearby spirals (e.g., Anderson et al. 2016); and a strong theoretical expectation of virialized infall (Rees & Ostriker 1977; White & Rees 1978; Kereš et al. 2005) suggest that the CGM is also filled by a much hotter gaseous component.

The physics permitting the coexistence of these two (or more) temperature components is not understood. The photoionization modeling analysis of Werk et al. (2014) demonstrated that if the low-ionization absorption associated with the low-redshift CGM is assumed to arise from a single gas phase in ionization equilibrium at $T < 10^5 \text{ K}$, the implied volume density of this phase is two orders of magnitude too low to be in pressure equilibrium with a $T \sim 10^6 \text{ K}$ hot medium. One plausible resolution to this tension is the invocation of a range of densities for the cool phase (as in Stern et al. 2016). Or, as discussed in Werk et al. (2014), this apparent failure of an equilibrium solution may instead suggest that cool structures in the CGM are transient.

Indeed, there are several lines of evidence pointing in this latter direction. First, the cool gas phase is frequently observed to trace large-scale gas outflows from $\sim L^*$, star-forming galaxies at velocities $\gtrsim 200 \text{ km s}^{-1}$ (e.g., Weiner et al. 2009; Martin et al. 2012; Rubin et al. 2014). Second, inflows from the IGM or of previously ejected gas are expected to persist to $z < 1$ and involve material over a broad range of temperatures (e.g., Kereš et al. 2005; Nelson et al. 2013, 2015; Ford et al. 2016). Third, recent theoretical works addressing this question posit that a CGM that is out of pressure equilibrium is a natural consequence of energy input from galactic outflows (Thompson et al. 2016; Fielding et al. 2017). In particular, the hydrodynamic simulations of CGM evolution discussed in Fielding et al. (2017) predict that the gas in dark matter halos less massive than $10^{11.5} M_\odot$ never reaches hydrostatic equilibrium and is instead supported by ram pressure and turbulence associated with galactic fountain flows. Above this critical mass scale, the simulated gaseous halos are indeed supported by thermal pressure and feature a stable virial shock. However, a cool phase arises from cold gas driven into the halo via feedback, this cool wind seeding additional cooling, and/or (untriggered) thermal instability of the hot component.

Moreover, many (perhaps all) of these processes involve the streaming of cool material through a surrounding hot medium and hence will likely give rise to Rayleigh–Taylor or Kelvin–Helmholtz instabilities (Agertz et al. 2007; Schaye et al. 2007; Crighton et al. 2015; Armillotta et al. 2017; Fielding et al. 2017). These phenomena tend to disrupt the cool phase (in the absence of additional stabilizing mechanisms) and occur on physical scales that are far smaller than those typically resolved in the CGM of halo-scale hydrodynamic simulations (Crighton et al. 2015; Nelson et al. 2015; Fielding et al. 2017). It is therefore quite challenging to predict the overall contribution of these instabilities to the phase structure of halo gas. They would seem to make it yet more difficult to explain the coexistence of hot and cool media.

Establishing the sizes of cool structures in the CGM is of primary importance to resolving this outstanding issue. First of all, empirical constraints on the spatial extent of absorbing material greatly improve estimates of (or limits on) the total masses of the structures. Second, the sizes and implied volume densities of these clouds are directly related to the timescales over which they will be disrupted by hydrodynamic

instabilities. Given that nonequilibrium ionization conditions may be typical for the CGM (e.g., Oppenheimer et al. 2018), such that size constraints from photoionization models are subject to substantial systematic uncertainties, direct observation (using, e.g., extended b/g sources or lensed QSOs) is the most promising method for assessing cloud morphology. We use the remainder of this section to discuss our constraints on the mass and survival time of cool, Mg II–absorbing circumgalactic structures.

3.4.1. Mg II–absorbing Structure Mass

We have established a lower limit on the coherence length of Mg II absorbers of $\ell_A > 1.9$ kpc. As discussed in Section 3.1.3, our analysis requires that W_{2796} remain approximately constant over this scale. This requirement may be satisfied either by individual cool clouds that extend over $\ell_A > 1.9$ kpc or by a spatial correlation in the numbers and velocity dispersions of smaller clouds (respectively, scenarios (1) and (2) in Figure 10). Here we offer a calculation of the lower limit on the mass of these structures implied by our constraint on ℓ_A under the assumption that they are best described by the former scenario.

As our Mg II absorption lines are saturated and their velocity components unresolved, we do not report column densities for these systems. However, large W_{2796} values of >0.5 Å (as are typical for our sample) commonly arise from absorbers with column densities $N(\text{Mg II}) > 10^{13.0} \text{ cm}^{-2}$ (as shown in Section 3.1.2). If we assume that Mg II is the dominant ion in this gas phase (Churchill et al. 2003b; Narayanan et al. 2008), adopt a solar abundance ratio ($\log \text{Mg}/\text{H} = -4.42$; Savage & Sembach 1996), and assume Mg is not depleted by dust (a very conservative assumption; Jenkins 2009), our limit on the total hydrogen column density along these sightlines is $N(\text{H}) > 10^{17.4} \text{ cm}^{-2}$. Following Crighton et al. (2015), we estimate the total mass of each structure with the relation

$$M_A = \ell_A^3 n_{\text{H}} \mu m_{\text{p}},$$

with the hydrogen number density $n_{\text{H}} = N(\text{H})/\ell_A$ and μm_{p} the mass per hydrogen atom with $\mu = 1.4$. This equation may be rewritten as

$$M_A = 1.1 \times 10^4 M_{\odot} \left(\frac{\ell_A}{1.9 \text{ kpc}} \right)^2 \left(\frac{N(\text{H})}{10^{17.4} \text{ cm}^{-2}} \right).$$

Each structure thus likely contains at a minimum $\sim 10^4 M_{\odot}$ of material. This mass is modest in the context of HVCs (which are estimated to contain $\sim 10^8 M_{\odot}$, including their ionized component; Lehner & Howk 2011) but is similar to the cool cloud mass estimated from the photoionization analysis of CGM absorption at $z \sim 2.5$ by Crighton et al. (2015).

3.4.2. Cool Cloud Survival

In the absence of stabilizing mechanisms (e.g., magnetic fields; McCourt et al. 2015), Kelvin–Helmholtz instabilities will destroy these clouds on a timescale similar to the cloud-crushing time expressed in Equation (1). Again assuming that each individual cloud extends over the length scale ℓ_A , we may

rewrite this timescale as in Crighton et al. (2015),

$$t_{\text{crush}} \approx 19.6 \text{ Myr} \left(\frac{\ell_A}{1.9 \text{ kpc}} \right) \left(\frac{v}{300 \text{ km s}^{-1}} \right)^{-1} \left(\frac{n_{\text{H}}/n_{\text{halo}}}{10} \right)^{1/2}$$

(see also Agertz et al. 2007 and Schaye et al. 2007), with v representing the relative velocity between the cool and hot phases and $(n_{\text{H}}/n_{\text{halo}})$ equal to the ratio of their number densities.¹¹ As noted in Section 1, this relation suggests that larger, higher-density, slow-moving clouds will last longer in the hot halo environment.

The density ratio $(n_{\text{H}}/n_{\text{halo}})$ is the most poorly constrained of the relevant quantities in this context; however, we may refer to Werk et al. (2014) for guidance. They demonstrated that for a halo of mass $M_h = 10^{12} M_{\odot}$, the virialized hot phase as modeled in Maller & Bullock (2004) has $n_{\text{halo}} \approx 10^{-3} \text{ cm}^{-3}$ within $R_{\perp} < 50$ kpc. This number density is very similar to their empirical estimate of the density of the cool, photoionized phase. We therefore contend that potential variations in this ratio would likely lead to only a marginal decrease in the value of t_{crush} estimated above.

Regarding the velocity v , we expect this to depend on the source of the material. Large-scale outflows traced by Mg II absorbers similar in strength to those observed toward our b/g sightlines occur ubiquitously among massive, star-forming galaxies at $z \sim 0.5$ and exhibit velocities of $\gtrsim 200\text{--}400 \text{ km s}^{-1}$ (Rubin et al. 2014). If these high velocities persist as the material flows away from the galaxies, structures with $\ell_A \sim 1.9$ kpc would be disrupted after $\lesssim 20$ Myr and after having traveled only ~ 6 kpc. Structures as large as $\ell_A \sim 4$ kpc (i.e., similar in size to the optical half-light radii of the host galaxy sample) moving at such high speeds would survive for up to ~ 41 Myr but would reach distances of only $R_{\perp} = 13$ kpc in that time. It is worth noting that unless there is a pervasive mechanism that acts to suppress hydrodynamic instabilities in galaxy halos of this size, the Mg II absorption observed to be moving at hundreds of km s^{-1} when viewed down the barrel toward star-forming systems cannot also give rise to strong Mg II absorption at $R_{\perp} > 20$ kpc, even if the sizes of the absorbing structures are well above our lower limit on ℓ_A .

Such stabilizing mechanisms have by no means been ruled out by empirical evidence. On the contrary, magnetic fields are frequently invoked to suppress cool cloud destruction (e.g., McClure-Griffiths et al. 2010; McCourt et al. 2015) and are beginning to be recognized as a common phenomenon in nearby spiral galaxy halos due to recent advancements in radio continuum surveys (e.g., Krause 2009; Wiegert et al. 2015). Analyses of high-resolution hydrodynamical simulations of the interaction of cool clouds with a hot wind fluid suggest that both radiative cooling and thermal conduction can inhibit the mixing of the two phases (Mellema et al. 2002; Cooper et al. 2009; Armillotta et al. 2017). These mechanisms may not always act in tandem, however; for instance, thermal conduction may itself also be suppressed by magnetic fields (McCourt et al. 2018). The internal structure of the clouds may also aid in prolonging their survival: as discussed in McCourt et al. (2018), such clouds may be prone to shattering into numerous overdense fragments as they cool, which could, in turn, make

¹¹ The length scale used in this equation by Crighton et al. (2015) refers to the radius of a spherical cloud. The factor required to adapt this relation to our assumption of cubical clouds is of order unity, and we ignore it here.

them more robust to disruption than a monolithic structure with uniform density.

On the other hand, the observed Mg II absorbers may instead originate in diffuse accretion streams and/or gas that is stripped from infalling satellite galaxies via tides or ram pressure (e.g., Fumagalli et al. 2011; Stewart et al. 2011, 2013; Ho et al. 2017). Qualitative examination of cosmological simulations predicting the presence of these inflows suggests that they would indeed be coherent over scales of $\gtrsim 2$ kpc (Stewart et al. 2011, 2013), satisfying our cloud size constraint. Moreover, several streams that overlap along the line of sight could yield the large observed values of W_{2796} . Yet another possible origin is the precipitation of Mg II-absorbing gas due to thermal instability of the hot virialized halo medium (Maller & Bullock 2004; Voit et al. 2015; Fielding et al. 2017). In either of these latter scenarios, we expect that the cool structures would be moving through their host halos at speeds comparable to the freefall velocity. For a halo with mass $M_h \sim 10^{11-12} M_\odot$ at $z \sim 1$, this velocity is predicted to be $\sim 50\text{--}150 \text{ km s}^{-1}$ (Goerdt & Ceverino 2015). From the equation for t_{crush} above, a cool structure with $\ell_A \sim 1.9$ kpc and a velocity of 100 km s^{-1} would survive without disruption over $\gtrsim 60$ Myr. As infalling streams may be significantly more extended than ~ 2 kpc, their survival time may plausibly be much longer.

Ideally, a hydrodynamic simulation of a forming galaxy would be used to test these competing scenarios. Indeed, insight gained from simulations of individual cool structures interacting with an ambient hot flow has grounded much of the foregoing discussion (e.g., Agertz et al. 2007; Heitsch & Putman 2009; McCourt et al. 2015, 2018; Armillotta et al. 2017). However, resolving the hydrodynamic instabilities on the surfaces of these structures may be beyond the capability of any current idealized or cosmological zoom simulation. Agertz et al. (2007) found that simulations using the adaptive mesh refinement (AMR) technique require at least seven grid cells per cloud radius to properly capture cold cloud evolution. Assuming the clouds have sizes near the limit of our coherence constraint ($\ell_A = 1.9$ kpc), this implies that cells < 140 pc across are needed. Recent AMR simulations do indeed have cells of approximately this size, but only in the highest-density regions (i.e., in the galactic disk; Hummels et al. 2013; Agertz & Kravtsov 2016). Simulations making use of smoothed particle hydrodynamics (SPH) codes require at least 7000 particles per cloud in order to adequately resolve surface instabilities (Crighton et al. 2015) and would therefore need to adopt particle masses $< 2 M_\odot$ to resolve a structure of $M_A = 10^4 M_\odot$. This limit is well below the baryonic particle mass of state-of-the-art SPH simulations (e.g., Hopkins et al. 2018; Muratov et al. 2017; Oppenheimer et al. 2018). The shortfall in our current capability to resolve diffuse structures in detail has led some to call for the implementation of “subgrid” recipes to approximate the complex phase structure of the CGM (e.g., Crighton et al. 2015). Advancements in both simulation techniques and in empirical characterization of these structures are needed to ultimately reveal their physical origin, lifetime, and role in feeding galactic star formation.

4. Conclusions and Future Directions

With the aim to characterize the small-scale structure of the cool, photoionized gas phase of the CGM, we have presented a detailed analysis of Mg II absorption measured along sightlines

to bright b/g galaxies probing the environments within projected distances $R_\perp < 50$ kpc of 27 star-forming f/g galaxies at redshifts $0.35 \lesssim z_{f/g} \lesssim 0.80$ and having stellar masses $9.1 < \log M_*/M_\odot < 11.1$. The galaxy pair sample, first described in GPG1, was drawn from PRIMUS (Coil et al. 2011) and includes b/g galaxies as faint as $B_{\text{AB}} \lesssim 22.3$ over a redshift range $0.4 < z_{b/g} < 1.3$. Rest-frame near-UV spectroscopy of the sample from GPG1 permits constraints on the Mg II $\lambda 2796$ absorption strength associated with the f/g systems to a limiting equivalent width $W_{2796} \gtrsim 0.5 \text{ \AA}$ in individual b/g galaxy spectra. Moreover, unlike the QSO sightlines typically used to probe circumgalactic material in absorption, the b/g galaxies for which *HST* imaging is available are spatially extended with half-light radii $1.0 \text{ kpc} < R_{\text{eff}}(z_{f/g}) < 7.9 \text{ kpc}$.

Our analysis also leverages a sample of W_{2796} values measured in b/g QSO sightline spectroscopy probing the halos of f/g galaxies with a similar range in stellar mass (M_*) at $z \sim 0.1\text{--}0.3$ from Chen et al. (2010a) and Werk et al. (2013). By making the assumptions that (1) the Mg II-absorbing CGM exhibits the same W_{2796} distribution for every host galaxy of a given M_* as a function of R_\perp and (2) the quantity $\log W_{2796}$ has a Gaussian distribution with a constant dispersion (i.e., with a dispersion that does not depend on M_* or R_\perp), we use this QSO–galaxy pair sample to construct a “fiducial” model for the $\log W_{2796}$ distribution in the CGM of low-redshift galaxies.

Then, by adopting the assumption that all Mg II absorbers have the same projected surface area (A_A) regardless of their strength, we use this fiducial model to show how the $\log W_{2796}$ distribution observed along a given set of sightlines depends on the ratio of the surface area of the b/g probes (A_G) to that of the absorbers ($x_A \equiv A_G/A_A$). We compare these model distributions to the sample of $\log W_{2796}$ values measured toward PRIMUS b/g galaxy sightlines, rejecting the null hypothesis that the observed and modeled distributions are drawn from the same parent population for values of the ratio $x_A \geq 15$ at 95% confidence. This limit, in combination with the observed distribution of b/g galaxy sizes, requires that the coherence scale of Mg II absorption—that is, the length scale over which W_{2796} does not vary—is $\ell_A > 1.9$ kpc. This is the first such constraint on the morphology of cool, photoionized structures in the inner CGM (within $R_\perp < 50$ kpc) of $\sim L^*$ galaxies at $z < 1$.

While this limit does not necessarily imply that the absorbing material occupies structures extending over this length scale, it requires that regions of the CGM giving rise to a particular observed W_{2796} are spatially correlated. Complementary experiments using gravitationally lensed QSO sightlines to probe strong Mg II absorbers (that are similar in strength to those observed in the inner CGM) likewise suggest that these systems exhibit similar velocity structures and absorption strengths over scales $\gtrsim 1\text{--}5$ kpc (e.g., Monier et al. 1998; Rauch et al. 2002; Chen et al. 2014). And although photoionization modeling of metal-rich absorbers identified in $z \gtrsim 2$ galaxy halos tends to yield smaller cloud sizes ($\lesssim 0.5$ kpc; e.g., Crighton et al. 2015; Lau et al. 2016), such modeling of circumgalactic absorption in the low-redshift universe is suggestive of more extended structures consistent with our coherence limit (Werk et al. 2014).

The primary factors limiting the strength of our constraints are the small sizes of the QSO–galaxy pair and projected galaxy pair samples and the relatively low S/N of our b/g

galaxy spectroscopy. Within the next few years, however, ongoing and prospective large-scale surveys promise dramatic improvements in these sample sizes. The SDSS-IV/eBOSS program will nearly quadruple the surface density of known QSOs at $0 < z < 1.6$ relative to that uncovered by earlier phases of the SDSS by 2020 (Dawson et al. 2016; Abolfathi et al. 2018; Lan & Mo 2018). Beginning in 2018, the Dark Energy Spectroscopic Instrument will pursue spectroscopy over $14,000 \text{ deg}^2$ of sky, obtaining samples of ~ 100 QSOs per deg^2 at $1 < z < 2$ and ~ 725 emission-line-selected galaxies per deg^2 at $0.6 < z < 1.0$ (DESI Collaboration et al. 2016). This unprecedented data set will therefore yield over 20,000 projected pairs of these objects with angular separations $< 10''$, effectively increasing the sample size of QSO–galaxy pairs included in the present work by more than two orders of magnitude.

Significant expansion of the projected galaxy pair sample will be more costly but nonetheless will also be within reach over the next few years. A crucial component will be the identification of numerous additional galaxies that are sufficiently bright to serve as effective b/g light sources within reasonable integration times with current facilities (i.e., Keck/LRIS and VLT/FORS2). The ongoing Dark Energy Survey (DES; Drlica-Wagner et al. 2018) and Dark Energy Camera Legacy Survey (DECaLS¹²) will permit the photometric selection of galaxy pair candidates over $14,000 \text{ deg}^2$ to a limiting magnitude of $g \lesssim 24.0$. The surface density of galaxy pairs with $R_{\perp} < 50 \text{ kpc}$ and b/g galaxies having $B_{\text{AB}} < 22.3$ identified in PRIMUS is $\sim 10 \text{ deg}^{-2}$, implying that at least 100,000 such targets could be discovered within the DES and DECaLS footprints. With a focus on those pairs with the brightest b/g objects, the present sample of ~ 30 pairs could be tripled within a modest Keck/LRIS allocation (Lee et al. 2014). Such improvements would add substantial leverage to constraints on models invoking a more realistic (and complex) relationship between absorber morphology and strength than that explored here.

Moreover, 30 m class ground-based telescopes will enable high-S/N spectroscopy of a significantly fainter galaxy population. With the Large Synoptic Survey Telescope poised to provide extremely deep, multiband imaging over $18,000 \text{ deg}^2$ (Ivezić et al. 2008), and because such galaxies have a much higher sky density than QSOs, this technique will ultimately become the primary method used to probe the $z \gtrsim 0.3$ CGM (Steidel et al. 2009). It therefore behooves us to carefully assess the effects of using extended b/g sources on the relevant observables (i.e., equivalent width, column density, velocity structure, etc.) by pursuing detailed comparisons with pencil-beam surveys. These efforts will both optimize future experiments and offer novel constraints on the small-scale structure of the cool phase of the gas that pervades galactic environments.

KHRR acknowledges support for this project from the Clay Postdoctoral Fellowship. ALC acknowledges support from NSF CAREER award AST-1055081.

These findings are in part based on observations collected at the European Organisation for Astronomical Research in the Southern Hemisphere under ESO programs 088.A-0529(A) and 090.A-0485(A). In addition, a significant portion of the data presented herein were obtained at the W. M. Keck

Observatory, which is operated as a scientific partnership among the California Institute of Technology, the University of California, and the National Aeronautics and Space Administration. The Observatory was made possible by the generous financial support of the W. M. Keck Foundation.

It is our pleasure to thank Brice Ménard, James Bullock, Nicolas Tejos, Jessica Werk, Andrew Fox, and Nicolas Bouché for enlightening discussions that improved this analysis. We gratefully acknowledge J. X. Prochaska for his reading of this manuscript and for sharing valuable insight on this topic. We also thank the anonymous referee, whose suggestions helped to improve this work.

Finally, the authors wish to recognize and acknowledge the very significant cultural role and reverence that the summit of Maunakea has always had within the indigenous Hawaiian community. We are most fortunate to have the opportunity to conduct observations from this mountain.

Appendix A

The Significance of f/g Galaxy Orientation

The analysis presented in Section 2 rests on the assumption that the fluctuation of W_{2796} within a halo hosting a galaxy of a particular M_* at a given R_{\perp} is independent of all other intrinsic galaxy properties. However, several studies now suggest that W_{2796} may depend to some degree on the placement of the b/g sightline relative to the orientation of the f/g host’s stellar disk. This issue was first addressed by Bordoloi et al. (2011), who studied Mg II absorption in stacked b/g galaxy sightlines probing the halos of f/g edge-on disk-dominated galaxies. These authors subdivided their sample by the azimuthal angle (Φ) of the b/g sightline, i.e., the angle between the f/g disk minor axis and the location of the b/g galaxy on the sky (taking the center of the f/g galaxy to be the origin). They found that in sightlines with $\Phi < 45^\circ$ (located along the minor axis of the f/g disk) and $R_{\perp} < 40 \text{ kpc}$, the median equivalent width of the blended Mg II doublet was enhanced by $\sim 0.8 \text{ \AA}$ relative to sightlines with $\Phi > 45^\circ$. This enhancement was interpreted as a signature of bipolar galactic winds.

Bouché et al. (2012) and Kacprzak et al. (2012) also investigated this issue using a sample of projected QSO–galaxy pairs, the latter reporting a 20%–30% enhancement in the covering fraction of $W_{2796} > 0.1 \text{ \AA}$ absorption close to both the minor and major axes of star-forming galaxies and finding that the strongest absorbers were detected in sightlines having Φ within $< 50^\circ$ of the minor axis at $R_{\perp} < 40 \text{ kpc}$. More recently, Lan et al. (2014) assessed the numbers of edge-on star-forming galaxies within $R_{\perp} < 50 \text{ kpc}$ of SDSS QSO sightlines exhibiting $W_{2796} > 1.5 \text{ \AA}$ absorbers, finding that there were more of these galaxies oriented such that the corresponding QSO sightline probed their minor axes. Taken together, these measurements suggest that the incidence of the strongest Mg II systems depends on the placement of the b/g sightlines relative to the f/g galaxy disks.

We therefore consider here whether this dependence should be included in our fiducial model for the Mg II–absorbing CGM. To determine whether a minor-axis enhancement of W_{2796} is exhibited by our QSO–galaxy comparison sample, we make use of the photometry of the f/g galaxies reported in SDSS Data Release 10 (DR10; Ahn et al. 2014).¹³ We use the results of the fits of an exponential disk model to the SDSS

¹² <http://legacysurvey.org/decamls>

¹³ <http://cas.sdss.org/dr10/en/home.aspx>




Figure 11. The $\log W_{2796}$ vs. the azimuthal angle of the b/g sightline for star-forming f/g hosts within $R_{\perp} < 50$ kpc and having inclinations $i > 50^{\circ}$ in our QSO–galaxy comparison sample. Pairs with f/g galaxies having low, intermediate, and high values of SFR are indicated with magenta circles, orange stars, and green squares. Sightlines with $\Phi \sim 0^{\circ}$ and $\sim 90^{\circ}$ are close to the major and minor axes, respectively.

r -band imaging of these galaxies as indicators of the disk axis ratios (b/a) and position angles. Of the 50 QSO–galaxy pairs with $R_{\perp} < 50$ kpc and star-forming f/g hosts in this sample, SDSS DR10 photometry is available for 48. We adopt the simple assumption that the inclination of each galaxy is given by $i = \arccos(b/a)$ and calculate Φ from the angle between each disk position angle and the corresponding QSO coordinate. Note that this reference frame is different from that used in Bordoloi et al. (2011), in that low values of Φ indicate sightlines along the disk major axes.

We show the distribution of $\log W_{2796}$ versus Φ for all star-forming f/g galaxies having inclinations $i > 50^{\circ}$ in Figure 11, with the symbols color-coded by SFR as in Figure 1. Only three of the 13 systems with $\Phi > 40^{\circ}$ have $W_{2796} > 1.0 \text{ \AA}$, while four of the 16 systems with $\Phi < 40^{\circ}$ meet this criterion. Indeed, we see no evidence for a significantly higher incidence of strong absorbers at high Φ in this sample.

There are several important differences between the experimental designs of the studies mentioned above and that of the analysis shown in Figure 11 that may give rise to the apparent inconsistency of these results. For instance, the Bordoloi et al. (2011) study focused on f/g hosts at $z \sim 0.7$, which are likely to have higher SFRs than the f/g systems in our QSO–galaxy pair sample and hence may drive stronger large-scale winds, leading to a stronger enhancement in bipolar Mg II absorption. Our sample may simply be too small to reveal a significant dependence of W_{2796} on Φ : whereas Bordoloi et al. (2011) included 54 pairs with edge-on f/g galaxies and with QSO sightlines passing within 45° of the disk minor axis at $R_{\perp} < 50$ kpc, our QSO–galaxy pair sample contains only ~ 11 systems with these properties. Moreover, the Bouché et al. (2012) study included only QSO–galaxy pairs exhibiting strong Mg II absorption and could not assess the azimuthal angle distribution of weak absorbers. The f/g galaxies discussed in Kacprzak et al. (2012) have a wide range of redshifts ($0.1 \lesssim z \lesssim 1.1$); in addition, about $\sim 50\%$ of their f/g sample is Mg II absorption-selected and hence may be biased to exhibit larger W_{2796} overall.

In any case, a larger sample of high-S/N QSO sightline spectroscopy probing f/g galaxy halos at $R_{\perp} < 50$ kpc and selected without regard for the halo absorption strength is needed before we may carry out a sensitive test of the azimuthal angle dependence of W_{2796} at $z \sim 0.2$. Given this state of affairs, we conclude that the inclusion of such a dependence for the W_{2796} distribution in our fiducial CGM model is unjustified.

ORCID iDs

Kate H. R. Rubin  <https://orcid.org/0000-0001-6248-1864>

Alison L. Coil  <https://orcid.org/0000-0002-2583-5894>

References

- Abolfathi, B., Aguado, D. S., Aguilar, G., et al. 2018, *ApJS*, 235, 42
- Adelberger, K. L., Shapley, A. E., Steidel, C. C., et al. 2005, *ApJ*, 629, 636
- Agertz, O., & Kravtsov, A. V. 2016, *ApJ*, 824, 79
- Agertz, O., Moore, B., Stadel, J., et al. 2007, *MNRAS*, 380, 963
- Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2014, *ApJS*, 211, 17
- Anderson, M. E., Churazov, E., & Bregman, J. N. 2016, *MNRAS*, 455, 227
- André, M. K., Le Petit, F., Sonnentrucker, P., et al. 2004, *A&A*, 422, 483
- Armillotta, L., Fraternali, F., Werk, J. K., Prochaska, J. X., & Marinacci, F. 2017, *MNRAS*, 470, 114
- Barger, A. J., Cowie, L. L., & Wang, W.-H. 2008, *ApJ*, 689, 687
- Ben Bekhti, N., Richter, P., Winkel, B., Kenn, F., & Westmeier, T. 2009, *A&A*, 503, 483
- Bergeron, J. 1986, *A&A*, 155, L8
- Berti, A. M., Coil, A. L., Behroozi, P. S., et al. 2017, *ApJ*, 834, 87
- Birnboim, Y., & Dekel, A. 2003, *MNRAS*, 345, 349
- Bordoloi, R., Lilly, S. J., Knobel, C., et al. 2011, *ApJ*, 743, 10
- Bouché, N., Hohensee, W., Vargas, R., et al. 2012, *MNRAS*, 426, 801
- Chen, H., Helsby, J. E., Gauthier, J., et al. 2010a, *ApJ*, 714, 1521
- Chen, H.-W., Gauthier, J.-R., Sharon, K., et al. 2014, *MNRAS*, 438, 1435
- Chen, H.-W., Wild, V., Tinker, J. L., et al. 2010b, *ApJL*, 724, L176
- Churchill, C. W., & Charlton, J. C. 1999, *AJ*, 118, 59
- Churchill, C. W., Mellon, R. R., Charlton, J. C., & Vogt, S. S. 2003a, *ApJ*, 593, 203
- Churchill, C. W., Rigby, J. R., Charlton, J. C., & Vogt, S. S. 1999, *ApJS*, 120, 51
- Churchill, C. W., & Vogt, S. S. 2001, *AJ*, 122, 679
- Churchill, C. W., Vogt, S. S., & Charlton, J. C. 2003b, *AJ*, 125, 98
- Coil, A. L., Blanton, M. R., Burles, S. M., et al. 2011, *ApJ*, 741, 8
- Cooke, J., & O’Meara, J. M. 2015, *ApJL*, 812, L27
- Cool, R. J., Moustakas, J., Blanton, M. R., et al. 2013, *ApJ*, 767, 118
- Cooper, J. L., Bicknell, G. V., Sutherland, R. S., & Bland-Hawthorn, J. 2009, *ApJ*, 703, 330
- Crighton, N. H. M., Hennawi, J. F., & Prochaska, J. X. 2013, *ApJL*, 776, L18
- Crighton, N. H. M., Hennawi, J. F., Simcoe, R. A., et al. 2015, *MNRAS*, 446, 18
- Davidson-Pilon, C., Kalderstam, J., Kuhn, B., et al. 2018, CamDavidsonPilon/lifelines, Zenodo, 10.5281/zenodo.1252342
- Dawson, K. S., Kneib, J.-P., Percival, W. J., et al. 2016, *AJ*, 151, 44
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036
- Diamond-Stanic, A. M., Coil, A. L., Moustakas, J., et al. 2016, *ApJ*, 824, 24
- Donahue, M., & Shull, J. M. 1991, *ApJ*, 383, 511
- D’Onghia, E., & Fox, A. J. 2016, *ARA&A*, 54, 363
- Drlica-Wagner, A., Sevilla-Noarbe, I., Rykoff, E. S., et al. 2018, *ApJS*, 235, 33
- Ellison, S. L. 2006, *MNRAS*, 368, 335
- Ellison, S. L., Ibata, R., Pettini, M., et al. 2004, *A&A*, 414, 79
- Faucher-Giguère, C.-A., Hopkins, P. F., Kereš, D., et al. 2015, *MNRAS*, 449, 987
- Ferland, G. J., Korista, K. T., Verner, D. A., et al. 1998, *PASP*, 110, 761
- Fielding, D., Quataert, E., McCourt, M., & Thompson, T. A. 2017, *MNRAS*, 466, 3810
- Foltz, C. B., Weymann, R. J., Roser, H.-J., & Chaffee, F. H., Jr. 1984, *ApJL*, 281, L1
- Ford, A. B., Davé, R., Oppenheimer, B. D., et al. 2014, *MNRAS*, 444, 1260
- Ford, A. B., Werk, J. K., Davé, R., et al. 2016, *MNRAS*, 459, 1745
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
- Fumagalli, M., Hennawi, J. F., Prochaska, J. X., et al. 2014, *ApJ*, 780, 74

- Fumagalli, M., Prochaska, J. X., Kereš, D., et al. 2011, *MNRAS*, **418**, 1796
- Goerdt, T., & Ceverino, D. 2015, *MNRAS*, **450**, 3359
- Haardt, F., & Madau, P. 2012, *ApJ*, **746**, 125
- Heitsch, F., & Putman, M. E. 2009, *ApJ*, **698**, 1485
- Ho, S. H., Martin, C. L., Kacprzak, G. G., & Churchill, C. W. 2017, *ApJ*, **835**, 267
- Hogg, D. W., Bovy, J., & Lang, D. 2010, arXiv:1008.4686
- Hopkins, P. F., Kereš, D., Oñorbe, J., et al. 2014, *MNRAS*, **445**, 581
- Kalberla, P. F. W., Wetzel, A., Kereš, D., et al. 2018, *MNRAS*, **480**, 800
- Hummels, C. B., Bryan, G. L., Smith, B. D., & Turk, M. J. 2013, *MNRAS*, **430**, 1548
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2008, arXiv:0805.2366
- Jenkins, E. B. 2009, *ApJ*, **700**, 1299
- Jones, T. W., Kang, H., & Tregillis, I. L. 1994, *ApJ*, **432**, 194
- Joung, M. R., Bryan, G. L., & Putman, M. E. 2012, *ApJ*, **745**, 148
- Kacprzak, G. G., Churchill, C. W., & Nielsen, N. M. 2012, *ApJL*, **760**, L7
- Kalberla, P. F. W., Burton, W. B., Hartmann, D., et al. 2005, *A&A*, **440**, 775
- Kereš, D., Katz, N., Fardal, M., Davé, R., & Weinberg, D. H. 2009, *MNRAS*, **395**, 160
- Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, *MNRAS*, **363**, 2
- Krause, M. 2009, *RMxAC*, **36**, 25
- Lan, T.-W., & Fukugita, M. 2017, *ApJ*, **850**, 156
- Lan, T.-W., Ménard, B., & Zhu, G. 2014, *ApJ*, **795**, 31
- Lan, T.-W., & Mo, H. 2018, *ApJ*, **866**, 36
- Lau, M. W., Prochaska, J. X., & Hennawi, J. F. 2016, *ApJS*, **226**, 25
- Lee, K.-G., Hennawi, J. F., White, M., et al. 2016, *ApJ*, **817**, 160
- Lee, K.-G., Hennawi, J. F., White, M., Croft, R. A. C., & Ozbek, M. 2014, *ApJ*, **788**, 49
- Lehner, N., & Howk, J. C. 2011, *Sci*, **334**, 955
- Lopez, S., Tejos, N., Ledoux, C., et al. 2018, *Natur*, **554**, 493
- Maller, A. H., & Bullock, J. S. 2004, *MNRAS*, **355**, 694
- Martin, C. L., Shapley, A. E., Coil, A. L., et al. 2012, *ApJ*, **760**, 127
- Lau, M. W., Shapley, A. E., Coil, A. L., et al. 2013, *ApJ*, **770**, 41
- McClure-Griffiths, N. M., Madsen, G. J., Gaensler, B. M., McConnell, D., & Schnitzler, D. H. F. M. 2010, *ApJ*, **725**, 275
- McCourt, M., Oh, S. P., O'Leary, R., & Madigan, A.-M. 2018, *MNRAS*, **473**, 5407
- McCourt, M., O'Leary, R. M., Madigan, A.-M., & Quataert, E. 2015, *MNRAS*, **449**, 2
- Mellema, G., Kurk, J. D., & Röttgering, H. J. A. 2002, *A&A*, **395**, L13
- Monier, E. M., Turnshek, D. A., & Lupie, O. L. 1998, *ApJ*, **496**, 177
- More, A., Lee, C.-H., Oguri, M., et al. 2017, *MNRAS*, **465**, 2411
- Moster, B. P., Naab, T., & White, S. D. M. 2013, *MNRAS*, **428**, 3121
- Morrissey, P., Matuszewski, M., Martin, D. C., et al. 2018, *ApJ*, **864**, 93
- Mshar, A. C., Charlton, J. C., Lynch, R. S., Churchill, C., & Kim, T.-S. 2007, *ApJ*, **669**, 135
- Muratov, A. L., Kereš, D., Faucher-Giguère, C.-A., et al. 2015, *MNRAS*, **454**, 2691
- Muratov, A. L., Kereš, D., Faucher-Giguère, C.-A., et al. 2017, *MNRAS*, **468**, 4170
- Narayanan, A., Charlton, J. C., Misawa, T., Green, R. E., & Kim, T.-S. 2008, *ApJ*, **689**, 782
- Nasouli-Shoar, S., Richter, P., de Boer, K. S., & Wakker, B. P. 2010, *A&A*, **520**, A26
- Nelson, D., Genel, S., Pillepich, A., et al. 2016, *MNRAS*, **460**, 2881
- Nelson, D., Genel, S., Vogelsberger, M., et al. 2015, *MNRAS*, **448**, 59
- Nelson, D., Vogelsberger, M., Genel, S., et al. 2013, *MNRAS*, **429**, 3353
- Nielsen, N. M., Churchill, C. W., & Kacprzak, G. G. 2013, *ApJ*, **776**, 115
- Noeske, K. G., Weiner, B. J., Faber, S. M., et al. 2007, *ApJL*, **660**, L43
- Oppenheimer, B. D., Crain, R. A., Schaye, J., et al. 2016, *MNRAS*, **460**, 2157
- Oppenheimer, B. D., Davé, R., Kereš, D., et al. 2010, *MNRAS*, **406**, 2325
- Oppenheimer, B. D., Segers, M., Schaye, J., Richings, A. J., & Crain, R. A. 2018, *MNRAS*, **474**, 4740
- Peek, J. E. G., Heiles, C., Douglas, K. A., et al. 2011, *ApJS*, **194**, 20
- Peeples, M. S., Werk, J. K., Tumlinson, J., et al. 2014, *ApJ*, **786**, 54
- Péroux, C., Rahmani, H., Arrigoni Battaia, F., & Augustin, R. 2018, *MNRAS*, **479**, L50
- Petitjean, P., Aracil, B., Srianand, R., & Ibata, R. 2000, *A&A*, **359**, 457
- Prochaska, J. X., & Hennawi, J. F. 2009, *ApJ*, **690**, 1558
- Prochaska, J. X., Werk, J. K., Worseck, G., et al. 2017, *ApJ*, **837**, 169
- Putman, M. E., de Heij, V., Staveley-Smith, L., et al. 2002, *AJ*, **123**, 873
- Rauch, M., Sargent, W. L. W., & Barlow, T. A. 1999, *ApJ*, **515**, 500
- Rauch, M., Sargent, W. L. W., & Barlow, T. A. 2001, *ApJ*, **554**, 823
- Rauch, M., Sargent, W. L. W., Barlow, T. A., & Simcoe, R. A. 2002, *ApJ*, **576**, 45
- Rees, M. J., & Ostriker, J. P. 1977, *MNRAS*, **179**, 541
- Richter, P., Westmeier, T., & Brüns, C. 2005, *A&A*, **442**, L49
- Rigby, J. R., Charlton, J. C., & Churchill, C. W. 2002, *ApJ*, **565**, 743
- Rogerson, J. A., & Hall, P. B. 2012, *MNRAS*, **421**, 971
- Rubin, K. H. R., Diamond-Stanic, A. M., Coil, A. L., Crighton, N. H. M., & Moustakas, J. 2018, *ApJ*, **853**, 95
- Rubin, K. H. R., Prochaska, J. X., Koo, D. C., et al. 2014, *ApJ*, **794**, 156
- Rubin, K. H. R., Prochaska, J. X., Koo, D. C., Phillips, A. C., & Weiner, B. J. 2010, *ApJ*, **712**, 574
- Rubin, K. H. R., Prochaska, J. X., Ménard, B., et al. 2011, *ApJ*, **728**, 55
- Saul, D. R., Peek, J. E. G., Grcevich, J., et al. 2012, *ApJ*, **758**, 44
- Savage, B. D., Kim, T.-S., Wakker, B. P., et al. 2014, *ApJS*, **212**, 8
- Savage, B. D., & Sembach, K. R. 1996, *ARA&A*, **34**, 279
- Schaye, J., Carswell, R. F., & Kim, T.-S. 2007, *MNRAS*, **379**, 1169
- Schechter, P. L., Morgan, N. D., Chehade, B., et al. 2017, *AJ*, **153**, 219
- Schneider, E. E., & Robertson, B. E. 2017, *ApJ*, **834**, 144
- Shen, S., Madau, P., Guedes, J., et al. 2013, *ApJ*, **765**, 89
- Smette, A., Robertson, J. G., Shaver, P. A., et al. 1995, *A&AS*, **113**, 199
- Smette, A., Surdej, J., Shaver, P. A., et al. 1992, *ApJ*, **389**, 39
- Smoker, J. V., Keenan, F. P., & Fox, A. J. 2015, *A&A*, **582**, A59
- Speagle, J. S., Steinhardt, C. L., Capak, P. L., & Silverman, J. D. 2014, *ApJS*, **214**, 15
- Steidel, C., Martin, C., Prochaska, J. X., et al. 2009, astro2010: The Astronomy and Astrophysics Decadal Survey, **286**
- Steidel, C. C., Erb, D. K., Shapley, A. E., et al. 2010, *ApJ*, **717**, 289
- Stern, J., Hennawi, J. F., Prochaska, J. X., & Werk, J. K. 2016, *ApJ*, **830**, 87
- Stewart, K. R., Brooks, A. M., Bullock, J. S., et al. 2013, *ApJ*, **769**, 74
- Stewart, K. R., Kaufmann, T., Bullock, J. S., et al. 2011, *ApJ*, **738**, 39
- Stocke, J. T., Keeney, B. A., Danforth, C. W., et al. 2013, *ApJ*, **763**, 148
- Thom, C., Peek, J. E. G., Putman, M. E., et al. 2008, *ApJ*, **684**, 364
- Thompson, T. A., Quataert, E., Zhang, D., & Weinberg, D. H. 2016, *MNRAS*, **455**, 1830
- Tripp, T. M., Sembach, K. R., Bowen, D. V., et al. 2008, *ApJS*, **177**, 39
- Tumlinson, J., Thom, C., Werk, J. K., et al. 2011, *Sci*, **334**, 948
- Tumlinson, J., Thom, C., Werk, J. K., et al. 2013, *ApJ*, **777**, 59
- van Loon, J. T., Bailey, M., Tatton, B. L., et al. 2013, *A&A*, **550**, A108
- van Loon, J. T., Smith, K. T., McDonald, I., et al. 2009, *MNRAS*, **399**, 195
- Voit, G. M., Bryan, G. L., O'Shea, B. W., & Donahue, M. 2015, *ApJL*, **808**, L30
- Wakker, B. P. 2001, *ApJS*, **136**, 463
- Wakker, B. P., York, D. G., Wilhelm, R., et al. 2008, *ApJ*, **672**, 298
- Weiner, B. J., Coil, A. L., Prochaska, J. X., et al. 2009, *ApJ*, **692**, 187
- Werk, J. K., Prochaska, J. X., Thom, C., et al. 2012, *ApJS*, **198**, 3
- Werk, J. K., Prochaska, J. X., Thom, C., et al. 2013, *ApJS*, **204**, 17
- Werk, J. K., Prochaska, J. X., Tumlinson, J., et al. 2014, *ApJ*, **792**, 8
- Weymann, R. J., & Foltz, C. B. 1983, *ApJL*, **272**, L1
- White, S. D. M., & Rees, M. J. 1978, *MNRAS*, **183**, 341
- Wiegert, T., Irwin, J., Miskolczi, A., et al. 2015, *AJ*, **150**, 81
- Young, P., Sargent, W. L. W., Oke, J. B., & Boksenberg, A. 1981, *ApJ*, **249**, 415
- Zahedy, F. S., Chen, H.-W., Rauch, M., Wilson, M. L., & Zabludoff, A. 2016, *MNRAS*, **458**, 2423