

The COS Absorption Survey of Baryon Harbors: Unveiling the Physical Conditions of Circumgalactic Gas through Multiphase Bayesian Ionization Modeling

Karl J. Haislmaier^{1*}, Todd M. Tripp¹, Neal Katz¹, J. Xavier Prochaska²,
Joseph N. Burchett³, John M. O’Meara⁴, Jessica K. Werk⁵

¹*Department of Astronomy, University of Massachusetts, Amherst, MA 01003*

²*Department of Astronomy, New Mexico State University, Las Cruces, NM 88001*

³*UCO/Lick Observatory, Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064*

⁴*W. M. Keck Observatory, Kamuela, HI 96743*

⁵*Department of Astronomy, University of Washington, Seattle, WA 98195-1580*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Quasar absorption systems encode a wealth of information about the abundances, ionization structure, and physical conditions in intergalactic and circumgalactic media. Simple (often single-phase) photoionization models are frequently used to decode such data. Using five discrete absorbers from the COS Absorption Survey of Baryon Harbors (CASBaH) that exhibit a wide range of detected ions (e.g., Mg II, S II – S VI, O II – O VI, Ne VIII), we show several examples where single-phase ionization models cannot reproduce the full set of measured column densities. To explore models that can self-consistently explain the measurements and kinematic alignment of disparate ions, we develop a Bayesian multiphase ionization modeling framework that characterizes discrete phases by their unique physical conditions and also investigates variations in the shape of the UV flux field, metallicity, and relative abundances. Our models require at least two (but favor three) distinct ionization phases ranging from $T \approx 10^4$ K photoionized gas to warm-hot phases at $T \lesssim 10^{5.8}$ K. For some ions, an apparently single absorption “component” includes contributions from more than one phase, and up to 30% of the H I is not from the lowest ionization phase. If we assume that all of the phases are photoionized, we cannot find solutions in thermal pressure equilibrium. By introducing hotter, collisionally ionized phases, however, we can achieve balanced pressures. The best models indicate moderate metallicities, often with sub-solar N/α , and, in two cases, ionizing flux fields that are softer and brighter than the fiducial Haardt & Madau UV background model.

Key words: galaxies: abundances — galaxies: evolution — galaxies: haloes — quasars: absorption lines — methods: data analysis

1 INTRODUCTION

An understanding of gas physics in the circumgalactic medium (CGM) of galaxies is vital for several areas of research. The spatial cycling of baryons and the fueling of star formation, the impact of active galactic nuclei (AGN) and stellar feedback on galaxy evolution, and the outcome of cosmological simulations all depend on the physical conditions in the CGM, and measurements of the abundances and characteristics of the CGM and IGM can provide insight into a variety of processes. Since halo gas is very challenging to study in emission with current facilities, ultraviolet (UV) absorption features imprinted by circumgalactic/intergalactic gases on the spectra of background QSOs provide the most practical means for probing the CGM in a wide variety of

contexts with statistically useful samples (e.g., Tumlinson et al. 2017; Rudie et al. 2019). Quasar absorption-line systems trace the dominant component of the baryon distribution at high redshifts (Weinberg et al. 1997; Rauch et al. 1997), and current evidence indicates that QSO absorbers reveal important baryon harbors at low redshifts as well (e.g., Shull et al. 1996; Tripp et al. 2000; Penton et al. 2004; Tumlinson et al. 2011; Werk et al. 2014; Burchett et al. 2019). In most CGM absorbers, the gas is optically thin and the hydrogen is >90% ionized, so ionization modeling is essential for deriving the chemical and thermodynamic properties of the gas/plasma, which are in turn critical variables for understanding its role in galaxy evolution.

To gain further insights there are several physical properties that are important to measure, including the gas density, temperature, pressure, and composition, as well as the ambient ionizing radiation field. The first three are related to the

* E-mail: khaislmaier@astro.umass.edu

physical structure of clouds in the CGM, including aspects such as the cloud size, pressure balance, mass, and energy content. The gas metallicity and *relative* abundance ratios (e.g., the abundance of nitrogen vs. α -group elements such as oxygen) are age-sensitive quantities, and hence trace the origin and evolution of gas and/or the presence of dust in the CGM, but interpreting abundances is complicated by the uncertain physics that govern the transport and mixing of galaxy ejecta into the CGM (see, e.g., [Frye et al. 2019](#)). The ionizing radiation field is significant cosmologically, but also for its systematic effect on the photoionization rate for all ions. However, typical ionization models only vary the overall metallicity and ionization parameter (a ratio between the number densities of ionizing photons and hydrogen atoms) and typically assume a solar gas composition with a single fixed incident radiation field from available models such as [Haardt & Madau \(2012\)](#), but it is well known that changes in the shape of the ionizing radiation field can significantly alter model results (e.g., [Crighton et al. 2015](#); [Chen et al. 2017](#); [Khaire & Srianand 2019](#); [Wotta et al. 2019](#)). The primary driver of the usual assumptions and simplifications of the models is not rooted in physics but rather in the inadequate number of absorption line diagnostics and model constraints that are typically recorded in a single absorption system. However, as we summarize in the following section, we are conducting a new survey (Tripp et al. 2021, *in prep.*) that provides a significant improvement in the number and quality of constraints available for individual QSO absorption systems. This larger number of constraints enables more detailed studies of absorber ionization models. In this paper we present a Bayesian Markov-Chain Monte-Carlo method for modeling these more observationally constrained absorbers, and we discuss the implications of these models for three absorption systems in particular, which present a variety of constraints and typify the data from our new program.

1.1 The COS Absorption Survey of Baryon Harbors

To investigate the physical nature and evolution of the CGM over a broad redshift range, we have initiated a program to study QSO absorption systems from $z = 0$ to $z \approx 1.5$ using ultraviolet spectra recorded with the Cosmic Origins Spectrograph (COS, [Green et al. 2012](#)) and the Space Telescope Imaging Spectrograph (STIS, [Woodgate et al. 1998](#)) on the *Hubble Space Telescope* (HST) as well as optical echelle spectra obtained with HIRES on Keck ([Vogt et al. 1994](#)). Our program — the COS Absorption Survey of Baryon Harbors (CASBaH, P.I. T. M. Tripp) — offers some unique advantages for probing the ionization, abundances, and physics of halo gas. This section reviews some of the motivations and strengths of CASBaH for the study of QSO absorption systems and the baryon cycle; additional information about the survey can be found in [Burchett et al. \(2019\)](#), [Prochaska et al. \(2019\)](#), and Tripp et al. (2021, *in prep.*).

The CASBaH program observed nine quasars with redshifts ranging from $z_{\text{QSO}} = 0.91573$ to $z_{\text{QSO}} = 1.47895$, and the spectra extend from observed wavelengths $\lambda_{\text{ob}} = 1152 \text{ \AA}$ to roughly the redshifted wavelength of the H I Ly α line at the quasar redshift [i.e., $(1 + z_{\text{QSO}}) \times 1215.67$]. The UV spectra in the full CASBaH dataset span *rest-frame* wavelengths extending from 465 to $> 3000 \text{ \AA}$, and the optical HIRES data extend the coverage farther into the red. As noted by [Verner](#)

et al. (1994) and illustrated in Figure 1, the UV rest-frame wavelength range below 912 \AA includes a rich suite of resonance absorption lines (see also Figure 2 in [Tripp 2013](#)), and access to this wavelength range provides unique diagnostics such as hot-gas tracers (e.g., Ne VIII) and, more importantly for this work, various banks of adjacent ions such as O I/O II/O III/O IV/O V/O VI, S I/S II/S III/S IV/S V/S VI, etc. that provide greatly improved constraints on ionization and physical condition models (see [Verner et al. 1996b](#) for a compilation of absorption line data $< 912 \text{ \AA}$). For comparison, most previous studies of oxygen in QSO absorbers have covered only O I and O VI. Without any information on O II – O V, it is difficult to draw precise conclusions about the gas ionization and modeling assumptions, which can lead to misleading results (see further discussion below). In this paper, we demonstrate the advantages gained in QSO absorber studies by 1) pushing deep into the far ultraviolet (FUV) and 2) using a Bayesian Markov-chain Monte-Carlo (MCMC) method for exploring the implications of the diagnostics afforded by the FUV coverage. In a follow-up paper, we will apply our methodology to the full set of CASBaH systems.

1.1.1 Challenges and Limitations of Typical Absorber Ionization Models

Since many studies are limited by the number of ions detected in an absorber (especially in the low-density CGM), necessary simplifications could lead to significant systematic errors. For example, in many studies the detected CGM lines are limited to species such as H I, C II, C IV, Si II, Si III, and Si IV. At lower redshifts, where there are fewer H I lines per unit redshift, transitions such as N II 1083.99 \AA , N III 989.80 \AA , C III 977.02 \AA , O VI 1031.93, 1037.62 can also be accessed. At higher redshifts (e.g., $z \gtrsim 2$), however, transitions at observed wavelengths $\lambda_{\text{ob}} < 1216 \text{ \AA}$ become increasingly difficult to measure because they are often severely blended with H I lines in the Ly α forest (or the spectrum is obliterated by an optically thick intervening Lyman limit system). While the typical sets of ions and elements employed in the majority of QSO absorber studies have established that the CGM is an important component of galaxies with some surprising characteristics (e.g., [Steidel et al. 2010](#); [Tumlinson et al. 2011](#)), many analyses of the ionization and physical conditions suffer from the following limitations.

1.1.2 Improved Constraints from Far-Ultraviolet Lines

First, important phases of the absorbing entity may be poorly constrained or not constrained at all. For example, if the only detected species are H I, C II, and C IV, the lack of knowledge of intermediate- and high-ionization stages can introduce serious systematic errors. Some papers choose to analyze such an ion set as a single-phase absorber modeled with a photoionization code such as CLOUDY ([Ferland et al. 2017](#)), but these systems are not necessarily single-phase “clouds”, and indeed many papers have presented evidence that QSO absorbers can arise in complex, multiphase media (e.g., [Charlton et al. 2003](#); [Ding et al. 2003](#); [Savage et al. 2005](#); [Tripp et al. 2008, 2011](#); [Burchett et al. 2015](#)). A single-phase photoionization model can typically fit the column densities of




Figure 1. Schematic summary of the redshift range over which detectable resonance lines of various species can be detected in intervening absorbers in the CASBaH spectra. The species shown are labeled on the right side of the diagram. For each species, a horizontal bar indicates the z range that is covered for a specific transition; the lower bound indicates the redshift at which the transition is redshifted into the COS FUV band (observed $\lambda > 1152$ Å), and the upper bound indicates the redshift where the line is redshifted beyond the long-wavelength end of the CASBaH NUV data; with one exception (see text), the NUV spectra all extend to at least 2400 Å. For example, we cover the C II 1036.34 Å line from $z_{\text{abs}} = 0.1097$ to at least 1.3168. However, for some of the sight lines, the NUV spectrum extends to even longer wavelengths, so that the upper redshift cutoff can be higher than this minimum range. The highest redshift that our survey can cover is set by the redshift of our highest-redshift QSO, PG1630+377 ($z = 1.4789$). For some species, we can access several resonance lines. For example, we have the ability to detect the C II lines at 543.26, 560.24, 594.80, 687.05, 903.63, 903.96, 1036.34, and 1334.53 Å. The coverage of each transition is indicated by a separate line unless the lines are quite close in wavelength (e.g., 903.63 and 903.96 Å), in which case a single line indicates the coverage.

H I, C II, and C IV quite nicely, but if the C IV actually arises in a separate phase, the assumption of a single-phase model will lead to an erroneous solution. The addition of species such as C III, Si III, and O VI to the analysis is helpful, but these additions can be problematic owing to line saturation and non-solar relative abundances (see below). Alternatively, it is sometimes assumed that only low-ionization stages can be safely assumed to be cospatial and well mixed with the H I, and some authors use photoionization models to fit H I, C II, Si II and ignore higher ionization stages that are also detected (e.g., C IV, Si IV) on the grounds that the higher ions are likely from a separate phase. This is also problematic because the C IV- and Si IV-bearing phases *can* contribute significantly to the total H I column density, and the assumption that 100% of the H I originates in the C II and Si II phase could lead to incorrect inferences about the physical conditions and metallicities. This problem is exacerbated by evidence that the absorption lines of intermediate ionization stages such as C III and Si III come from yet another phase that is distinct from both the C II phase and possibly also the C IV phase (see below). It is necessary to consider how the results might change if the detected H I and other ions are distributed amongst two, three, or even more phases.

Second, when only a small number of lines are covered, some potentially valuable constraints can be ruined by line saturation or blending. Column densities of species like C III and Si III can provide insights on multiphase absorbers, but most observations only record a single line from C III and Si III and, moreover, these single lines are very strong and prone to strong/severe saturation, which limits their usefulness. This problem is especially insidious in spectra from the Cosmic Origins Spectrograph on HST because the modest spectral resolution and broad wings of the COS line-spread function (Kriss 2011) can cause saturated lines to appear to be unsaturated, and the presence or absence of unresolved saturation can be difficult to establish with the coverage of only a single strong line. Similarly, a single line of an important species can be unmeasurable if it is badly blended with a strong absorption feature from a different redshift.

Third, many studies assume that the relative abundances of elements follow the solar pattern (e.g., from Lodders 2003 or Asplund et al. 2009). However, there is ample evidence of departures from solar abundances. In various contexts, including QSO absorbers arising in galaxy halos, there are clear indications that nitrogen can be underabundant by various amounts (e.g., Vila Costas & Edmunds 1993; van Zee & Haynes 2006; Pettini et al. 2008; Battisti et al. 2012; Berg et al. 2012), the alpha-group elements have well-known non-solar abundances compared to iron-group elements (e.g., McWilliam 1997; Prochaska et al. 2000; Wolfe et al. 2005), and the relative carbon abundance exhibits a complex dependence on metallicity in both stars and QSO absorbers (Akerman et al. 2004; Pettini et al. 2008; Fabbian et al. 2009; Penprase et al. 2010). Thus the assumption of solar relative abundances could also introduce substantial systematic errors if intercomparison of species such as C II, N II, Si II, and Fe II provide the primary model constraints.

Finally, ionizing models require other assumptions regarding, for example, the shape of the ionizing flux field impinging on the gas and the basic ionization mechanism (whether the absorber is photoionized or collisionally ionized or both). Often a single shape for the ionizing flux is assumed, but plau-

sible variations in the shape of the radiation field can significantly change the model outcomes, and the recent controversy regarding the shape and intensity of the UV background light (Kollmeier et al. 2014; Shull et al. 2015; Khaire & Srianand 2015, 2019; Khaire et al. 2019) illustrates the difficulties and uncertainties encountered in models of the extragalactic UV background.

To compactly summarize the ions, elements, and individual transitions that are available at any given redshift in the CASBaH spectra, Figure 1 shows the redshift range over which the UV spectra cover a given *detectable* transition of a given ion. Often, multiple transitions with a range of line strengths can be detected from certain species, which is advantageous for dealing with saturation and blending.¹ The UV spectra extend from $\lambda_{ob} = 1152 \text{ \AA}$ to at least 2400 \AA , and for the higher-redshift targets, the spectra extend to at least 2780 \AA . Thus, for example, both lines of the O VI doublet at $\lambda_r = 1031.93$ and 1037.62 \AA are covered from $z_{abs} = 0.11636$ to at least $z_{abs} = 1.31299$, as indicated in the bar labeled “O VI” in Figure 1 (for many species, the maximum absorber redshift in a given sight line is set by the QSO redshift, not the wavelength range of the data, so some of the sight lines cover less redshift path than shown in the figure). For some ions (e.g., C II and N II), many resonance lines are covered, and each line is indicated by a separate horizontal bar (unless the resonance-transition wavelengths are very close, e.g., C II 903.63 and 903.96, in which case both lines are indicated with a single bar).

A few remarks and caveats about Figure 1 and the survey data are worth noting here. First, as the absorber redshift (z_{abs}) increases, the number of species increases, but many of the ions that are often studied in the QSO absorber literature (e.g., Mg II and C IV) are redshifted beyond the long-wavelength limit of the CASBaH UV spectra at higher z_{abs} (Mg II is not shown in Fig. 1 because it is not covered at all in most of the UV spectra). By design, these favorite lines are often covered in our Keck HIRES data. Second, while some very highly ionized species in Figure 1 such as Na IX and Mg X have been detected at high significance in the CASBaH data, these mostly occur in “proximate” absorbers with $z_{abs} \approx z_{QSO}$ (Muzahid et al. 2013, 2016). Only Qu & Bregman (2016) have reported the detection of weak Mg X in one intervening ($z_{abs} \ll z_{QSO}$) CASBaH absorber. The Ne VIII doublet is the most highly ionized species frequently detected in the CASBaH spectra (Meiring et al. 2013; Burchett et al. 2019, Tripp et al. 2021 *in prep.*). Third, while species such as C I, N I, and O I have a large number of resonance transitions in the far-UV including lines with a broad range of strengths (see, e.g., Sembach et al. 2004), we do not show these species in Figure 1 because they are rarely detected in CASBaH spectra. This may owe to the “inner CGM” bias (Prochaska et al. 2017) – the CASBaH targets were selected to cover large paths for species such as Ne VIII, so sight lines with known optically thick Lyman-limit absorbers or low FUV GALEX fluxes were excluded from consideration.

For guidance in interpreting detections of various ions, Fig-

ure 2 graphically portrays the ionization potentials and photoionization rates of species that are commonly detected in CASBaH spectra. In each panel, the black squares represent the labeled species of interest (e.g., C II), and the red dot indicates the next lower ionization stage (e.g., C I). All of the points are presented in order of increasing ionization potential (upper panel) or photoionization rate (lower panel), and pairs of adjacent ions of a given element are connected with a line. This figure shows some factors that might lead certain species to coexist in the same gas phases. For example, there is a set of species with ionization potentials just above 1 Ryd (N II, O II, Al III, Si III, and S III), and similar energies are required to ionize all of these ions. Thus, we might expect these ions to trace similar conditions. The information in Figure 2 is useful for understanding the outcomes of the various ionization models discussed below.

2 DATA AND MEASUREMENTS

To develop our methods and begin to exploit the rich information provided by the CASBaH program, this paper focuses on three line-rich UV absorption systems found in high-quality QSO spectra from CASBaH in the range $0.4 < z < 0.7$. These are taken from a larger (>35) sample of H I-selected $\log N_{H\text{ I}} \geq 15 \text{ cm}^{-2}$ CASBaH intervening absorption systems that will be analyzed in a subsequent paper. Our primary goal is to select systems well-suited for developing and testing the methodology presented herein. To that end, we randomly selected three systems that (1) have robust, but optically thin ($\log N_{H\text{ I}} \lesssim 17 \text{ cm}^{-2}$), $N_{H\text{ I}}$ measurements, (2) are detected in a variety of metals, including C, N, and Fe in order to constrain models probing their relative abundance variations, (3) exhibit species across a range of ionization stages, including one system with only low-ions as well as two featuring both low- and high-ions, and (4) present some variety of component structure (to test how well the software performs in cases with simple vs. complex component structure). Some properties of the selected absorption systems are presented in Table 1. For convenience, we will often refer to the absorbers by the shorthand name in the first column of Table 1, which is just the rough redshift of the absorption system plus ‘A’ or ‘B’ to refer to individual components at that redshift. The large number of well-detected lines in these systems (see below) enables us to test the software in detail and to demonstrate some initial results. The high quality of the data, including high resolution detections of Mg II and Fe II with the ground-based Keck HIRES spectrograph, allow us to not only model for the physical parameters outlined above, but to do so separately for individual velocity components within each absorption system. Before proceeding, we emphasize our adopted parlance by which the terms *component* and/or *absorber* denote a group of kinematically aligned absorption lines, while *phases* are spectroscopically unresolved partitions of a given velocity component/absorber corresponding to distinct sets of physical conditions.

2.1 Observations

The COS FUV spectra covering $1152\text{--}1800 \text{ \AA}$ were taken with the G130M and G160M gratings, which provide a spectral resolution of $\approx 15\text{--}20 \text{ km s}^{-1}$ per resolution element (Fis-

¹ The species shown in Figure 1 have many more resonance lines than indicated, but the lines that are not shown are too weak to be detected at the typical signal-to-noise (S/N) of CASBaH data or are outside the wavelength range of the spectra.

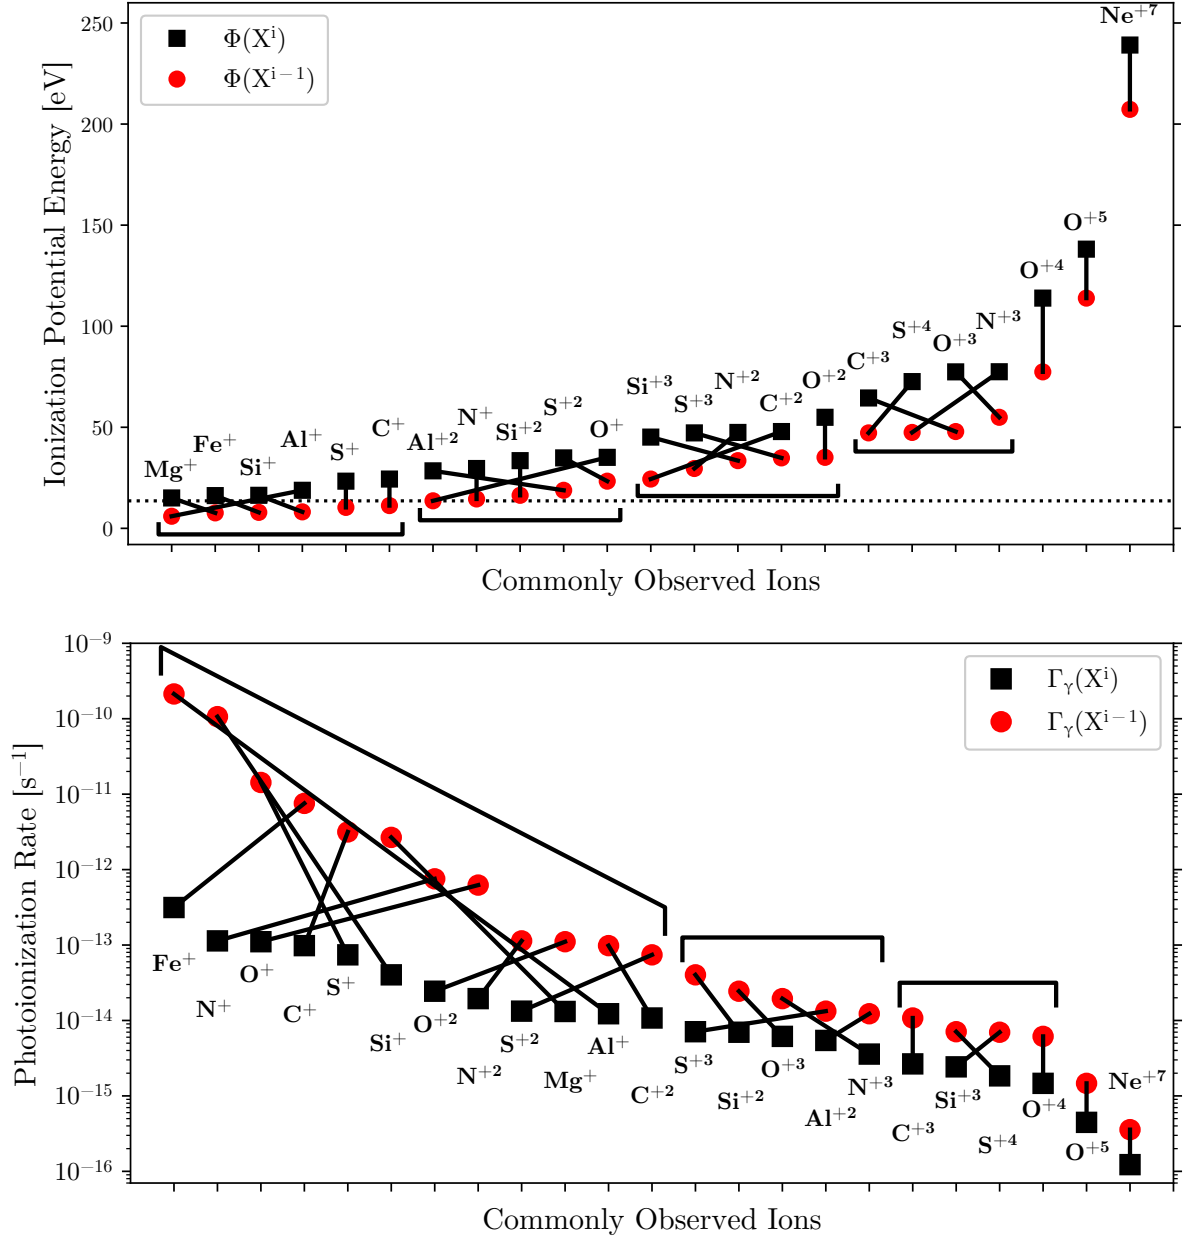


Figure 2. *Top:* A comparison of ionization potentials (IP) of ions that are commonly detected in CASBaH spectra. For each element, the black square indicates the IP of the labeled ion (X^i), which is connected with a black line to a red circle that shows the IP of the next lower stage (X^{i-1}). Thus, the red dot indicates the energy required to create the ion, and the black square plots the energy needed to destroy it. The points are plotted in order of increasing ionization energy. Horizontal black brackets group species with overlapping ionization creation and destruction energies (overlapping slanting black lines). Isolated species (i.e., S⁺, C⁺, and O⁺² in the top panel and C⁺³ and O⁺⁴ in the bottom panel) were assigned to the bracketed groups with the most similar ionization potentials. Heuristically these categorize groups of ions that might be expected to coexist in a specific gas phase. The horizontal dotted line at 1 Rydberg shows the ionization potential of H I. *Bottom:* Photoionization rates (Γ) of the same ions, calculated for an optically thin gas cloud that is ionized by the UV background from [Haardt & Madau \(2012\)](#) at redshift $z=0.5$, with photoionization cross sections from [Verner et al. \(1996a\)](#). As in the top panel, red circles and black squares indicate X^{i-1} and X^i , respectively, and the labeled ions and their adjacent lower ionization stages are connected with a black line. Here the points are plotted in order of decreasing Γ . Note the very low photoionization rates expected for the highest ionization stages.

[cher et al. 2019](#)). The data were recorded in the time-tagged photon-counting mode, and two grating tilts were used with both gratings to fill in the chip gap in the COS FUV detector (see Table 1.2 in [Rafelski 2018](#)). In addition, multiple exposures at three or four focal-plane split positions were

used at each grating tilt, so the flux at any given wavelength was recorded at six to eight different locations on the detector. This greatly mitigates detector fixed-pattern noise when the exposures are aligned in wavelength space and coadded. Initial data reduction steps were completed using the CAL-

Table 1. Characteristics of the three absorption systems (five total absorbers) under consideration here.

Absorber ^a	QSO Sightline	z_{abs}	v^b [km s ⁻¹]	$\log N_{\text{H I}}^c$ [cm ⁻²]	Detected Ions ^d
0.4A	PG1630+377	0.41760	0	15.46 ± 0.02	H I (full Lyman series), C II/III, N II/III, O II/III/VI, Mg II, Si II/III, Fe II
0.4B	PG1630+377	0.41760	44	15.33 ± 0.01	H I (full Lyman series), C II/III, N III, O II/III/VI, Mg II, Si III
0.5A	PG1522+101	0.51850	3	16.22 ± 0.01	H I (full Lyman series), C II/III, N III, Si II/III O II/III, Mg II, Fe II
0.6A	PG1338+416	0.68606	0	16.52 ± 0.02	H I (full Lyman series), C II, N II/III/IV, O II/III/VI, Mg I/II, Ne VIII, Al III, Si II, S II/III/V, Fe II
0.6B	PG1338+416	0.68606	46	15.90 ± 0.06	H I (full Lyman series), C II/III, N III/IV, O II/III/IV/VI, Mg II, Ne VIII, S III/V/VI

^a Throughout this work, we reserve the term “absorber” to mean a distinct *velocity component* of a given *absorption system*. Refer to Section 3.2 for more details.

^b The velocity of H I (from Voigt profile fitting) relative to the systemic redshift of the absorption system.

^c The column density uncertainties are derived from Voigt profile fits to numerous, high S/N H I Lyman Series lines simultaneously, hence the high precision.

^d These lists include lower limits, but exclude upper limits, which are presented in Tables 2–4.

COS pipeline (version 3.1.7) to produce extractions of the one-dimensional spectra (see Rafelski 2018 for details). The exposures were then aligned using well-detected lines of comparable strength from the ISM or extragalactic systems (e.g., higher H I Lyman series lines) and coadded as described in Meiring et al. (2011) and Tripp et al. (2021, *in prep.*). Finally, the combined spectra were binned to the Nyquist sampling rate of two pixels per resolution element (by default, the COS pipeline delivers 6 pixels per resolution element). The coadded and binned FUV spectra of the the three CASBaH QSOs in this paper have signal-to-noise ratios (per resolution element) ranging from 15 – 43 in the G130M range and 12 – 28 in the G160M range.²

To obtain NUV spectra of the targets, we used either the NUV mode of COS or the E230M echelle mode of STIS, depending on the target wavelength range of the spectrum. At $\lambda_{\text{ob}} \gtrsim 2200$ Å, we generally preferred to use STIS because it provides somewhat higher spectral resolution and a better, more Gaussian line-spread function; the COS NUV modes provide a spectral resolution of ≈ 15 – 20 km s⁻¹ while the STIS E230M resolution is ≈ 10 km s⁻¹. In the 1800–2200 Å range the COS NUV mode is much more sensitive and efficient than the corresponding STIS configuration, but this COS mode has two large gaps in the covered wavelengths at any given grating tilt (see Figure 13.17 in Fischer et al. 2019). To obtain complete COS NUV spectra without gaps, we used the G185M and G225M gratings with several grating tilts to fill in the gaps (see Tripp et al. 2021, *in prep.*). The COS NUV data were reduced in a manner similar to the FUV data but with an updated version of CALCOS (version 3.1.8) that was needed to correct a wavelength calibration problem that we encountered when we initially reduced the NUV data. We also used a 9-pixel extraction box rather than the default 57-pixel box. The smaller extraction box leads to a small loss of flux owing to the broad wings of the COS line-spread function but it significantly improves the S/N of the extracted spectrum.

The STIS data were reduced using the method of Tripp et al. (2001), including a weighted coaddition of the overlapping regions of adjacent orders as well as a weighted coaddition of the individual exposures. The NUV spectra were obtained to detect stronger lines such as H I *Ly α* and O VI doublets redshifted into the NUV range. Consequently, the NUV spectra were not required to attain as high S/N as the FUV observations. The final NUV spectra used here have S/N per resolution element ranging from 4 to 19.

As noted above, we also obtained high-resolution optical spectroscopy with the HIRES instrument on the Keck 1 telescope. Most of the CASBaH targets were observed with the HIRES C1 decker, which provides a spectral resolution of 6 km s⁻¹, and were recorded with the 3-chip mosaic CCD detector.³ The data were reduced and continuum normalized using the procedures described by O’Meara et al. (2015). Additional information about these observations can be found in Tripp et al. (2021, *in prep.*).

Figure 3 shows an example of the CASBaH FUV COS spectra. This small portion of the PG1630+377 spectrum demonstrates some of the advantages and challenges of high-S/N and high-resolution observations of $z \approx 1$ quasars. The absorption feature in Figure 3 is a complicated blend of lines from various species at various redshifts, and one might reasonably question whether this blend can be usefully decomposed. However, since there are many resonance transitions available in the far-UV (see Figure 1), we find that we can obtain meaningful constraints by exploiting all of the information in the data and the broad wavelength coverage of the survey; the velocity centroids and line widths of blended features are often constrained by other transitions of the same species recorded elsewhere in the spectra. Indeed, this is evident in Figure 3: much of the absorption in this example is caused by O II and O III ions in the absorption system at $z_{\text{abs}} = 0.41760$ (see also Figure 4), and by comparing the optical

² These S/N ratios apply to the majority of the FUV spectra; the S/N ratios drop in small wavelength ranges at $\lambda_{\text{ob}} < 1160$ and $\lambda_{\text{ob}} > 1750$ Å.

³ Two of the targets in the broader CASBaH program were observed with the C5 decker, and those spectra have a somewhat lower resolution of 8 km s⁻¹. However, all of the data presented in this paper employed the higher resolution C1 aperture.




Figure 3. A small portion of the far-ultraviolet COS spectrum of PG1630+377 (black histogram) obtained as part of the CASBaH survey. This example shows the CASBaH data quality and complexity and illustrates some of the advantages and challenges presented by the data, as discussed in the text. In this wavelength range, absorption lines from the O II multiplet at 832.76, 833.33, and 834.47 Å as well as the O III 832.93 line at $z_{\text{abs}} = 0.41760$ are detected but are blended with each other and with absorption lines from other absorption systems at different redshifts. To extract absorption-line measurements from these data, we have fit Voigt profiles to the entire blend, indicated with a thin purple line. In these O II and O III transitions, we find evidence of three components at $v = 0, 43$ and 197 km s^{-1} in the $z_{\text{abs}} = 0.41760$ system. These absorption components from the O II multiplet are shown with a thick red line, and the O III absorption components are plotted with a dashed blue line. The O II and O III lines of interest are blended with H I Lyman series lines at $z_{\text{abs}} = 0.2153, 0.2741$, and 0.2780 as well as multicomponent Ne VIII 770.409 Å absorption at $z_{\text{abs}} = 0.5334$. The identification of the H I lines, and their detailed contributions to this blend, are highly secure because they are constrained by many other H I Lyman series lines at the same redshift. Likewise, the Ne VIII 770.409 features are corroborated by their kinematical alignment with O VI and other metals at the same redshift.

depths at the expected wavelengths of the various transitions, one finds that the O II and O III features are robustly identified and well constrained.

2.2 Column Density Measurements

We identified absorption metal lines associated with the prominent H I Lyman Series lines of the systems in Table 1 and then measured their column densities and limits. Figures 4–6 show the absorption profiles in these systems, and the measured line parameters are listed in Tables 2–4. The unusual quality and richness of the CASBaH QSO spectra enhances our ability to detect absorption and useful non-detections with lower uncertainties for a larger number of species compared with average quality QSO spectra (e.g., $z_{\text{QSO}} \approx 0.2$, $S/N \lesssim 10$). For many ions, the uncertainty is reduced by the presence of multiple resonance transitions (transitions that are unavailable at lower redshift) as well as by the decrease in shot noise. The improved confidence in our measurements stems from a) an enhanced ability to identify and constrain systematics owing to intervening absorption lines from other redshifts, b) careful decomposition of absorbers' underlying velocity structures, and c) co-fitting multiple components with shared parameters. Below we relate our insights related to the absorbers in this study and to the analysis of high-quality, line-rich QSO spectra in general.

2.2.1 Line Identification and Deblending

Prior to making line measurements, we identified all absorption systems in our spectra based on their H I Lyman Series lines and/or corroborating metal line transitions at the same redshift (Tripp et al. 2021 *in prep.*). A full census of lines in each sightline is essential for identifying and deblending contaminating lines, i.e., absorption features from various redshifts that coincidentally overlap in observed wavelength with an absorption line of interest. Potentially unidentified contamination, whether from weak Ly α lines ($N_{\text{H I}} \lesssim 10^{13.5} \text{ cm}^{-2}$) or metal lines associated with an obscured or weak H I line (see, e.g., Figure 7 in Tripp et al. 2008 and Figures 4 and 7 in Meiring et al. 2013) could lead to an overestimation of the contaminated line's absorption strength. This potential for contamination increases with the redshift of the observed QSO, owing to the steadily increasing density of UV transitions that are redshifted to observable wavelengths as well as the increase in path length along the line-of-sight.

Additional checks for hidden contamination can be performed for species with an observed coverage of doublet and multiplet transitions by comparing their apparent column density (ACD) profiles (Savage & Sembach 1991; Sembach & Savage 1992), which should match in the absence of contamination. For species with only one observable transition (e.g., C III $\lambda 1977.02 \text{ Å}$, N IV $\lambda 765.15 \text{ Å}$, Si III $\lambda 1206.05 \text{ Å}$, and S V $\lambda 786.48 \text{ Å}$), a comparison of velocity profiles (absorption or ACD) with those of the full set of detected species will often help to confirm a detection by virtue of their similarity, or it may reveal some contaminating absorption line as a




Figure 4. Absorption velocity profiles and best-fit Voigt Profiles for detections and selected non-detections at $z=0.41760$ in the PG1630+377 sightline (includes Absorbers 0.4A and 0.4B at velocities ≈ 0 and ≈ 43 km s $^{-1}$, respectively). These draw from observations using multiple instruments with significant differences in S/N and resolution between them. The flux has been continuum-normalized as described in the text. Dashed red (yellow) lines illustrate individual Voigt profiles for the labeled species (unlabeled contamination), with solid red (yellow) lines indicating the total modeled absorption. The Voigt profile fits to contaminating lines from absorbers at other redshifts (yellow curves) are all reasonably well constrained through co-fitting with another transition (at the contaminating absorber’s redshift) of the same or similar species observed in a separate region of the spectrum.

spurious bulge in the velocity profiles. Two examples of this are the high velocity O III and S III lines at $v \approx 200$ km s $^{-1}$ in the $z=0.41760$ absorber from the PG1630+377 sightline. Figure 4 shows absorption profiles in red for species at this redshift, including Si III, as well as contaminating lines in yellow. We can confirm Si III absorption at $v \approx 200$ km s $^{-1}$ by noting its similarity to the C III and H I absorption features present at that velocity, and we can simultaneously conclude that a strong Si III feature at $v \approx 145$ km s $^{-1}$ is unlikely because it is not corroborated in C III, nor in H I. O III, on the other hand, is apparent in the flux decrement remaining after accounting for all the adjacent, well-constrained contaminating lines. In the same figure, the O VI doublet profiles provide an example of contamination (centered at $v \approx 30$ km s $^{-1}$ in the O VI $\lambda 1037\text{\AA}$ profile) that becomes apparent when comparing multiplet species. These examples underscore the complexity of the data at moderate redshifts and the importance of high quality spectra, and extensive coverage of the Ly α forest region, for de-blending contamination. Spectra with insufficient resolution or S/N, or an inadequate wavelength range to cover corroborating lines, would not be nearly as effective

at constraining such large numbers of coincident absorption features.

2.2.2 Continuum Normalization

As a precursor to Voigt Profile fitting, we continuum-normalized our COS and STIS spectra, using continuum placements determined by fitting polynomials to arbitrary 10–50 \AA windows of continuum pixels (identified by eye) in the spectrum. We prefer this approach, which directly emphasizes human intuition, versus a more automated approach, which typically combine indirect human intuitions with additional poorly characterized algorithmic systematics. In general, we expect the continuum placement uncertainty (typically ≈ 0.01 dex) to have a minimal impact on the final absorption line column density measurement, based on a comparison of line measurements using continuum placements estimated independently by two of the authors (Haislmaier and Tripp). The precision of our continuum estimates is a consequence of the excellent spectral S/N, which significantly narrows the continuum placement ambiguity. The HIRES spectra were normalized using the methods of O’Meara et al. (2015).

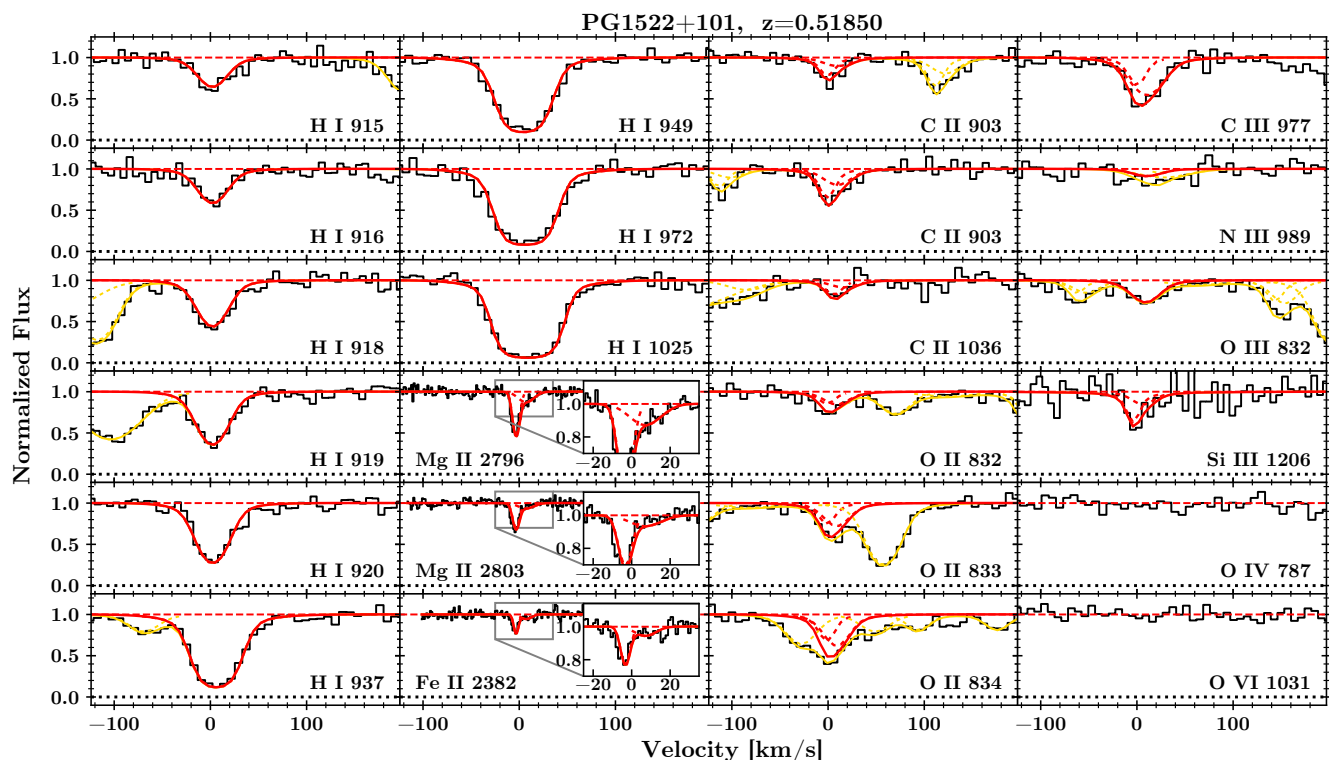


Figure 5. Absorption velocity profiles and best-fit Voigt Profiles for detections and selected non-detections at $z=0.51850$ in the PG1522+101 sightline (Absorber 0.5A). See Figure 4 for a description of the colored lines.

2.2.3 Voigt Profile Fitting

We measured the column density, linewidth, and velocity of all absorption lines associated with our selected H I absorbers by fitting Voigt Profiles (VP) to the relevant portions of the continuum-normalized spectra; the results are listed in Tables 2 - 4. In this paper, we chose the redshift of the strongest H I component in each of the three absorption systems as the systemic redshift. A visual comparison of ACD profiles suggests that all ionization stages are consistent with having a common velocity structure (i.e., the velocity separation between components), a point we discuss further below. We leverage this information to minimize VP uncertainties by co-fitting velocities for various groupings of ions.

Unfortunately, the absolute kinematic alignment of lines separated by as little as 1 \AA can be obscured by highly localized distortions of the wavelength scale caused by the type of detector used by COS (Dixon et al. 2007). For example, N III 989 \AA , O II 834 \AA , and O III 832 \AA in Figure 4 all have the same two prominent components separated by $\approx 43 \text{ km s}^{-1}$, but the velocities of the two O II components are offset by $\approx 6 \text{ km s}^{-1}$ from the O III components, while N III exhibits a $\approx 4 \text{ km s}^{-1}$ shift. A similar effect also plagues the C IV and O VI in the same figure. While in these examples the offsets could be real, similar offsets are found for different lines of the same species that should have exactly the same centroid. These distortions are typically largest near the detector edges and the COS detector gap. To model this source of error, we add a free parameter (labeled as v_{corr} in Tables 2 - 4) for each $\approx 1 \text{ \AA}$ segment of each spectrum that, in effect, adjusts the wavelength calibration of each segment independently. v_{corr} is mathematically redundant with the VP velocity parameter,

so we only need it when the VP component velocities, v , are shared or “co-fit” between lines in different segments of the spectrum. While this issue with the COS wavelength scale is somewhat annoying (and possibly impossible to fully rectify), we underscore that these calibration errors are small enough so that they do not profoundly affect the conclusions in this paper.

The fits for C II and O II at $z=0.68606$ in PG1338+416 are an excellent example of how co-fitting can improve the measurement. Fitted individually, the stronger C II VP converges to an unrealistic column density of $N_{\text{C II}}=10^{17.62} \text{ cm}^{-2}$ and a spurious broadening parameter value of $b=0.45 \text{ km s}^{-1}$. This poor fit primarily results from poorly constrained velocity centroids for C II. A similar problem occurs with O II. Meanwhile, the $v \approx 0$ and $v \approx 43 \text{ km s}^{-1}$ Mg II doublet components in this absorption system can be precisely measured in the high resolution HIRES spectrum. Since all three species are low-ions and follow similar absorption profiles, i.e., strong absorption at $v \approx 0 \text{ km s}^{-1}$ relative to the weaker $v \approx 43 \text{ km s}^{-1}$ component, they likely arise from similar physical processes and thus ought to make good candidates for co-fitting.⁴ With this technique, we manage to recover useful constraints on the C II and O II column densities (reported in Table 4).

In some cases it is also advantageous to fit two species with shared broadening parameters in addition to shared veloci-

⁴ Note that tying velocity centroids together does not necessarily require the respective species to be associated with the same phase (see discussion in Section 3.1). The phase structure present in a given velocity component can only be inferred by ionization modeling, such as we perform below.




Figure 6. Absorption velocity profiles and best-fit Voigt Profiles (VPs) for detections and selected non-detections at $z=0.68606$ in the PG1338+416 sightline (includes Absorbers 0.6A and 0.6B at velocities ≈ 0 and ≈ 43 km s^{-1} , respectively). Ne VIII 780Å is located in an unconstrained blend and could not be fitted, but is consistent with the fit to Ne VIII 770Å. The red dashed-line VPs shown for Ne VIII 780Å are drawn based on best fit parameters from Ne VIII 770Å for comparison. The velocities and b -values for VP fits to saturated lines were constrained through co-fitting with non-saturated lines following the co-fitting considerations described in the text (note that the column densities listed in Table 4 include those derived from the fits shown here as well as those for conservative upper and lower limits). See Figure 4 for a description of the colored lines.

ties. This must be done with caution, however, since several factors influence the line-widths. First, ions produced in different temperature gas phases (e.g., low and high-ions in a multiphase medium) will have distinct thermal broadening solutions. Additionally, the degree of thermal broadening varies with atomic weight, and this should be considered when comparing different species. We refer to this as “thermally-scaled” broadening. Second, we must account for the possibility of non-thermal broadening that has the potential to smooth out linewidth differences between different-mass species, what we call “evenly-scaled” broadening. Finally, weak, narrow, and unresolved lines can appear to

broaden lines after smearing by the line spread function⁵. It is reasonable to question when it is possible to share broadening parameters in light of these considerations. In practice, however, as long as one restricts shared broadening to species with similar ionization potentials (i.e., expected to be in the same gas “phase”), broadening parameters from the co-fitted VPs are predominantly consistent with the results of independent VP fits (but typically with improved uncertainties), and the difference has an even smaller impact on the column density. The exception to this statement is in the presence of narrow, saturated lines (see the discussion of limits below for an

⁵ We use the line spread functions provided by the Space Telescope Science Institute for COS (Kriss 2011) and a 2 pixel Gaussian FWHM for HIRES spectra.

example from the PG1630+377, $v=0.41760$ system). When co-fitting broadening parameters, we take care to compare fits based on “thermally-scaled” and “evenly-scaled” broadening. The bracketed numerical superscripts in Tables 2 - 4 mark line groups that have a shared parameter. The type of shared parameter varies from case-to-case; notes about these measurements are provided in Sections 2.2.4 and 2.3.

2.2.4 Upper and Lower Limits

Limits were measured for saturated lines and non-detections in various ways. We adopted 3σ upper limits for lines that did not meet a 3σ equivalent width (EW) threshold for detection. In addition, occasionally upper limits were measured from VP fitting for absorption features suspected to be affected by an unidentified blend (e.g., S VI at $z=0.68606$, $v \approx 43 \text{ km s}^{-1}$ in PG1338+416, Figure 6). Lower limits to saturated lines were estimated from their equivalent widths. In some cases intrinsically narrow, saturated lines can deceptively appear to be optically thin (e.g., the line core is well above the zero-flux level) as a consequence of the COS line spread function’s broad wings (Kriss 2011). For example, the C III for the $z=0.41760$ system in PG1630+377 is well-fit by two apparently optically thin VPs. Our ionization modeling (see below) persistently overpredicts the C III column density in both components, however, so we explored its absorption more carefully and found that the same lines were well-fit with a separate set of VP parameters having optically thick column densities consistent with those predicted by the ionization model. This alternative fit for C III requires significantly narrower (i.e., smaller broadening parameter) lines than the original fit, however. To ensure that we adopt a self-consistent set of VP parameters, we re-fit the N III, O III, and Si III lines along with the C III, choosing to co-fit their velocities and thermally-scaled broadening parameters for various fixed choices of broadening parameter since all four ions have similar ionization potentials. In fact, we found in this case that all four ions are consistent with having saturated (or nearly saturated), narrow VPs, and that a double-bounded limit is more accurate to report than a single best-fit. We report lower limits estimated from the naive fit, i.e., assuming the lines are optically thin, and upper limits from a fit where the broadening parameter was fixed to match the low-ion broadening.

2.3 Notes for Individual Absorption Systems

2.3.1 $z=0.41760$, PG1630+377 (Absorbers 0.4A and 0.4B)

As shown in Figure 4, in this absorption system we identify four components (Components A–D in Table 2), including two line-rich components with relatively high H I column densities at $v=0$ and 43 km s^{-1} as well as two weaker components with lower $N_{\text{H I}}$ values at $v=110$ and 197 km s^{-1} (see Table 2). Interestingly, the three components with low/mid-ion absorption (i.e., all except the one at $v \approx 110 \text{ km s}^{-1}$) exhibit a gradient in ionic column densities and in velocity: the degree of ionization in each velocity component appears to be correlated with increasing velocity. The strongest H I, C II, O II, and Mg II absorption appears in the lowest velocity component, while C III, C IV, O III, and O VI are stronger in the higher-velocity features. This ionization gradient is most

Table 2. Column Density Measurements for the $z=0.41760$ Absorption System in PG1630+377 (including Absorbers 0.4A and 0.4B)

Ion	λ_0^a (Å)	$\log N(X)^{b,c}$ (cm^{-2})	$b^{c,d}$ (km s^{-2})	$v^{d,e}$ (km s^{-2})	$v_{\text{corr}}^{d,e}$ (km s^{-2})
Component A ($v \approx 0 \text{ km s}^{-1}$)					
HI Ly $\alpha - \epsilon, \eta$		15.46 ± 0.016	13.6 ± 0.4	0.0 $\pm 1.0^{[1]}$	-1.0 $\pm 1.3^{[1]}$
CII 903, 903		13.46 ± 0.059	5.6 $\pm 0.7^{[1]}$	0.0 $\pm 1.0^{[1]}$	-0.2 $\pm 1.1^{[2]}$
CIII 977		$\geq 13.52 \pm 0.029$	11.7 $\pm 0.2^{[2]}$	0.4 $\pm 0.5^{[2]}$	-0.4 $\pm 0.6^{[3]}$
CIII 977 ^g		$< 14.20 \pm 0.000$	7.5 $\pm 0.2^{[2]}$		
CIV 1548, 1550		13.48 ± 0.175	11.7 $\pm 0.2^{[2]}$	0.4 $\pm 0.5^{[2]}$	1.3 $\pm 3.0^{[4]}$
NII 915		13.05 ± 0.104	5.2 $\pm 0.7^{[1]}$	0.0 $\pm 1.0^{[1]}$	-4.4 $\pm 2.8^{[5]}$
NIII 989		13.62 ± 0.034	10.8 $\pm 0.2^{[2]}$	0.4 $\pm 0.5^{[2]}$	3.4 $\pm 1.1^{[6]}$
OII 832, 833, 834		13.70 ± 0.051	4.8 $\pm 0.7^{[1]}$	0.0 $\pm 1.0^{[1]}$	-3.8 ± 0.0
OIII 832		$\geq 13.99 \pm 0.031$	10.1 ± 0.6	0.4 ± 0.6	-3.8 ± 0.0
OIII 832 ^g		$< 14.11 \pm 0.045$	6.5 $\pm 0.2^{[2]}$		
OVI 1031, 1037		13.17 ± 0.156	22.8 ± 11.3	0.4 $\pm 0.5^{[2]}$	-1.5 $\pm 3.0^{[7]}$
MgII 2796, 2803		12.17 ± 0.013	4.0 $\pm 0.2^{[3]}$	0.0 $\pm 1.0^{[1]}$	-1.3 $\pm 1.1^{[8]}$
SiIII 1206		$\geq 12.79 \pm 0.077$	7.6 $\pm 0.2^{[2]}$	0.4 $\pm 0.5^{[2]}$	-5.5 $\pm 1.3^{[9]}$
SiIII 1206 ^g		$< 13.05 \pm 0.121$	4.9 $\pm 0.2^{[2]}$		
SVI 933 ^f		$< 12.43 \pm 0.117$			
FeII 2344, 2382, 2600		11.80 ± 0.042	3.5 ± 0.5	0.0 $\pm 1.0^{[1]}$	-1.3 $\pm 1.1^{[8]}$
Component B ($v \approx 44 \text{ km s}^{-1}$)					
HI Ly $\alpha - \epsilon, \eta$		15.33 ± 0.012	17.4 ± 0.5	43.7 $\pm 1.0^{[3]}$	-1.0 $\pm 1.3^{[1]}$
CII 903, 903		13.16 ± 0.031	14.6 $\pm 1.7^{[4]}$	43.7 $\pm 1.0^{[3]}$	-0.2 $\pm 1.1^{[2]}$
CIII 977		$\geq 13.83 \pm 0.035$	15.4 $\pm 0.2^{[5]}$	46.5 $\pm 0.4^{[4]}$	-0.4 $\pm 0.6^{[3]}$
CIII 977 ^g		$< 14.40 \pm 0.000$	11.0 $\pm 0.2^{[5]}$		
CIV 1548, 1550		14.19 ± 0.136	17.5 ± 4.2	46.5 $\pm 0.4^{[4]}$	1.3 $\pm 3.0^{[4]}$
NII 915		11.82 ± 1.872	13.5 $\pm 1.7^{[4]}$	43.7 $\pm 1.0^{[3]}$	-4.4 $\pm 2.8^{[5]}$
NIII 989		$\leq 13.53 \pm 0.038$	14.3 $\pm 0.2^{[5]}$	46.5 $\pm 0.4^{[4]}$	3.4 $\pm 1.1^{[6]}$
NIII 989 ^g		$> 13.50 \pm 0.039$	10.2 $\pm 0.2^{[5]}$		
NV 1238 ^f		$< 12.99 \pm 0.204$			
OII 832, 833, 834		13.48 ± 0.048	12.6 $\pm 1.7^{[4]}$	43.7 $\pm 1.0^{[3]}$	-3.8 ± 0.0
OIII 832		$\geq 14.35 \pm 0.023$	13.4 ± 0.6	46.5 ± 0.5	-3.8 ± 0.0
OIII 832 ^g		$< 14.52 \pm 0.040$	9.5 $\pm 0.2^{[5]}$		
OVI 1031, 1037		13.58 ± 0.056	18.6 ± 3.6	46.5 $\pm 0.4^{[4]}$	-1.5 $\pm 3.0^{[7]}$
MgII 2796, 2803		11.53 ± 0.039	4.7 $\pm 0.6^{[6]}$	43.7 $\pm 1.0^{[3]}$	-1.3 $\pm 1.1^{[8]}$
SiIII 1206		12.56 ± 0.295	10.1 $\pm 0.2^{[5]}$	46.5 $\pm 0.4^{[4]}$	-2.1 $\pm 1.3^{[9]}$
SiIV 1402 ^f		$< 13.02 \pm 0.200$			
SVI 944 ^f		$< 12.56 \pm 0.180$			
FeII 2344, 2382, 2600		10.91 ± 0.238	3.1 $\pm 0.6^{[6]}$	43.7 $\pm 1.0^{[3]}$	-1.3 $\pm 1.1^{[8]}$
Components A + B					
NV 1238 ^f		< 13.35			
CIV 1550 ^f		$\geq 14.12 \pm 0.052$			
OI 988 ^f		< 13.56			
OVI 1031 ^f		$\geq 13.66 \pm 0.033$			
MgI 2852 ^f		< 11.03			
AlIII 1670 ^f		< 12.22			
AlIII 1854 ^f		< 12.76			
SiII 1193 ^f		$< 12.78 \pm 0.070$			
SiIV 1402 ^f		< 13.49			
Component C ($v \approx 109 \text{ km s}^{-1}$)					
HI Ly $\alpha - \epsilon$		13.51 ± 0.076	45.4 $\pm 10.6^{[7]}$	109.4 $\pm 3.3^{[5]}$	-1.7 ± 0.0
OVI 1031, 1037		13.09 ± 0.094	11.4 $\pm 10.6^{[7]}$	109.4 $\pm 3.3^{[5]}$	-1.7 ± 0.0
Component D ($v \approx 199 \text{ km s}^{-1}$)					
HI Ly $\alpha - \epsilon$		14.64 ± 0.011	17.4 ± 0.5	198.5 $\pm 0.9^{[6]}$	-1.0 $\pm 1.3^{[1]}$
CIII 977		12.87 ± 0.029	12.9 $\pm 1.4^{[8]}$	198.5 $\pm 0.9^{[6]}$	-0.2 $\pm 0.6^{[3]}$
OIII 832		13.61 ± 0.071	11.2 $\pm 1.4^{[8]}$	198.5 $\pm 0.9^{[6]}$	-3.8 ± 0.0
SiIII 1206		12.12 ± 0.142	8.5 $\pm 1.4^{[8]}$	198.5 $\pm 0.9^{[6]}$	-2.1 $\pm 1.3^{[9]}$

^a Only atomic transitions used to produce the measurements are listed. Some transitions covered by the observations had to be excluded as they did not provide meaningful constraints to the fit, e.g., due to severe blending or low signal-to-noise.

^b Limits preceded by partial inequality signs (\leq or \geq) indicate solid detections affected by significant systematics (e.g., saturation or complex covariances) in one direction. Non-detections with a single numeric (i.e., no extra uncertainty term) were derived from 3σ equivalent width upper limits.

^c The Voigt Profile (VP) broadening parameter, or b -value.

^d All parameters sharing the same bracketed superscript number in a given column derive their value from a single free parameter, i.e., that parameter was shared by several different Voigt profiles.

^e Section 2.2.3 describes the relation between “ v ”, the component velocity, and “ v_{corr} ”, a wavelength dependent velocity correction. The systemic redshift is listed in the table header.

^f These measurements were derived from the equivalent width.

^g See the text regarding the estimation of these upper limits.

Table 3. Column Density Measurements for the $z=0.51850$ Absorption System in PG1522+101 (Absorber 0.5A)

Ion	λ_0^a (Å)	$\log N(X^i)^b$ (cm ⁻²)	$b^{c,d}$ (km s ⁻²)	$v^{d,e}$ (km s ⁻²)	$v_{corr}^{d,e}$ (km s ⁻²)
Component A1 ^g ($v \approx -5$ km s ⁻¹)					
CII 903.6, 903.9, 1036	13.22 ± 0.121	6.4 ± 0.3 ^[1]	-4.6 ± 0.2 ^[1]	4.0 ± 1.6 ^[2]	
CIII 977	12.85 ± 0.289	6.4 ± 0.3 ^[1]	-4.6 ± 0.2 ^[1]	4.0 ± 2.2 ^[1]	
OII 832, 833, 834	13.67 ± 0.166	5.5 ± 0.3 ^[1]	-4.6 ± 0.2 ^[1]	3.2 ± 2.0 ^[3]	
MgII 2796, 2803	12.13 ± 0.036	4.5 ± 0.3 ^[1]	-4.6 ± 0.2 ^[1]	0.0 ± 0.0	
SiII 1193, 989	12.38 ± 0.443	4.2 ± 0.3 ^[1]	-4.6 ± 0.2 ^[1]	4.0 ± 2.2 ^[1]	
SiIII 1206	12.65 ± 0.648	4.2 ± 0.3 ^[1]	-4.6 ± 0.2 ^[1]	1.4 ± 4.9 ^[4]	
FeII 2382	11.89 ± 0.051	3.0 ± 0.3 ^[1]	-4.6 ± 0.2 ^[1]	0.0 ± 0.0	
Component A2 ^g ($v \approx +6$ km s ⁻¹)					
CII 903.6, 903.9, 1036	13.04 ± 0.157	16.0 ± 1.8 ^[2]	5.8 ± 1.7 ^[2]	4.0 ± 1.6 ^[2]	
CIII 977	13.18 ± 0.094	18.7 ± 2.1 ^[3]	5.8 ± 1.7 ^[2]	4.0 ± 2.2 ^[1]	
NIII 989	13.09 ± 0.220	17.3 ± 2.1 ^[3]	5.8 ± 1.7 ^[2]	4.0 ± 2.2 ^[1]	
OII 832, 833, 834	13.82 ± 0.096	13.8 ± 1.8 ^[2]	5.8 ± 1.7 ^[2]	3.2 ± 2.0 ^[3]	
OIII 832	13.76 ± 0.049	16.2 ± 2.1 ^[3]	5.8 ± 1.7 ^[2]	3.2 ± 2.0 ^[3]	
MgII 2796, 2803	11.79 ± 0.092	11.2 ± 1.8 ^[2]	5.8 ± 1.7 ^[2]	0.0 ± 0.0	
SiIII 1206	12.11 ± 0.713	12.2 ± 2.1 ^[3]	5.8 ± 1.7 ^[2]	1.4 ± 4.9 ^[4]	
FeII 2382	11.59 ± 0.122	7.4 ± 1.8 ^[2]	5.8 ± 1.7 ^[2]	0.0 ± 0.0	
Component A = A1 + A2 ^g					
HI Lyβ - ε, θ - ν	16.22 ± 0.011	16.9 ± 0.2	3.2 ± 0.4	0.0 ± 0.0	
CII 903.6 ^f	≥ 13.34 ± 0.042				
CIII 977 ^f	≥ 13.23 ± 0.025				
CIV 1548 ^f	< 13.62				
NII 915 ^f	< 13.29				
NIV 765 ^f	≤ 12.97 ± 0.092				
OI 988 ^f	< 13.83				
OIV 787 ^f	< 13.56				
OVI 1031 ^f	< 13.40				
SiIII 1206 ^f	≥ 12.42 ± 0.133				
SiIV 1393 ^f	< 13.21				
SII 765 ^f	≤ 12.41 ± 0.181				
SV 786 ^f	< 12.48				

^{a-f} Tablenotes a-f are given in Table 2^g Components A1 and A2 represent our best attempt to resolve the absorption features into two components. As discussed in the text, we use the total column densities in component A (i.e., A1 + A2) as constraints on our ionization models.

evident in the H I lines; the H I column densities clearly decrease with increasing velocity. This could simply be a consequence of less gas in the higher-velocity clouds, but the fact that the column densities of more highly ionized metals do not decrease as well suggests that this owes to the gas becoming increasingly more ionized with increasing velocity.

It is also interesting to note that Component C at $v = 110$ km s⁻¹ is only detected in O VI and H I. In their blind STIS survey, [Tripp et al. \(2008\)](#) report many absorbers only detected in O VI and H I (no C III), and similarly [Werk et al. \(2016\)](#) found a category of absorbers (their “No-low” sample) with a similar absence of low-ionization stages, yet having O VI aligned with moderate N(H I). We measure a narrow linewidth for O VI in Component C, and our best fit to the H I feature at this velocity returns a doppler broadening parameter $b(\text{H I}) \approx 4 \times b(\text{O VI})$. Both measurements are consistent with the O VI and H I arising in a single, thermally-broadened gas phase. While the best-fit line widths for this component imply a temperature slightly below the T range where O VI is expected to peak in collisional ionization equilibrium, the line width uncertainties are substantial, and detailed analysis of this feature would require a higher spectral resolution and/or better S/N.

In Component D at $v \approx 200$ km s⁻¹, there is enough degeneracy in the fit to the spectrum to allow for some weak O VI absorption, but this component is also highly uncertain ow-

Table 4. Column Density Measurements for the $z=0.68606$ Absorption System in PG1338+416 (including Absorbers 0.6A and 0.6B)

Ion	λ_0^a (Å)	$\log N(X^i)^b$ (cm ⁻²)	$b^{c,d}$ (km s ⁻²)	$v^{d,e}$ (km s ⁻²)	$v_{corr}^{d,e}$ (km s ⁻²)
Component A ($v \approx -83$ km s ⁻¹)					
HI Lyβ, ε	14.45 ± 0.048	40.4 ± 6.0	-82.9 ± 4.0	-0.3 ± 0.0 ^[1]	
NIV 765	12.71 ± 0.129	27.2 ± 9.8 ^[1]	-62.9 ± 6.1 ^[1]	-0.3 ± 0.0 ^[1]	
OIV 787	13.64 ± 0.146	27.2 ± 9.8 ^[1]	-62.9 ± 6.1 ^[1]	-4.1 ± 0.0	
OVI 1031, 1037	13.42 ± 0.201	27.2 ± 9.8 ^[1]	-62.9 ± 6.1 ^[1]	12.3 ± 0.0 ^[2]	
NeVIII 770 ^f	< 13.40 ± 0.230				
Component B = B1 ($v \approx -3$ km s ⁻¹) ∪ B2 ($v \approx 13$ km s ⁻¹)					
HI Lyβ, ε, ι, μ - π	16.52 ± 0.022	15.3 ± 0.9	0.0 ± 2.7	-0.3 ± 0.0	
CII 903.6, 903.9	≥ 13.75 ± 0.093	9.2 ± 1.3 ^[2]	-2.9 ± 0.6 ^[2]	7.7 ± 1.2 ^[3]	
CII 903.6, 903.9 ^g	< 14.04 ± 0.106	6.1 ± 0.0 ^[2]			
NII 915	13.38 ± 0.100	9.2 ± 1.3 ^[2]	-2.9 ± 0.6 ^[2]	2.0 ± 3.0 ^[4]	
NIII 763	13.71 ± 0.092	17.6 ± 2.4 ^[3]	12.5 ± 1.9 ^[3]	-0.3 ± 0.0 ^[1]	
NIV 765	13.13 ± 0.077	22.0 ± 3.2 ^[4]	12.5 ± 1.9 ^[3]	-0.3 ± 0.0 ^[1]	
OII 832, 833, 834	≥ 14.10 ± 0.056	9.2 ± 1.3 ^[2]	-2.9 ± 0.6 ^[2]	-2.1 ± 0.0	
OII 832, 833, 834 ^g	< 14.32 ± 0.081	6.1 ± 0.0 ^[2]			
OIII 702, 832	≥ 14.66 ± 0.071	17.6 ± 2.4 ^[3]	12.5 ± 1.9 ^[3]	-2.1 ± 0.0	
OIV 787	14.61 ± 0.084	22.0 ± 3.2 ^[4]	12.5 ± 1.9 ^[3]	-1.1 ± 2.7 ^[5]	
OVI 1031, 1037	14.19 ± 0.098	31.4 ± 5.4	12.5 ± 1.9 ^[3]	12.3 ± 0.0 ^[2]	
NeVIII 770 ^f	≤ 13.42 ± 0.221				
MgI 2852	11.01 ± 0.097	4.0 ± 1.1	-2.6 ± 0.8	2.3 ± 0.7 ^[6]	
MgII 2796, 2803	12.95 ± 0.019	5.7 ± 0.1	-2.9 ± 0.6 ^[2]	2.3 ± 0.7 ^[6]	
MgII 2796, 2803	11.77 ± 0.121	13.0 ± 4.0	12.5 ± 1.9 ^[3]	2.3 ± 0.7 ^[6]	
AlIII 1854, 1862	12.42 ± 0.071	8.9 ± 1.7	-2.9 ± 0.6 ^[2]	2.3 ± 0.7 ^[6]	
SII 763, 765	12.71 ± 0.084	5.1 ± 0.2 ^[5]	-2.9 ± 0.6 ^[2]	-0.3 ± 0.0 ^[1]	
SiII 724	13.38 ± 0.063	18.9 ± 5.1 ^[6]	12.5 ± 1.9 ^[3]	-6.4 ± 3.4 ^[7]	
SIV 748	12.78 ± 0.187	18.9 ± 5.1 ^[6]	12.5 ± 1.9 ^[3]	0.9 ± 2.5 ^[8]	
SV 786	12.44 ± 0.156	22.0 ± 3.2 ^[4]	12.5 ± 1.9 ^[3]	-1.1 ± 2.7 ^[5]	
SVI 944 ^f	< 13.00 ± 0.050				
FeII 2382, 2586, 2600	12.72 ± 0.014	5.1 ± 0.2 ^[5]	-2.9 ± 0.6 ^[2]	1.3 ± 0.7 ^[9]	
FeII 2382, 2586, 2600	≤ 11.52 ± 0.137	7.8 ± 0.0	12.5 ± 1.9 ^[3]	1.3 ± 0.7 ^[9]	
Component C = C1 ($v \approx -39$ km s ⁻¹) ∪ C2 ($v \approx 56$ km s ⁻¹)					
HI Lyβ, ε, ι, μ - π	15.90 ± 0.057	33.5 ± 3.0	46.4 ± 5.5	5.0 ± 0.0	
CII 903.6, 903.9	13.24 ± 0.080	19.0 ± 0.0 ^[7]	39.4 ± 0.8 ^[4]	13.0 ± 1.2 ^[3]	
NIII 763	13.77 ± 0.076	19.0 ± 2.0 ^[8]	55.8 ± 1.8 ^[5]	5.0 ± 0.0 ^[1]	
NIV 765	13.80 ± 0.057	19.0 ± 2.0 ^[8]	55.8 ± 1.8 ^[5]	5.0 ± 0.0 ^[1]	
OII 832, 833, 834	13.78 ± 0.047	19.0 ± 0.0 ^[7]	39.4 ± 0.8 ^[4]	3.2 ± 0.0	
OIII 702, 832	≥ 15.02 ± 0.126	19.0 ± 2.0 ^[8]	55.8 ± 1.8 ^[5]	3.2 ± 0.0	
OIV 787	≥ 15.26 ± 0.219	19.0 ± 2.0 ^[8]	55.8 ± 1.8 ^[5]	4.2 ± 2.7 ^[5]	
OVI 1031, 1037	14.68 ± 0.053	25.8 ± 2.8	55.8 ± 1.8 ^[5]	17.6 ± 0.0 ^[2]	
NeVIII 770	14.19 ± 0.062	78.8 ± 14.1	55.8 ± 1.8 ^[5]	5.0 ± 0.0 ^[1]	
MgII 2796, 2803	11.79 ± 0.046	5.9 ± 0.8	39.4 ± 0.8 ^[4]	7.6 ± 0.7 ^[6]	
AlIII 1854	< 12.34				
SII 765 ^f	< 13.42 ± 0.225				
SiII 724	12.96 ± 0.136	8.4 ± 3.3 ^[9]	55.8 ± 1.8 ^[5]	-1.1 ± 3.4 ^[7]	
SIV 748	13.27 ± 0.074	8.4 ± 3.3 ^[9]	55.8 ± 1.8 ^[5]	6.2 ± 2.5 ^[8]	
SV 786	13.03 ± 0.056	13.2 ± 2.9 ^[10]	55.8 ± 1.8 ^[5]	4.2 ± 2.7 ^[5]	
SVI 933, 944	12.76 ± 0.340	13.2 ± 2.9 ^[10]	55.8 ± 1.8 ^[5]	16.9 ± 6.9 ^[10]	
FeII 2382 ^f	< 12.01				
Components B + C					
CIII 977 ^f	≥ 14.11 ± 0.014				
OI 877 ^f	< 13.42 ± 0.221				
OIV 787 ^f	≥ 15.05 ± 0.006				
OVI 1037 ^f	≥ 14.75 ± 0.023				
SVI 944 ^f	≤ 13.16 ± 0.060				

^{a-f} Footnotes a-f are given in Table 2^g See the text regarding how and why these upper limits were estimated.

ing to its location in a complex blend, and we do not address it further in this work.

2.3.2 $z=0.51850$, PG1522+101 (Absorber 0.5A)

In this system, plotted in Figure 5 and summarized in Table 3, we mainly detect singly-ionized species, along with C III, O III, and possibly Si III, with no clear evidence for a hot phase. Indeed, we selected this system for this paper to test how

our code would work for absorbers with no affiliated highly ionized species; such absorbers are not uncommon.

The two species detected in the higher resolution HIRES spectra provide interesting insights into this absorber. The Mg II doublet requires a weaker and broader component at $\Delta v \approx +6 \text{ km s}^{-1}$ in addition to the stronger and narrower component at $\approx -5 \text{ km s}^{-1}$ (Figure 5). The same component structure is evident in the Fe II absorption profile as well, though it does not appear in the remaining detected species, presumably owing to the lower spectral resolution at which they were observed (i.e., with COS). While a single-component VP fit adequately matches the absorption profiles of the COS detections, a two-component fit encompasses a wider range of plausible solutions. For example, a two-component fit for C III, where both component velocities are co-fitted with the higher resolution Mg II lines, includes a narrower, slightly saturated C III line that is kinematically aligned with the narrower Mg II component.

Although we report two-component VP parameters in Table 3 for several detected species (see also Figure 5), we only list one component for H I, N III, and O III. In the case of O III, a two-component fit would not converge; the column density of the lower-velocity component is pushed to a negligibly low value, effectively leaving a single component fit. The weak N III absorption feature is too noisy and crowded with intervening lines from other redshifts to reasonably constrain a multiple component fit. Finally, we fit a single component to the H I Lyman series lines to get a tight measurement of the total $N_{\text{H I}}$.

In a later section, we reproduce the measured column densities of the detected species through various ionization models. We find that the uncertainties of the forced two-component fits to the COS detections are too high to meaningfully constrain a two-component or two-phase ionization model, and that a single-component, single-phase ionization model suffices to explain the total measured column densities. To be clear, we derive the total column density by summing the measured column densities of each component VP and summing their uncertainties in quadrature.

2.3.3 $z=0.68606$, PG1338+416 (Absorbers 0.6A and 0.6B)

This system (Figure 6 and Table 4) is a Ne VIII absorber that exhibits strong absorption in all ionization stages (low though high-ions). As such, it represents an excellent test for the multiphase modeling discussed below (Section 4). In this system we firmly identify three components at $v \approx -83$, -3 , and 39 km s^{-1} . There is room for a fourth component at $\approx 150 \text{ km s}^{-1}$ in the N IV and O IV profiles, though both are also contaminated by blending from other redshifts and we are unable to search for corroborating H I absorption at this velocity owing to a lack of adequate⁶ coverage in the NUV for this particular sightline. We do not include any VP-fitting

⁶ The H I in this candidate component is likely to be quite weak, so we would need to use the Ly α transition to search for H I at this velocity. At this z , Ly α is redshifted into the NUV. We do not have NUV COS or STIS spectra of PG1338+416, and the resolution of archival NUV spectra captured by HST's Faint Object Spectrograph is too low to reliably de-blend and identify weak H I lines.

results for this potential fourth component in Table 4, as we focus our analysis on the three well-detected components.

As with Absorber 0.5A, the higher resolution Mg II and Fe II spectral profiles reveal sub-components unresolved in the COS spectra. Component B splits into two sub-components, separated by $\approx 15 \text{ km s}^{-1}$. Additionally, a comparison in Component C of the HIRES vs. COS spectra suggests the presence of two sub-components. The low-ions dominate sub-components B1 and C1, while the mid- and high-ions seem to be mostly concentrated in B2 and C2. It is less likely that these misalignments between low and mid/high-ions are artifacts of COS wavelength calibration issues since we detect these velocity offsets in several pairs of low-ion lines adjacent in the spectrum to observed mid/high-ion lines. These include O II ($\lambda\lambda 832, 833$, and 834 \AA) vs. O III ($\lambda 832 \text{ \AA}$), S II ($\lambda 765 \text{ \AA}$) vs. N IV ($\lambda 765 \text{ \AA}$), and H I⁷ (Ly β $\lambda 1025 \text{ \AA}$) vs. O VI ($\lambda 1031 \text{ \AA}$). A similar offset is seen between N IV ($\lambda 765 \text{ \AA}$) and Ne VIII ($\lambda 770 \text{ \AA}$), though these ions are not as close in wavelength as the pairs previously mentioned. As is also the case for Absorber 0.5A, we do not have sufficient constraints from species detected in the higher resolution HIRES spectrum to explore ionization models for each of these sub-components separately. Instead, we use the multiphase models developed below to infer the fraction of each species in a particular phase.

While we focus our ionization modeling efforts in this work on the two main components, there is also a substantial apparent O VI absorption feature at $v \approx -63 \text{ km s}^{-1}$ that appears to coincide with possible H I, N IV, O IV, and Ne VIII features, although the velocity alignment is not perfect between these species (possibly owing to hidden substructure as evidenced above). Independent fits of N IV prefer a significantly narrower feature than O VI and a velocity offset of $\approx +16 \text{ km s}^{-1}$. O IV is severely contaminated, but the contaminating lines are positioned such that there is some inexplicable optical depth in the spectrum if O IV is removed. Hence, we can obtain reasonable measurements for the O IV column density by co-fitting its broadening parameter and velocity parameter with N IV and O VI.

3 THE IONIZATION MODELING SCHEME

As discussed in § 1.1.1, physical conditions in QSO absorption systems are traditionally inferred from simple (often single-phase) photoionization models that best fit the measured ionic column densities. Motivated by the multiphase kinematic alignment of detected species discussed below, we leverage the CASBaH spectra to explore multiphase models that expand on basic, two/three parameter models in several ways: 1) by parameterizing individually resolved absorption features (i.e., velocity components) as distinct gas clouds, which we refer to henceforth as separate *absorbers* even if they share the same systemic redshift (labels assigned to each are listed in Table 1), 2) by allowing for multiple discrete *phases* within each absorber, 3) through the addition of physically motivated parameters for the shape of the ultraviolet (UV) ionizing radiation field and gas-phase metal abundances, and

⁷ The small offset between H I and the low-ions is probed by N II ($\lambda 915 \text{ \AA}$) vs. H I ($\lambda 915.3$ and 916 \AA).

4) by applying Bayesian models to make probabilistic parametric inferences that incorporate prior knowledge. We use Markov Chain Monte Carlo (MCMC) sampling to robustly characterize the posterior distributions of the most probable models.

3.1 Kinematic Alignment of Disparate Ionization Stages

All but one (Absorber 0.5A) of the absorption systems examined here exhibit detections of high-ions such as C IV, S V, O VI, and Ne VIII alongside low-ions such as Mg II, Fe II, and C II (as well as many intermediate ionization stages, see Section 2). As summarized graphically in Figure 2, the energies required to create and destroy these species through ionization span a tremendous range extending from 7.6 to 239.1 eV. Even without any ionization modeling, this wide range of ionization stages on its own implies the presence of a multiphase entity, since no single gas temperature/density state can produce appreciable abundances of low, intermediate, and high ionization stages simultaneously. However, there is a puzzling aspect of these absorption systems. As we and others have shown in several previous studies (e.g., Tripp et al. 2000, 2006, 2008, 2011; Muzahid et al. 2012; Meiring et al. 2013; Savage et al. 2014; Burchett et al. 2015; Crighton et al. 2015; Werk et al. 2016; Rudie et al. 2019) these disparate ionization stages are often remarkably well aligned in velocity space. Their column density ratios seem to vary from component to component, but the velocity centroids and line widths of the various ions are quite similar (typically well within the measurement uncertainties). This kinematic alignment is not naturally expected in a simplistic picture of the CGM. For example, if the O VI arises in a large, volume-filling phase in the halo of a galaxy while the Mg II originates in tiny clouds distributed throughout the disk and halo of that galaxy, one might expect normal galaxy kinematics to cause these ions to have significant (and easily measured) differences in their velocity centroids and absorption-profile shapes.

Instead, the profiles of low-, intermediate-, and high-ionization stages often have remarkably similar profiles. In some instances, one can recognize small differences in centroids and/or line widths when comparing diverse ions, but overall, a wide range of ions often exhibit very similar and distinctive kinematics, which suggests that there is a direct relationship between the various ionization stages. To demonstrate this, Figures 7 and 8 show examples of the kinematic velocity alignment of the wide range of ions detected in absorbers 0.4A, 0.4B, 0.6A, and 0.6B. To compare the various ion profiles in these figures, we have constructed apparent column density profiles (e.g., Savage & Sembach 1991; Jenkins 1996), which provide linear presentations of the absorption.⁸ Briefly, in this method the apparent optical depth in each pixel as a function of velocity, $\tau_a(v)$, is used to calculate the apparent column density per unit velocity, $N_a(v) = (m_e c / \pi e^2) (f \lambda)^{-1} \tau_a(v)$, where f is the oscillator strength,

λ is the rest wavelength of the transition, and the other symbols have their usual meanings. While $N_a(v)$ profiles can be used in various ways (Savage & Sembach 1991; Jenkins 1996), our main application in this paper is to use them to directly overlay the profiles of various ions, as we have done in Figures 7 and 8.

Beginning with the PG1630+377 system in Figure 7, we see that in both Absorbers 0.4A and 0.4B, while the absorption profiles are not identical, there is a striking correspondence of the detected ions. There is an obvious difference between Mg II and the other lines, primarily on account of the higher spectral resolution provided by the Keck HIRES spectrograph that recorded the Mg II data, yet it is still useful to compare the Mg II centroids to those of the other ions (all other species in Fig. 7 were recorded with COS and thus have the same resolution). Especially in Absorber 0.4B, there is a profound similarity of the line shapes and velocity structure (i.e., the velocity offsets between the centroids of the two dominant components) of species with ionization potentials ranging from 13.6 eV (H I) to 64.5 eV (C IV) to 138.1 eV (O VI). This begs a question: how can the similarity of the profiles be reconciled with the fact that the physical conditions that will maximize one set of ions should lead to undetectable amounts of other sets of ions at the same velocity in this system? Moreover, some of the aligned species have significantly different atomic masses, so their lines should have detectably different widths if they arise in the same gas and are thermally broadened (see section 4.1 in Tripp et al. 2008). Similar line widths of, e.g., H and O, implies that the lines are *not* predominantly thermally broadened, which in turn indicates that the gas is relatively cool (Tripp et al. 2008). Could the gas be photoionized? This is an important question for a variety of reasons. The O VI ion, which is relatively easy to detect and is often assumed to be collisionally ionized, would have very different implications if it is produced by photoionization.

In Absorber 0.4A we find more indications of complexity. First, the peaks of several metals appear to have a small (2–10 km s^{−1}) offset to the blue compared to the H I peak, and there is marginal evidence that the low-ionization metal peaks are offset by the same amount compared to the intermediate- and high-ionization metals. This could reflect unresolved component structure *within* 0.4A, with a lower ionization component on the blue side of 0.4A, a higher ionization component on the red side, and H I from both components that smears together to give the appearance of a single H I feature. Alternatively, these subtle misalignments could be coincidental effects of the aforementioned COS wavelength distortions (Section 2.2). To sort this out would require higher resolution and higher S/N. Such spectra will not be available in the near future, so we must find a way to interpret the data in hand. Second, the intermediate- and high-ionizations stages are weaker, relative to the low-ionization stages, in 0.4A. Nevertheless, the more highly ionized gas is still kinematically aligned with the lower ionization gas, albeit with different $N(\text{low})/N(\text{high})$ patterns. That is, there is some type of physical relationship between the low- and high-ionization materials; these are highly unlikely to be random patches of low- and high-ionization gases that have no connection whatsoever. This is reminiscent of the behavior seen in other complex but kinematically aligned QSO absorption systems (see, e.g., Fig. 3 in Tripp et al. 2011).

⁸ The “raw” absorption profiles are exponential attenuations of the QSO light and thus are more difficult to directly scale for comparison among different elements and ions, which can have vastly different abundances and thus require some scaling to overplot their profiles.




Figure 7. Comparison of apparent column density profiles $[N_a(v)]$, see text] of several ions detected in the PG1630+377 absorption system at $z_{\text{abs}} = 0.4176$. The species that are compared are indicated in the legend of each panel. In all panels, the H I 930.804 Å line (gold diamonds) is compared to a metal ion that is scaled up by the factor listed in the legend, and regions contaminated by interloping lines from other redshifts are plotted in gray. Vertical dotted lines mark the velocities of Absorbers 0.4A and 0.4B (see Table 1).

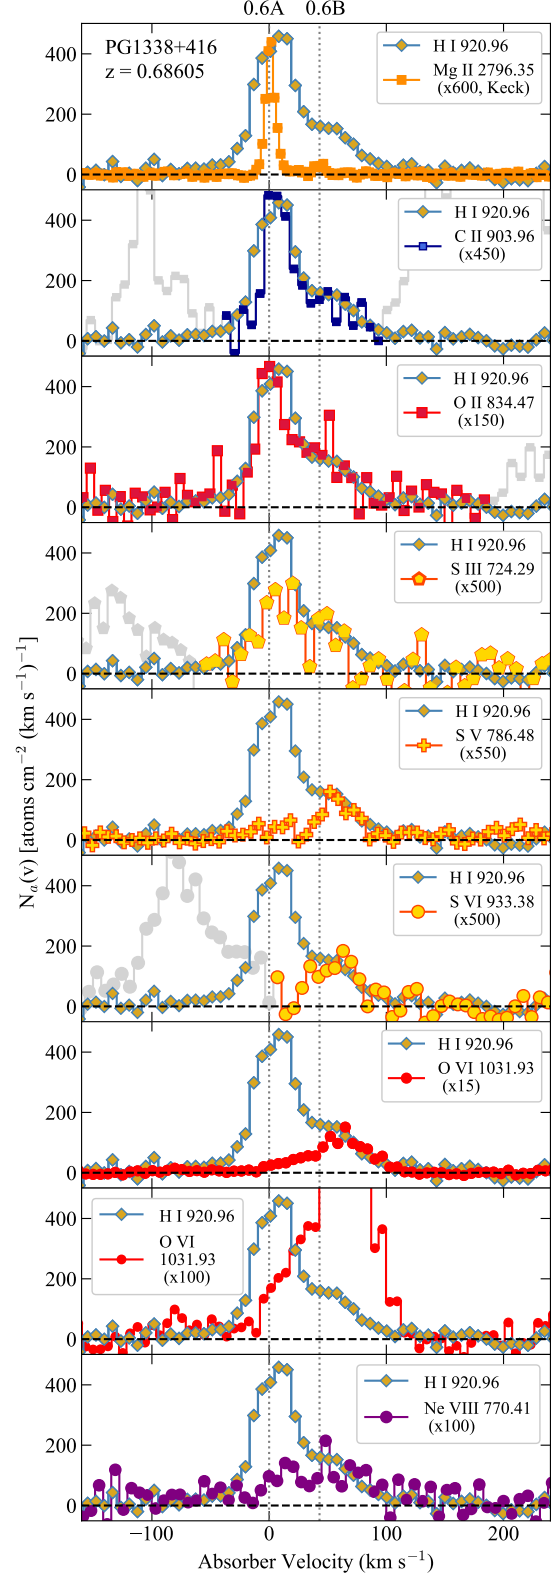


Figure 8. Comparison of H I vs. metal $N_a(v)$ profiles, as in Figure 7, for the absorption system at $z_{\text{abs}} = 0.68605$ in the spectrum of PG1338+413. In this case, the metal profiles are compared to the H I 920.963 Å line in each panel. Dotted lines indicate the velocities of Absorbers 0.6A and 0.6B (Table 1).

Absorbers 0.6A and 0.6B in the PG1338+416 system (Figure 8) exhibit very similar patterns. However, this system is arguably even more interesting because its higher redshift provides access to a greater number of more highly ionized species. Indeed, we detect S V, S VI, and Ne VIII in addition to O VI, and Figure 8 shows that the shapes of the $N_a(v)$ profiles of these high ions are remarkably similar to the shapes of the H I and lower-ionization metals in 0.6B. Like the PG1630+377 system, the lower-ionization absorption is relatively stronger in 0.6A than in 0.6B, but the high-ionization species are still present in 0.6A (see also Figure 6). Again, there are some hints that the component structure may be more complex than just two components, but better information on the component structure must await a future space telescope with greater sensitivity and spectral resolution.

The kinematic alignment of different ionization phases makes it impossible to directly determine the fraction of a given ion's absorption in each gas phase *from the spectrum alone*. While it is often assumed that the low-ion columns originate entirely in the lowest ionization gas phase (i.e., the other, more ionized phases contribute only negligible additional amounts to the low-ion column densities), in fact the ionization potentials of “low ions” are often well above 1 Rydberg (see Figure 2), and this assumption may be spurious unless it can be verified by multiphase ionization modeling that accurately models the full set of observed ions simultaneously. Indeed, this concern strongly motivates our detailed attempts to formulate a model that consistently accounts for *all* the detected absorption lines.

In many cases though, it is difficult or even impossible to identify individual velocity components in systems with densely packed component structures, particularly in spectra of moderate signal-to-noise and/or moderate spectral resolution (such as is afforded by the COS G130M and G160M gratings). Consequently, it is common practice to ignore inter-component variances altogether by integrating all absorption from a given species near a common systemic redshift into a single representative column density. This practice implicitly assumes that the absorbing medium is homogeneous in composition and physical conditions. To avoid making this assumption, we require an ionization model (along with sufficient observational constraints) that can account for the component-to-component variances in ion column density ratios.

3.2 Implementation of Multiple Phases and Components

We designed our ionization modeling scheme to flexibly fit an arbitrary combination of velocity components and ionization phases simultaneously. The ionic column densities associated with distinct velocity components are measured through absorption line fitting (see Section 2.2). Each measured column density is then resolved probabilistically into one or more phases using a Bayesian multiphase ionization model. For example, suppose O III, O IV, O V, and O VI are all detected at $v = 0 \text{ km s}^{-1}$ in a hypothetical absorption system. We could assume, for example, that this $v = 0 \text{ km s}^{-1}$ absorber consists of two discrete *phases*, including one that produces O III and O IV absorption, and a second (more highly ionized) phase that produces O V and O VI absorption. Unfortunately, this is not a sound assumption because the column density for some species (e.g., O IV) might include appreciable contri-

butions from each phase, or there might be more than two phases present. Ostensibly, it is difficult to make correct assumptions about what fraction of a given ion's column density arises in each of the phases, but we can use Bayes' Theorem coupled with a Monte Carlo Markov Chain (MCMC) method to explore the range of physical conditions in each component that are consistent with the data, and we can use metrics such as Bayes factors to judge the relative probability of, e.g., a 2-phase vs. a 3-phase model.

For each absorber in this study (Table 1), we compute models with one, two, and three ionization phases and use their Bayes factors to select the most probable model. While each phase can, in principle, be modeled with a unique set of parameters, it can be advantageous and logical to reduce the dimensionality of the parameter space significantly by assuming that the metallicity and gas composition parameters (e.g., C/α , N/α , and α/Fe) remain uniform across all the phases within a given absorber. The kinematic alignment of low- and high-ions discussed above lends credence to this idea that multiphase ions arise in gas phases that have some type of physical relationship and are likely to have similar abundances. We can also assume that the various phases in an (optically thin) absorption system are exposed to the same ionizing radiation field (i.e., in a given system, the optically thin gas is permeated and photoionized by the same UV background and/or the same light escaping from nearby galaxies). We also explore the effect of parameterizing the UV ionizing radiation field at the level of individual absorbers (one set of UV parameters per *absorber*) *versus* at the systemic level (one set of shared UVB parameters for all components in a given *absorption system*). After sharing various parameters across phases in this manner, only the gas density, temperature, and H I column density⁹ remain as parameters uniquely specified at the level of individual phases.

3.3 Bayesian Modeling

3.3.1 Parameteric Inference via Bayes Theorem

Following similar studies of QSO absorbers (Crighton et al. 2015; Fumagalli et al. 2016; Prochaska et al. 2017), this work adopts a Bayesian modeling perspective, which offers two key advantages over traditional χ^2 fitting techniques: 1) a probabilistic inference of the model parameters, and 2) a natural framework for incorporating prior knowledge. In the Bayesian paradigm, $p(\theta|\mathbf{N}, M)$ is the probability of the model parameters θ given the data (here \mathbf{N} is used to denote column densities) and a particular model M . Bayes Theorem states that $p(\theta|\mathbf{N}, M)$ is proportional to the product of the likelihood distribution $L(\mathbf{N}|\theta, M)$ and the prior probability of the parameters $p(\theta|M)$:

$$p(\theta|\mathbf{N}, M) = \frac{L(\mathbf{N}|\theta, M)p(\theta|M)}{p(\mathbf{N}|M)}. \quad (1)$$

where the prior probability of the data $p(\mathbf{N}|M)$, also known as the marginal likelihood, is a normalizing constant that depends on the particular model M . In the context of our ionization modeling application, a model $M \equiv$

⁹ The total H I column density is well-constrained in these systems, but the way the H I is distributed among the kinematically-aligned phases is unknown and is allowed to vary in our models.

$M(n_{\text{phases}}, I_{\text{mech}}^1, \dots, I_{\text{mech}}^{n_{\text{phases}}})$ is specified by the number of model phases n_{phases} , the ionization mechanism I_{mech}^i used for the i th phase, and the order in which the phases are assigned. The order of the phases is only important if the multiphase model relies on multiple types of ionization mechanisms, since it affects various “ranking priors” that we introduce below (i.e., for the density and temperature parameters). We condense this information into a compact notation by using the letters “P” and “C” to refer to photoionization equilibrium (PIE) only and a combination of photoionization *and* collisional ionization mechanisms (defined below), respectively, and by combining letters so that each phase is represented by a single letter and their order reflects their use in the ionization model (lowest ionization stage first). Thus, a “PPC” model has three phases and only the third (most highly ionized) phase incorporates P+C ionization. Sometimes, column density measurements that are not kinematically resolved across individual absorbers will necessitate joint modeling of two (or more) absorbers. For example, in the PG1338+416 system, the C III components from Absorbers 0.6A & 0.6B are strongly blended and cannot be disentangled (see Fig. 6), but we can still use the C III information by jointly modeling the two absorbers. We refer to absorbers linked in this manner as *component groups*. We can extend this notation to account for this scenario. For example, “PPCa-PCb,” indicates that Absorber 0.6A has a “PPC” model and Absorber 0.6B takes a “PC” model, specified by the joint posterior $p(\theta_{\text{PPC}}, \theta_{\text{PC}} | N_{0.6A}, N_{0.6B}, N_{0.6AB}, M_{\text{PPC}}, M_{\text{PC}})$.

The parameters inferred using Bayes Theorem are not point estimates, but fully specified probability densities. One derives the marginal posterior density of individual parameters $\theta_{i,p}$ from the full posterior $p(\theta | N, M)$ by integrating (i.e., *marginalizing*) over all the other parameters:

$$p(\theta_i | N, M) = \int p(\theta | N, M) d\theta_1 \dots d\theta_{i-1} d\theta_{i+1} \dots d\theta_n. \quad (2)$$

Wherever possible, we make use of the full parameter probabilities (i.e., the marginal posterior densities), but in some figures that rely on point estimates for clarity, we represent the corresponding marginal posterior densities by the primary mode of the distribution with uncertainties that encompass the interior 68% of the distribution in each direction. For example, a parameter p with a distribution resembling the standard log-normal distribution ($\mu=0$, $\sigma=1$) would be represented as $p=0.37_{-0.19}^{+1.86}$. We prefer this notation since it imparts some information about the skew in the underlying distribution.

The Bayesian likelihood distribution describes the probability of observing the data given a set of model parameters. We treat each measured column density as an independent random variate of a zero-truncated normal (ZTN) distribution (Equation A4) with the mean given by the multiphase ionization model and the variance set by the uncertainty in the corresponding Voigt profile fit. Empirically, the ZTN distribution appears to best describe simulated measurements of a single, noisy absorption feature across a wide range in column density (from non-detection to saturated) and signal-to-noise ratios. We model upper and lower limits as lower and upper tail probabilities, respectively, of the corresponding ZTN distribution.

The likelihood function that we employ in this paper is described in detail in Appendix A1. This likelihood function de-

pends on two key assumptions: 1) that ionic column densities derived from Voigt-profile fitting are essentially independent, and 2) that our adopted error model with fixed variance (per species) adequately reflects the uncertainty in the column density measurements. These same assumptions have been used in most of the papers in the literature that use Bayesian methodology (e.g., Crighton et al. 2015; Fumagalli et al. 2016; Glidden et al. 2016; Prochaska et al. 2017) to model quasar absorption line spectra. We expect that our likelihood function is at least accurate to first order, especially for “clean” (unblended and not saturated) measurements; the measurement of severely overlapping Voigt profiles (e.g., with velocity separations on the order of their linewidths $\delta v_{12} \lesssim b_1, b_2$) are not fully independent. We are currently investigating the relative significance of these effects in a separate project. For the present study, we will follow the literature and assume that the two assumptions are accurate to first order.

A key part of any Bayesian model is the prior distribution $p(\theta | M)$. As the ratio of the number of data points to model parameters decreases, the weight of the prior increases (relative to the likelihood). In many of our models, the number of parameters is close (within a factor of two) to the number of strong column density constraints, so the prior distribution plays an important role in our inferences. We strive to incorporate *informative* priors wherever possible since so-called *non-informative* priors can introduce undesired biases. In the following parameter overview sections, we describe various priors on both the model parameters and derived quantities. Appendix A2 formally describes $p(\theta | M)$.

3.3.2 MCMC Sampling

We sample the posterior densities of our models using an ensemble Monte Carlo Markov Chain (MCMC) algorithm, specifically the affine invariant “Parallel Tempered Ensemble Sampler” (PTES) implemented in the popular Python package EMCEE v2.2.1¹⁰ (Foreman-Mackey et al. 2013). The core (non-tempered) algorithm initializes an ensemble of “walkers”, which are similar to Metropolis-Hastings chains with the key difference being that the proposal distribution for a given walker depends on the other walkers’ positions. In our application, the tempered version of the algorithm is faster and more likely to converge than the basic algorithm. The ensemble proposal distribution also includes a scale parameter, a , that primarily affects the acceptance rate. We find that scale parameters in the range $a=1.2$ – 2 optimally balance the algorithm’s convergence rate against its ability to explore parameter space. The tempered version initializes multiple ensembles of walkers and allows each ensemble to wander through a power-law scaled version of the posterior (including one ensemble dedicated to the true, non-scaled posterior). These tempered chains can swap positions in a manner that preserves the detailed balance of the MCMC. This results in a more robust sampling of multimodal posteriors, and frequently a boost in their convergence rate as well. We initialize

¹⁰ Note, the PTES has been removed from the most recent version of EMCEE (v3.0.2) and is now the stand-alone package PTMCEE: <https://github.com/willvousedn/ptmcee>. These program changes did not alter the underlying mathematical algorithm, so it is consistent in both versions.

the sampler with 70–150 walkers and 2–3 temperings (using the provided default power-law scaling). Each walker’s initial position is uniformly sampled across the allowed parameter volume (i.e., where $p(\theta|M) > 0$), an initialization technique that we find leads more reliably to convergence for complex models than initializing all the walkers near a single best-guess point in parameter space.

MCMC algorithms are guaranteed to draw samples from the stationary distribution (e.g., the posterior) if run for an infinite duration. For finite MCMC runs, however, one must rely on heuristics to assess when the chain has moved away from the initial positions and converged on the stationary distribution. The Gelman-Rubin (GR) statistic (Gelman & Rubin 1992; Brooks & Gelman 1998) is a commonly used heuristic that tracks how well multiple independent, co-evolved chains have mixed. Unfortunately, since the PTES walkers are not individually Markovian (each ensemble is, in fact, a single large Markov Chain), it is incorrect to use the GR statistic on the ensemble walkers. One could, in principle, compute the statistic for several co-evolving ensembles, but this is typically impractical. Instead, Foreman-Mackey et al. (2013) advocate using the integrated autocorrelation time (IACT) τ_{iact} of the ensemble to determine when to cease sampling. The idea behind using the IACT as a stopping criterion is rooted in the observation that Monte Carlo integration errors decrease as a function of the number of samples: $\sigma_{\text{Monte Carlo}} \propto N^{-0.5}$. The effective sample size is reduced, however, by the IACT between samples: $N_{eff}^i = N/\tau_{iact}^i$, where $i = 1, 2, \dots, n_{par}$ and n_{par} is the number of model parameters. It makes sense to follow $N_{eff}^{min} = N/\max\{\tau_{iact}^i\}$ for determining convergence. One may stop the simulation once N_{eff}^{min} reaches a target threshold (we target $N_{eff}^{min} > 10^4$). In practice, τ_{iact}^i can be poorly estimated when N_{eff}^i is low. Hence, we wait until $N_{eff}^{min} > 100$ to begin collecting samples. In practice, we find that the IACT used in this way works well in an empirical sense: the sampled posterior density does not evolve significantly after the IACT has converged and the results are consistently reproducible.

3.3.3 Model Selection Metrics

In subsequent sections we will seek to quantify the performance of individual models relative to a set of competing models. We utilize two common (Bayesian) metrics for this purpose: 1) Bayesian evidence ratios or Bayes factors (Kass & Raftery 1995), which express the relative probabilities or “odds” of pairs of models, and 2) posterior predictive p -values (Gelman et al. 1996), which are Bayesian versions of classical p -values.

The Bayes factor K for two models M_1, M_2 is the ratio of their marginal likelihoods multiplied by the prior odds ratio of the models. This gives the posterior odds, i.e., the relative probabilities of the two models:

$$(\text{posterior odds}) = (\text{Bayes factor}) \times (\text{prior odds}) \quad (3)$$

or equivalently,

$$\frac{p(M_2|N)}{p(M_1|N)} = \frac{\int p(N|\theta_2, M_2)p(\theta_2|M_2)d\theta_2}{\int p(N|\theta_1, M_1)p(\theta_1|M_1)d\theta_1} \times \frac{p(M_2)}{p(M_1)} \quad (4)$$

Lacking information that favors any particular model over the others *a priori*, we assign a factor of unity for the prior odds of all pairs of our models and treat the Bayes factor as the

de facto posterior odds. We compute the marginal likelihood by identifying and resampling posterior modes following the approach outlined in Weinberg (2013). Details are given in Appendix A3.

By virtue of their definition as probabilities, Bayes factors allow one to consider how strongly the data support a particular model. Unlike classical hypothesis tests, which only consider evidence *against* the preferred model (the null hypothesis), Bayes factors also consider evidence *in favor of* the model. Thus, using Bayesian evidence and Bayes factors, one may rank a set of models probabilistically and *accept* those above some “decisive” odds ratio threshold rather than merely rejecting those below the threshold.

In practice, Bayes factors are sensitive to the choice of prior and the decision threshold is somewhat subjective. Additionally, accurate estimates of the marginal likelihood are often computationally challenging, particularly for models with moderate to high dimensionality and/or multimodal posteriors. To address these shortcomings, we 1) set informative priors wherever possible, and 2) model mock absorbers to assess the sensitivity and accuracy of our Bayes factor estimates. Our test results (see Appendix A4) indicate a decision threshold of $K > 10^{1.5}$ suffices to choose one model over another. Note that this corresponds to “very strong” evidence on the oft-invoked interpretation scale of Jeffreys (1961).

Bayes factors provide a useful metric for comparing two models, but unless one has computed the complete set of competing models, it is impossible to state the *absolute* probability of any given model. Thus, we would also like some criterion to reject models in an absolute sense when discrepancies between the model and data are unlikely to have arisen by chance. The posterior predictive p -value (PP p -value), is a Bayesian version of the classical p statistic that probabilistically evaluates (in simulated datasets) a model’s ability to replicate some summary statistic $T(y^{obs}, \theta)$ of the data y^{obs} and model parameters θ . For example, a PP p -value of 0.35 means that, assuming the model is true, we can expect simulated datasets y^{sim} to produce a statistic $T(y^{sim}, \theta) > T(y^{obs}, \theta)$ 35% of the time. Extending this concept, models that result in PP p -values near 0 are deemed unlikely and rejected. Conversely, it is appropriate to be suspicious about models with PP p -values very close to 1 – this could indicate that such models are overfitting the data, which compromises the insight that they provide. Models with PP p -values ≈ 1 should be subjected to extra scrutiny to evaluate whether, e.g., the model has too many parameters. The rejection threshold or significance level α and the choice of summary statistic are subjective and depend on the problem at hand. We set $\alpha=0.05$ and define $T(y, \theta)$ to be the scaled and summed residuals between observed/simulated and modeled column densities. Appendix A5 expounds on the details of this statistic.

3.4 Overview of Model Parameters

3.4.1 Ionization Mechanism

To infer the physical conditions in a given absorber, we need to know the parameters that are most likely to produce the measured ionic column densities. This necessarily involves differencing the measured and model-predicted columns, where we define the model prediction as the sum of column densities across all ionization phases assigned by the model for

each absorber (Equations A4–A6). We use version 17.00 of CLOUDY, last described by (Ferland et al. 2017), to compute the ionization equilibrium and radiative transfer through the absorbing gas. Assuming that the gas is optically thin, we can build a basic multiphase model as the composite of several single-phase models, e.g., as layers or shells of a gas cloud. For a given slab of gas, CLOUDY iteratively adjusts the thickness until $N_{\text{H I}}$ reaches a prescribed value, while simultaneously solving the radiative transfer equations and balancing the heating and cooling processes. The resulting CLOUDY solution reports the abundances of all the ionization states for the first thirty elements (through Zn). CLOUDY also has an option to fix the slab temperature to an arbitrary value. This effectively forces CLOUDY to function as a collisional ionization model while simultaneously accounting for the energy supplied by a background radiation field. Hence we label this mode photo-plus collisional-ionization (P+C). Note that in P+C models, heating and cooling are not necessarily balanced. CLOUDY’s prevalence in the literature makes it a natural choice as our core ionization model, to enable a straightforward comparison between our multiphase models and prior work. We note in passing, however, that there are other intriguing alternative ionization models that might benefit from the same multiphase framework we develop in this work (e.g., Gnat & Sternberg 2007; Oppenheimer et al. 2018; Buie et al. 2020).

Since CLOUDY models are relatively expensive to compute on-the-fly, we precomputed large grids of ionization models for both the PIE-only and P+C ionization modes at the grid points specified in Table 5. To facilitate arbitrary sampling across the parameter space these grids were then interpolated into continuous functions via 1-D Akima¹¹ interpolants applied iteratively along each dimension.

3.4.2 The UV Radiation Field

The impinging ultraviolet light is expected to be the dominant source of ionization in diffuse IGM/CGM absorbers, but local variations in its shape and intensity are poorly characterized. The UV background (UVB) contributes a substantial fraction toward the UV ionizing photon budget, yet models of the cosmic mean extragalactic radiation field (e.g., Haardt & Madau 1996, 2001, 2012; Faucher-Giguère et al. 2009) vary year-to-year in their predictions (e.g., see Figure 11 in Fumagalli et al. 2017). Unknown intrinsic uncertainties in the UVB estimates are not the only concern, however, as Upton-Sanderbeck et al. (2018) list several ionizing flux sources local to galaxies that may significantly enhance the density of ionizing photons. Additionally, the ionization in some intervening absorbers is likely impacted by their proximity to AGN (foreground to the target QSO and adjacent to the sight line), which can exhibit a wide variation in UV spectral slopes and luminosity (Scott et al. 2004; Stevans et al. 2014). This consideration may be more significant given the potential for “flickering” AGN (Oppenheimer et al. 2018) to produce ionization fossils that persist while the AGN is inactive.

To account for a wide range of physically plausible local radiation fields, we introduce two parameters that adjust the spectral shape of the UV radiation field, illus-




Figure 9. An illustration of the various UV ionizing radiation field parameterizations explored by our models. (a) In the top panel we vary the spectral slope α_{UV} (with minimal and maximal values illustrated by the red solid curves) and the He II Lyman Limit absorption strength $\beta_{\text{H I}}$ (curves with different linestyles) relative to the fiducial Haardt & Madau (2012) UVB (black curve). Ionization potential energies are marked for some commonly detected ions, to underscore the potential impact of variations in the shape of the UV spectral energy distribution on their photoionization rates. (b) The bottom panel includes a wider view of the ionizing radiation field, with each curve’s intensity normalized to the fiducial H I photoionization rate $\Gamma_{\text{H I}}^0$ as described in the text.

trated in Figure 9. First, borrowing from the approach of Crighton et al. (2015) (see also Fumagalli et al. 2016) we vary the approximately powerlaw slope of the fiducial radiation field (in this study, Haardt & Madau 2012) by applying a “hinge” at 1 Ryd and modifying the slope of the spectrum on the higher energy side (of the hinge) via the power law spectral index parameter, α_{UV} . These variations are implemented as shown in Equation 5 through the function $f(E)$. We apply a truncated normal prior to α_{UV} , specifically $\mathcal{PT}_{\text{Trunc. Normal}}(\alpha_{\text{UV}}, -1.41, 0.7, -2.5, -0.28)$ (see Equation A1), where the bounds are tied to the range of our precomputed grid of ionization models and the values for μ and σ are derived from the distribution of AGN spectral indices α' measured by Stevans et al. (2014) at rest-frame FUV (500–1000 Å) (see the blue histogram in their Figure 10).

Moderately and highly ionized species (see tickmarks in the upper panel of Fig. 9) are also sensitive to the strength of the He II Lyman Limit absorption trough ($\lambda_{\text{rest}} \approx 228 \text{ Å}$ corresponding to $\approx 10^{1.74} \text{ eV}$), and estimates of this feature in

¹¹ 1-D Akima interpolation is less prone to strong spurious oscillations with small numbers of data points (Akima 1970).

models of the cosmic mean UVB have varied significantly over the past two decades. Hence, we introduce a second scaling parameter $\beta_{\text{He II}}$ that makes the helium dip deeper or shallower.¹² Given the wide range in $\beta_{\text{He II}}$ over the limited set of UVB estimates in the literature, we place a uniform prior $p_{\text{Uniform}}(\beta_{\text{He II}}, 0, 3)$ (Equation A17) on $\beta_{\text{He II}}$ designed to encompass the variations seen in the literature (see Fig. 9a).

Given the two parameters α_{UV} and $\beta_{\text{He II}}$, the following equation defines the shape $F_\nu(E)$ of our variable radiation field for energy E :

$$\log_{10} F_\nu(E) = g(E, f(E)) \quad (5)$$

where

$$f(E) = \begin{cases} H(E), & E \leq E_0 \\ H(E) + (\alpha_{\text{UV}} + 1.4)E, & E > E_0 \end{cases}$$

and

$$g(E, x) = \begin{cases} \beta_{\text{He II}} x + (1 - \beta_{\text{He II}})(mE + b), & 1\text{Ryd} < E < 18.4\text{Ryd} \\ x, & \text{otherwise} \end{cases}$$

where $H(E)$ is the base 10 logarithm of the fiducial [Haardt & Madau \(2012\)](#) radiation field and $mE + b$ is the line between the two $(E, H(E))$ coordinate pairs evaluated at $E = 1$ Ryd and 18.4 Ryd. The value $\alpha_{\text{UV}} = -1.4$ corresponds to the approximate spectral index of $H(E)$ near redshift $z \approx 0.5$.

To fully specify the UV spectral energy distribution, we need a prescription for the intensity of the ionizing radiation. The ionization state of the optically thin gas, however, is degenerate with respect to the total ionizing intensity and the gas density, which makes it impossible to directly constrain both properties simultaneously. Instead, it is common to present model results as a function of the ionization parameter $U = n_\gamma/n_{\text{H}}$, where n_γ is the number density of ionizing photons, which scales with the total ionizing intensity of the radiation for a fixed spectral shape. In our parameterization of the UV radiation field the spectral shape is not fixed, however, so we must first normalize by the total ionizing intensity. This poses a problem for the physical interpretation of n_γ , and by extension, U , since the total ionizing intensity is sensitive to the distribution of X-ray photons in a given SED while the ionization state of the gas is not, since the tail of the H I photoionization cross-section is negligible at such high energies. As a solution to this problem, we instead define a pseudo-ionization parameter Q_Γ based on the H I photoionization rate $\Gamma_{\text{H I}}$ rather than on the photon number density:¹³

$$Q_\Gamma \equiv \frac{\Gamma_{\text{H I}}}{n_{\text{H}}}. \quad (6)$$

Like n_γ , $\Gamma_{\text{H I}}$ also scales linearly with the total ionizing intensity for a fixed spectral shape, and similarly needs to be normalized. The normalized rate is:

$$\tilde{\Gamma}_{\text{H I}} = f(\alpha_{\text{UV}}, \beta_{\text{He II}}) \cdot \Gamma_{\text{H I}}, \quad (7)$$

¹² More precisely, it adjusts the total integrated flux decrement from the unabsorbed powerlaw in the vicinity of the He II Lyman limit, with a scaling of unity returning the fiducial curve, and a maximum scaling reflecting the largest estimate of such a flux decrement found in the literature (i.e., [Haardt & Madau \(1996\)](#)).

¹³ Note, unlike U , Q_Γ is not dimensionless.

Table 5. List of Parameters, Priors, and Pre-Computed Ionization Model Grid Details

Parameter	Bounds	Step	Prior	Reference
$\log_{10} n_{\text{H}}(z)$ (cm^{-3})	$(-6.3, -1.3)^{a=0}$	a	Eqn. A15	
$\log_{10} T$ (K) ^b	(4.0, 6.4)	0.1	Eqns. A14, A17	
$[\alpha/\text{H}]^c$	(-3.0, 0.7)	d	Eqns. A13, A17	
α_{UV}	(-2.5, -0.28)	0.317	Eqns. A10, A1	
$\beta_{\text{He II}}$	(0, 3.0)	1.0	Eqns. A11, A17	
$[\text{C}/\alpha]^e$	(-0.8, 0.8)	0.4	Eqns. A12, A17	
$[\text{N}/\alpha]^e$	(-1.2, 0.4)	0.4	Eqns. A12, A17	
$[\alpha/\text{Fe}]^e$	(-1.2, 0.4)	0.4	Eqns. A12, A17	
$\log_{10} N_{\text{H}}$ (cm^{-2}) ^f	$(0, \gtrsim 23)$	—	Eqns. A16, A18	
$\sigma_{\ln P, \text{true}}^g$	(0, 3.5)	—	Eqn. A7	
$s_{\ln P}^g$	Latent Parameter	—	Eqn. A7	

^a The i -th redshift-dependent grid value of $\log n_{\text{H}}(z)$ is given by $-3.846 \times \ln(e^{1.638} - i \times e^{-2.047}) + \ln(1+z)^3$ for $i = 0, 1, \dots, 29$.

^b $\log_{10} T$ is only relevant for P+C models.

^c See Section 3.4.4 regarding why $[\alpha/\text{H}]$ instead of $[\text{Z}/\text{H}]$.

^d The i -th grid value of $[\alpha/\text{H}]$ is $3.7 \times \ln(i \times e^{-2.8344} + e^{-0.811})$ for $i = 0, 1, \dots, 13$.

^e In P+C models these parameters are applied as post-CLOUDY adjustments to shorten the computation time. Real-time processing of non-solar gas composition produces slight differences (typically $\lesssim 0.01$ dex) in the modeled electron number density and ionic column densities.

^f N_{H} is not gridded because the solution scales linearly with N_{H} for an optically thin gas (see Section 3.4.5).

^g $\sigma_{\ln P, \text{true}}^g$ is a hyperparameter and $s_{\ln P}^g$ is a property of the other parameters. Both relate to pressure equilibrium in multiphase models (see Section 3.4.6).

which depends on the normalizing function

$$f(\alpha_{\text{UV}}, \beta_{\text{He II}}) = \frac{\Gamma_{\text{H I}}^0}{\Gamma_{\text{H I}}(\alpha_{\text{UV}}, \beta_{\text{He II}})}. \quad (8)$$

For consistency with the fiducial values of α_{UV} and $\beta_{\text{He II}}$ (i.e., those which reproduce the fiducial UV ionizing radiation field), we set $\Gamma_{\text{H I}}^0 = \Gamma_{\text{H I}}^{\text{HM12}}(z)$, where $\Gamma_{\text{H I}}^{\text{HM12}}(z)$ is the redshift-dependent H I photoionization rate of the UVB from [Haardt & Madau \(2012\)](#). Compared to alternative UVB models (e.g., see [Fumagalli et al. 2017](#), Fig. 11), $\Gamma_{\text{H I}}^{\text{HM12}}(z)$ would seem to be an excellent choice as a lower limit on the true H I photoionization rate. Indeed, given the potential for extra radiation intensity from local ionizing sources (e.g., AGN), it makes sense to work with such a lower limit.

3.4.3 Gas Density and Temperature

The UV intensity normalization discussed above affects the gas density as well, since the two quantities are related by Equation 6. In conjunction with Equation 7, the adjusted density corresponding to the normalized H I photoionization rate is:

$$\tilde{n}_{\text{H}} = \frac{\tilde{\Gamma}_{\text{H I}}}{Q_\Gamma} = \tilde{\Gamma}_{\text{H I}} \left(\frac{n_{\text{H}}}{\Gamma_{\text{H I}}} \right) = f(\alpha_{\text{UV}}, \beta_{\text{He II}}) \cdot n_{\text{H}}. \quad (9)$$

The unadjusted density, n_{H} , in conjunction with an unnormalized H I photoionization rate, $\Gamma_{\text{H I}}$, describes the ionization state of the gas (Q_Γ). Equation 9 shows that upon normalizing the total ionizing intensity, now characterized by $\tilde{\Gamma}_{\text{H I}}$, we can recover the same ionization parameter Q_Γ only

if we also multiply the density n_{H} by $f(\alpha_{\text{UV}}, \beta_{\text{He II}})$. If one takes the fiducial $\Gamma_{\text{H I}}^0$ as a lower limit, then $f(\alpha_{\text{UV}}, \beta_{\text{He II}})$ becomes a lower limit by Equation 8, as suggested above, and the adjusted density is also a lower limit: $\tilde{n}_{\text{H}} \lesssim n_{\text{H}}^{\text{true}}$. Note that treating the density as a lower limit propagates to a lower limit on the derived gas pressure and an upper limit on the inferred cloud size, two properties discussed below.

In multiphase models, the *relative* density between individual phases irradiated by a common UV ionizing radiation field can be uniquely specified, unlike¹⁴ the *absolute* gas density, which depends on one’s choice of $\Gamma_{\text{H I}}^0$ (to normalize the radiation field). Throughout this work, we assume that kinematically aligned absorption components are exposed to a common radiation field. Under this assumption, we can directly compare gas densities, pressures, and cloud sizes across ionization phases inferred by multiphase models for a given absorber (or multiple absorbers if they share a common set of UV parameters) without worrying about the true normalization of the gas density.

For single-phase models, the only constraints on the density are the boundaries of the parameter interpolation grid, so we place a uniform prior on the logarithm of the density across the accessible region of parameter space. In multiphase models, however, the prior also enforces a density ranking among the phases (of a given absorber) so that a particular model phase always takes densities with the same ranking index i . For example, in a three-phase model, the first phase is always assigned the highest density, the second phase gets the intermediate density, and the lowest density goes to the third phase. When the ionization mechanism (PIE-only or P+C) is consistent across all phases, then this prior is essentially a bookkeeping exercise with no physical implications. When both PIE-only and P+C phases are intermixed, however, the inferred gas densities *are* impacted by the ranking prior. The prior is still warranted in this case though, since we generally expect that higher temperature (P+C) gas has a lower density than cooler (PIE-only) gas. The exception to this prior would be if high density, warm-hot gas were embedded in a cooler photoionized cloud, but the large implied pressure gradient should make any such state short-lived. Without this phase ranking prior, the model typically prefers the exception, but for non-physical reasons: the density is poorly constrained at higher temperatures and the cloud size prior (see below) prefers higher densities (smaller clouds). Equation A15 gives the probability density of the ranking prior.

The CLOUDY P+C models additionally include the gas temperature as a free parameter. This is particularly useful for modeling hot gas phases in a multiphase absorber. We place a uniform prior on the ‘hot’ phase temperature $\log T$ between the boundaries of the CLOUDY interpolation grid (Table 5). Note that at the lowest gas temperatures the P+C model is largely photoionized, and collisional ionization only becomes dominant at higher temperatures. In the latter regime, the

density dependence of the heating and cooling processes cancel, so the gas density has little effect on the model. Even at temperatures of $\approx 10^{5.5}$ K, however, the density dependence is not completely erased. A two-phase P+C model with a ranked density structure will elicit a ranked temperature structure (highest temperature corresponding to lowest density) without any informative temperature prior. If the temperature were independent of density, this would not occur.

3.4.4 Gas Metallicity and Composition

In addition to the standard metallicity parameter normalized to Asplund et al. (2009) solar abundances, we also include a set of parameters governing the relative abundances of individual elements. For the absorbers examined here, we tagged all the α -elements to directly follow the metallicity parameter (i.e., at solar relative abundance ratios), while the carbon, nitrogen, and iron abundances were allowed to vary with respect to the α abundances. All other species were required to follow solar relative abundances except in one absorber, which also demands a non-solar aluminum abundance (see below). The motivation for this individual treatment of C, N, and Fe is their differing nucleosynthetic origins; the relative abundances of these species are observed to depart from their solar values in various contexts (e.g. McWilliam 1997; Fabbian et al. 2009; Lehner et al. 2013; Cooke et al. 2017). To all metallicity and relative abundance parameters we applied non-informative, uniform priors across the range of the precomputed CLOUDY grid.

We detect aluminum in Absorber 0.6A, and, in exception to the scheme above, we find that an adequate fit requires tuning the $[\text{Al}/\alpha]$ relative abundance. We include $[\text{Al}/\alpha]$ as an extra parameter in our models for Absorber 0.6A and assign a truncated normal prior $p_{\text{Trunc. Normal}}([\text{Al}/\alpha], 0, 0.2, -0.5, 0.5)$.

We do not explicitly attempt to include parameters that allow for depletion of some elements by incorporation into dust, but an underabundance of Fe in the gas phase could be interpreted as a result of dust depletion instead of as the result of nucleosynthetic origins. In the absorption systems studied in this paper, we do not detect many of the refractory elements that are highly sensitive to the presence of dust, so it can be difficult in some cases to distinguish between nucleosynthesis and dust depletion. Our approach is simply to explore the relative abundances allowed by the models, which can then be interpreted and discussed in various contexts.

3.4.5 $N(\text{H})$ Scaling

All our absorbers fall in the optically thin regime $N(\text{H I}) \lesssim 10^{17}$ where the ionic column densities scale linearly with $N(\text{H I})$. This allows us to precompute a single grid of ionization models at a generic, though optically thin, value of $N(\text{H I})$ (we chose $\log N(\text{H I}) = 10^{15.5}$) and to scale the column densities to fit any of our absorbers. Once the grid has been computed, there is no reason to prefer $N(\text{H I})$ as a scaling parameter over any other ionic column density. Instead, we take the total hydrogen column ($N_{\text{H}} = N(\text{H I}) + N(\text{H II})$) as a more useful scaling parameter since it allows us to establish a direct prior on the physical extent of the absorption ‘cloud’. Each modeled ionization phase is assigned a scaling parameter, so that the column densities of the p -th phase $N(\chi_p^{\text{ion}})$ are

¹⁴ Alternatively, if one has strong prior information on $\Gamma_{\text{H I}}^0$, then $\Gamma_{\text{H I}}^0$ could be included within a Bayesian model as a ‘nuisance’ parameter. This would enable a full specification of the gas density. For example, one could use measurements of the UVB value for $\Gamma_{\text{H I}}^0$ as a prior, but only if one has high confidence that the absorber is not partially ionized by local galaxy/AGN sources. We do not take this approach given the limited number of constraints on $\Gamma_{\text{H I}}^{\text{UVB}}$ and the uncertainty in the local environment.

set by scaling the precomputed column densities $N(X_{\text{grid}}^{\text{ion}})$:

$$N(X_p^{\text{ion}}) = \frac{N_{\text{H}}^p}{N_{\text{H}}^{\text{grid}}} N(X_{\text{grid}}^{\text{ion}}). \quad (10)$$

The total hydrogen column density is closely related to the line-of-sight path-length, or “cloud size” l_{abs} of an absorber, so we encode some reasonable assumptions about the cloud size as a prior on N_{H}^p . We assume that absorbers are smaller than ≈ 100 kpc so that these absorption systems, which are often clearly composed of several components and multiple phases, can plausibly fit within a galaxy halo. At low redshifts, these assumptions are supported by CGM studies that have probed *single* galaxies with multiple sightlines to background QSOs (e.g., [Keeney et al. 2013](#); [Bowen et al. 2016](#)); these studies indicate that the size of metal-enriched CGM clouds is $\ll 100$ kpc. At intermediate and high redshifts, several groups have similarly used the multiple sightlines to images of gravitationally lensed QSOs to show that metal absorbers change on scales smaller than 100 kpc ([Rauch et al. 2001](#); [Chen et al. 2014](#)). [Rauch et al. \(2001\)](#) probed the spatial density distribution of C IV absorbers along pairs of sightlines to three gravitationally lensed quasars ($z_{\text{QSO}} = \{2.32, 2.72, 3.62\}$) and found little difference in the multi-sightline comparison of C IV column density or velocity shear at scales less than a few hundred parsecs and increasingly greater differences on larger scales, with simple estimates suggesting the absorption likely extends over ≈ 10 kpc regions. More recently, [Rubin et al. \(2018\)](#) calculated a coherence scale of $\gtrsim 2$ kpc for strong ($W_{2796} > 0.1 \text{ \AA}$) Mg II absorbers. We model these assumptions with a Gamma distribution prior over the multiphase path length l_{abs} , which we define as the total path length through all the discretely modeled gas phases p that compose a given absorber. In this work we explore $n_{\text{phases}} \leq 3$ and we approximate the path length through a single phase p by its modeled hydrogen column density to number density ratio.

$$l_{\text{abs}} = \sum_{p=1}^{n_{\text{phases}}} \frac{N_{\text{H}}^p}{\langle n_{\text{H}}^p \rangle} \quad (11)$$

where $\langle n_{\text{H}}^p \rangle = f_v n_{\text{H}}^p$. Since the hydrogen number density volume filling factor $f_v \leq 1$ is unknown, we proceed by setting it to unity and treating Equation 11 as a lower limit on the multiphase path length.

A convenient property of the Gamma distribution ensures that the sum $X + Y$ of two independent Gamma variables, $X \sim \text{Gamma}(k_1, \theta)$ and $Y \sim \text{Gamma}(k_2, \theta)$, will also be a Gamma variable $X + Y \sim \text{Gamma}(k_1 + k_2, \theta)$. Since we treat the total path-length as a Gamma variable, we can model path-lengths of individual phases (in a multiphase absorber) as independent Gamma variables by assigning a common scale parameter $\theta_{\text{path-length}}$. Furthermore, if we consider the hydrogen column density of an individual phase N_{H}^p as the path-length of the phase l_{abs}^p scaled by its density n_{H}^p ,

$$N_{\text{H}}^p = n_{\text{H}}^p l_{\text{abs}}^p, \quad (12)$$

then the scaling property of the Gamma distribution¹⁵ allows us to recast the prior on l_{abs}^p as a Gamma distribution over N_{H} , i.e., $P(N_{\text{H}}^p | n_{\text{H}}^p) = \text{Gamma}(k^p, n_{\text{H}}^p \theta_{\text{path-length}})$. We find that setting $k^p = 0.2$ and $\theta_{\text{path-length}} = 15$ kpc produces good results, though we do not attempt to rigorously optimize these

values (e.g., through sensitivity testing or by setting hyper-priors.)

3.4.6 Gas Thermal Pressure

The multiphase nature of our models allows us to investigate differences in the physical conditions between individual model phases. One important property is the variation in thermal pressures among the n_{phases} model phases, which can act as an indicator for thermal pressure equilibrium (TPE) or lack thereof across the modeled absorber. The question of whether or not an absorber exists in TPE is important because it equates to a test for hydrostatic equilibrium in the ambient medium under two assumptions: every phase in the absorber 1) exhibits the same ratio of thermal vs. non-thermal pressure support and 2) follows a similar spatial distribution. Regarding the first point, [Lochhaas et al. \(2020\)](#) conducted a detailed investigation of halo pressures in two simulated galaxies and found that thermal pressure support dominates at radii larger than $0.1 R_{\text{vir}}$. In support of the second assumption, one could argue that the simplest explanation for the observed kinematic correlations discussed in Section 3.1 and exemplified in Figures 7-8 is that the gas in different ionization phases *does* follow a similar spatial distribution.

We aim to address two questions concerning TPE and, by extension, hydrostatic equilibrium: 1) is a TPE solution *possible*, and 2) *how likely* is TPE? Suppose we treat the set of n (where typically $n=2$ or 3) thermal pressures, which we derive for each sampling of the posterior of our multiphase absorber models, as samples of an underlying log-normal pressure distribution with variance $\sigma_{P, \text{true}}^2$. This assumes that predominantly stochastic processes determine the density and temperature distributions in the underlying medium. If the underlying medium is a galaxy halo, then an approximately log-normal thermal pressure distribution is supported by [Lochhaas et al. \(2020\)](#). At the other extreme, we cannot discount the possibility that a single clump of gas plunging through the halo might exhibit non-thermal pressure variations owing to interface physics (e.g., shocks and instabilities). In this case, the methodology presented here will likely underestimate the probability of hydrostatic equilibrium. For a log-normal thermal pressure distribution, the *logarithmic* pressure in each phase, $\ln P_i$, samples a normal distribution with $\sigma_{\ln P, \text{true}}^2$, and hence the sample variance of the set of n phase pressures $s_{\ln P}^2 = \text{Var}\{\ln P_1, \dots, \ln P_n\}$ has the known sampling distribution

$$\frac{(n-1)s_{\ln P}^2}{\sigma_{\ln P, \text{true}}^2} \sim \chi^2(n-1). \quad (13)$$

Equation 13 implicitly defines $P(\sigma_{\ln P, \text{true}}^2 | s_{\ln P}^2)$, i.e., the probability that the underlying log-normal pressure distribution has a variance $\sigma_{\ln P, \text{true}}^2$ given the sample variance $s_{\ln P}^2$ calculated from the set of model phase pressures. To make a probabilistic interpretation about TPE one needs to condition $P(\sigma_{\ln P, \text{true}}^2 | s_{\ln P}^2)$ on the data, i.e., the column densities \mathbf{N} . The posterior predictive distribution for $\sigma_{\ln P, \text{true}}^2$

$$P(\sigma_{\ln P, \text{true}}^2 | \mathbf{N}) = \int P(\sigma_{\ln P, \text{true}}^2 | s_{\ln P}^2) P(s_{\ln P}^2 | \mathbf{N}) d\theta \quad (14)$$

provides a straightforward way for this to be accomplished.

¹⁵ If $X \sim \text{Gamma}(k, \theta)$, then for $c > 0$: $cX \sim \text{Gamma}(k, c\theta)$.

Here $P(s_{\ln P}^2|\mathbf{N})$ is the posterior distribution over $s_{\ln P}^2$, which in turn is computed from the posterior distributions of the phase densities and temperatures ($\ln P_i = \ln n_i + \ln T_i$ for $i = 1, \dots, n$). If we define TPE as a constraint on the underlying pressure dispersion of the modeled gas phases, i.e., $\sigma_{\ln P, \text{true}}^2 < \alpha_{\text{TPE}}$ for some threshold value α_{TPE} , then a multiphase TPE solution is 1) *possible* if $P(\sigma_{\ln P, \text{true}}^2 < \alpha_{\text{TPE}}|\mathbf{N}) > 0$ and 2) *likely* if $P(\sigma_{\ln P, \text{true}}^2 < \alpha_{\text{TPE}}|\mathbf{N}) > 0.5$. We take $\alpha_{\text{TPE}} = 0.18$, which is the value of $\sigma_{\ln P, \text{true}}^2$ when the full width at half maximum (FWHM) of its corresponding *log-normal* distribution is half its mean. The reader can assess how our interpretation of TPE might change with different values of α_{TPE} based on the information presented in Section 5.1.

The disadvantage of the posterior predictive distribution is that it precludes setting a prior constraint on $\sigma_{\ln P, \text{true}}^2$ directly. To allow for such a prior we first set a prior on $s_{\ln P}^2$ and make $\sigma_{\ln P, \text{true}}^2$ a hyperparameter of $P(s_{\ln P}^2)$. This allows us to set a hyperprior on $\sigma_{\ln P, \text{true}}^2$ and use the marginal posterior distribution $P(\sigma_{\ln P, \text{true}}^2|\mathbf{N})$ of the hyperparameter $\sigma_{\ln P, \text{true}}^2$ as a replacement for the posterior predictive distribution expressed in Equation 14. This approach has the added advantage of requiring us to specify a prior on $s_{\ln P}^2$, a derived parameter of the model that would not normally require a prior. In the absence of an explicit prior, $s_{\ln P}^2$ would align with the implicit prior *induced* by the complex intersection of priors on explicit model parameters (indeed, implicit priors are unavoidable for all derived model parameters). Setting an explicit prior on a derived parameter first requires canceling the induced one. Handley & Millea (2019) prove that the maximum entropy prior, i.e., the one that assumes the least information, is achieved with a canceling step where the joint probability of the explicit model parameters is divided by the induced prior on the derived quantity of interest, in this case $s_{\ln P}^2$ (see their Equation 3). We use kernel density estimation (KDE) to capture the prior induced on $s_{\ln P}^2$ in 10^5 random uniform samples of the explicit parameters.¹⁶ After canceling, we multiply by our desired explicit prior, in this case $P(s_{\ln P}^2|\sigma_{\ln P, \text{true}}^2)$, which is implicitly defined by the inverse of Equation 13:

$$\frac{\sigma_{\ln P, \text{true}}^2}{(n-1)s_{\ln P}^2} \sim \text{Inv-}\chi^2(n-1). \quad (15)$$

Now, $\sigma_{\ln P, \text{true}}^2$ is a hyperparameter of the prior on $s_{\ln P}^2$, and we are free to assign a hyperprior to it as intended. Further details are provided in the appendix (see Equation A7).

Several considerations affect the bounds, weight, and shape of our prior on $\sigma_{\ln P, \text{true}}^2$. First, we restrict $\sigma_{\ln P, \text{true}}^2$ to $0 \leq \sigma_{\ln P, \text{true}}^2 \leq 3.5^2$. This range corresponds to the finite support of our KDE for $s_{\ln P}^2$, outside of which $P(s_{\ln P}^2|\sigma_{\ln P, \text{true}}^2)$ quickly becomes negligible. Next, we set the shape of the prior as a mixture of a Uniform and a Half-Normal distribution and weight them so that $P(\sigma_{\ln P, \text{true}}^2 < \alpha_{\text{TPE}}|\mathbf{X}) = P(\sigma_{\ln P, \text{true}}^2 > \alpha_{\text{TPE}}|\mathbf{X}) = 0.5$, i.e., the prior probabilities of TPE and non-TPE are equal. This results in a high probability density

within the narrow TPE threshold and a low probability density across the much wider non-TPE region (see Section 5.1).

4 MULTIPHASE IONIZATION MODELING

In this section, we investigate the ionization structure of the absorbers presented in Section 2 using the Bayesian multiphase ionization modeling scheme we describe in Section 3 and Appendix A1. Our goal is to find the most probable, physically motivated, ionization equilibrium model that plausibly fits *all* the data while minimizing the number of model parameters. Beginning with a model consisting of a single phase of photoionized gas, we incrementally add discrete gas phases or alter the ionization mechanism (either PIE-only or P+C, see Section 3) of a given phase until we achieve an adequate fit. We use posterior predictive p -values to reject models that cannot plausibly reproduce the measured column densities, and we compute Bayes factors to weight and select the “best” (i.e., most probable) of the remaining models. We show that a single-phase model cannot reproduce the measured column densities in four of the five discrete absorption components that we probe in detail, and we explore various two- and three-phase alternatives. Crucially for what follows, we refer to each resolved velocity component as a unique “absorber,” reserving the term “absorption system” for groups of absorbers at a common systemic redshift along a given line-of-sight. While this is somewhat different from typical usage in the literature, it is helpful in this paper because a given component can consist of multiple phases that are modeled together and it is convenient to refer to this entity of kinematically aligned phases as an “absorber”.

We will use Absorber 0.5A (PG1522+101, see Fig. 5) to illustrate some of our analysis methods in the next section. For convenient reference in this discussion, we summarize in Table 6 some of the details of the models we will consider for this case.

4.1 Model Selection

In Section 3.3.3 we presented two quantitative model selection statistics, Bayes factors and PP p -values. Here we introduce posterior predictive plots (PP-plots) as graphical aids in evaluating the quality of our models. These PP-plots are designed to visually demonstrate the significance of discrepancies between measured and simulated columns for individual ions.

Figure 10 introduces this type of plot using real data and models for Absorber 0.5A. The panels on the left side of the figure show marginal posterior distributions, plotted in the typical “corner plot” fashion, of parameters from a 3-parameter model derived for this absorber. Parameter vectors randomly sampled from the model posterior are passed to the (potentially multiphase) ionization model to derive modeled column densities. On the upper-right side of Figure 10, we plot orange color-coded histograms of simulated column densities drawn from the 3-parameter posterior predictive distributions (see Appendix A5 for details) of the observed species. The histograms’ horizontal axes are centered at the *measured* column density values corresponding to each observed species and scaled by the corresponding measurement 1σ uncertainty. The variation in color intensity along

¹⁶ The thermal pressure variance depends on just the temperature and density, but in PIE-only models the temperature is itself a derived quantity of the model determined by all the CLOUDY ionization modeling parameters. Hence, in general the induced prior depends on all the bounded parameters included in Table 5.

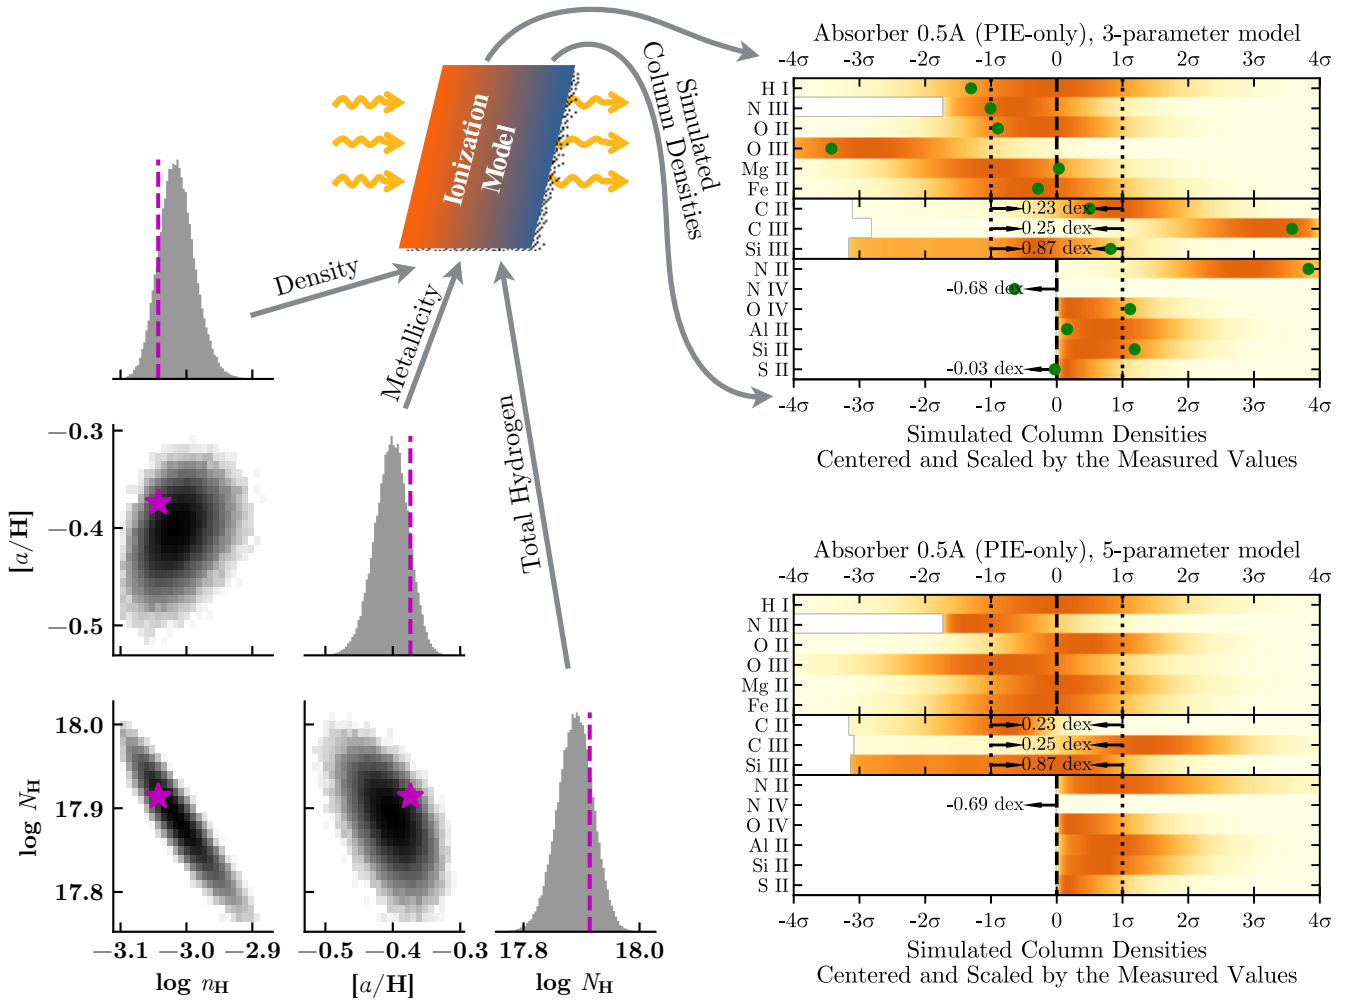


Figure 10. A comparison between PP-plots generated for an inadequate 3-parameter single-phase, PIE-only model of Absorber 0.5A (*upper-right*) and a better (see text) 5-parameter version (*lower-right*). The whole diagram also illustrates how the upper-right PP-plot is related to the posterior distribution of the 3-parameter model. The magenta stars and dashed lines in the “corner plot” (*left*) highlight a random parameter vector sampled via MCMC from the posterior of the 3-parameter model. Using the ionization model specified by these parameters, we simulate a vector of mock column densities, located as green dots in the upper-right PP-plot. The horizontal axis for the simulated columns has been standardized by subtracting the actual, measured column density values (i.e., those listed in Table 3) and scaling by their measurement uncertainties. Thus it measures the number of (measurement error) standard deviations by which the simulated column densities are in excess of the measured values. We repeat this procedure 10^5 times for independent draws from the posterior to obtain the posterior predictive distributions for individual ions, which we plot as color-coded histograms (horizontal bars) in the upper-right PP-plot. See the text and Figure 11 for a detailed description of the color mapping and the procedure for generating simulated column densities. Simulated column densities are only plotted on the constrained side of limits, but we numerically label the logarithmic offset between the peak of the posterior predictive distribution and the measured limit if the peak lies on the unconstrained side of the limit. For double bounded limits, the region interior to $\pm 1\sigma$ scales logarithmically and the interior numeric gives the logarithmic column density difference between $\pm 1\sigma$. In some cases (e.g., N III), there are sharp edges on the left side of the color-coded histograms; this happens when the model reaches the column density = 0 level.

each horizontal bar maps to the shape of the respective histogram (of simulated column densities), with the darkest regions corresponding to peak prediction (probability) density. In an ideal fit the darkest orange region should be centered at zero and predominantly fall within the $\pm 1\sigma$ measurement uncertainties. Figure 11 provides three examples to illustrate how the color maps to probability density in various scenarios (including limits).

Figure 10 also illustrates the process of generating simulated column densities from the posterior predictive distribution. A single simulated column density begins with a draw

from the posterior distribution (left side of Figure 10), which is then passed through the ionization model and converted into a vector of column densities (one for each species of interest). While the model-predicted column densities are indeed useful (e.g., for interpreting fractional degrees of ionization, as we do below), they lack a meaningful standard of comparison. Simulated column densities, however, do have a clear basis of comparison, i.e., they should look “similar” to the true measured ones. In Appendix A5, for example, we introduce one possible metric to quantify this similarity in order to compute PP p -values. To sample the posterior pre-

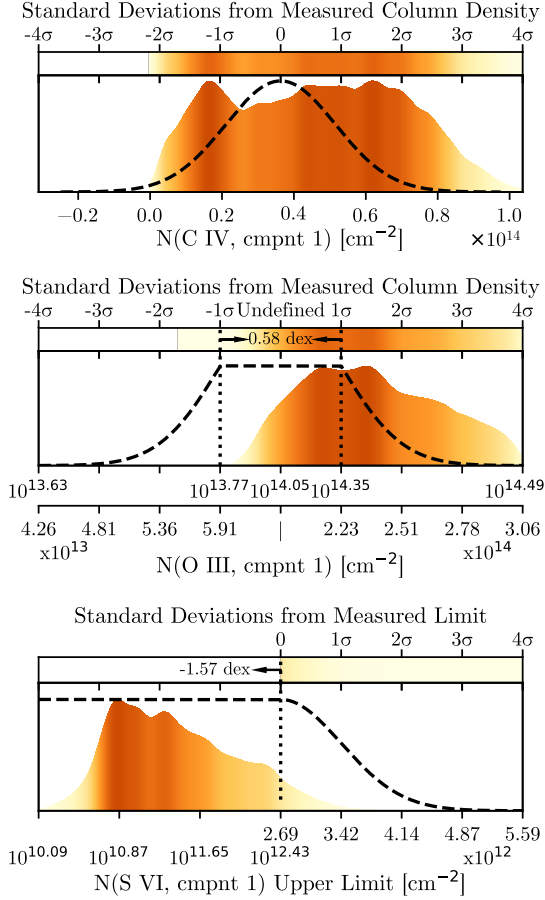


Figure 11. Expanded probability densities for arbitrary C IV (a), O III (b), and S VI (c) examples. These are intended to clarify the interpretation of the compact-bars in posterior predictive plots, such as the right side of Figure 10 and similar figures throughout this paper and how they compare model-predicted column densities to measurements and limits. The color distribution encodes peaks and valleys from the smoothed histogram of scaled column density residuals onto a compact color scale. Dashed black lines illustrate the shape of the assumed uncertainty distribution corresponding to the (mock) column density measurements and limits. Each panel includes both raw (*bottom*) and scaled (*top*) column density coordinates along the pair of horizontal axes. The black arrows in the lower two examples are explained in Figure 10. Column densities on the unconstrained side of the S VI histogram ($<0\sigma$) are logarithmically binned and are ignored by the colorbar. The O III example includes a joint constraint from both a lower and upper limit simultaneously. In this case, the measured column density has a uniform chance of falling between $\pm 1\sigma$ (vertical black dotted lines). Inside these bounds the simulated column densities are logarithmically binned, while the values outside this region are linearly binned. The histograms have been normalized to reflect a consistent probability value at any given location along the axis, although the relative area under the curve is not conserved due to the differential stretching of the horizontal axis.

dictive distribution of a given species, we draw one simulated column density from the ZTN error model (see Section 3.3 and Appendix A1) that takes a model-predicted column density and the measurement uncertainty of the corresponding measured column as its location and scale parameters. We repeat this step 10^5 times, each time using a different model-

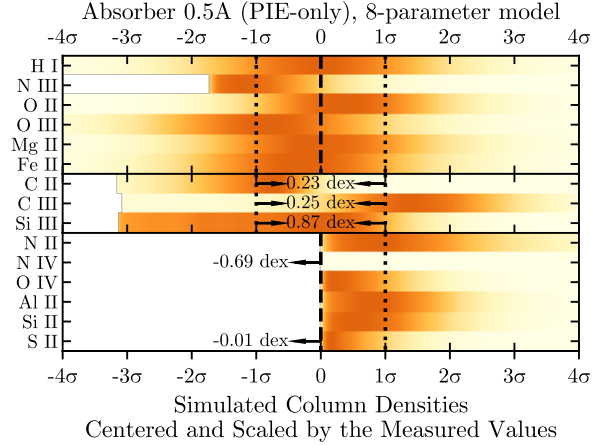


Figure 12. Scaled residuals for an 8-parameter single-phase, PIE-only model of Absorber 0.5A. This model varies the same density, metallicity, and total hydrogen column density parameters as the 3-parameter model (see Figure 10) plus the C/α and N/α relative abundance parameters added by the 5-parameter model (also in Figure 10), as well as the α/Fe , α_{UV} , and $\beta_{\text{He II}}$ parameters.

predicted column density derived as a random draw from the model posterior. The resulting simulated column densities are drawn from the posterior predictive distribution.

Inspection of the upper-right panel of Figure 10 reveals that for Absorber 0.5A, the 3-parameter PIE model's posterior predictive distributions for N II and C III both *overestimate* their respective measured columns by $\approx 3\sigma$ of the measurement uncertainty, while the prediction for O III skews to the left by $\approx 3\sigma$ (i.e., *underestimating* the measured columns). These discrepancies do not extend to the posterior predictive distributions for the remaining detected species (where we count C II and Si III as detections), nor to the Al II, Si II, S II, N IV, and O IV columns, which are all less than or approximately equal to their respective upper limits¹⁷. The bottom-right panel of Figure 10 shows the PP-plot for a slightly more complex, 5-parameter model of the same absorber. This PP-plot improves noticeably in predictive consistency compared to the three-parameter model. An 8-parameter model, shown in Figure 12, fails to improve over the 5-parameter model, however, despite its additional model flexibility. The degree to which one of these competing models is preferred over the others can be discerned using the quantitative selection metrics outlined in Section 3.3.3, but the advantage of the PP-plots is their ability to pinpoint patterns underlying the quantitative assessments.

4.2 The Single Phase Model

Among the five absorbers in this study, only Absorber 0.5A (PG1522+101, see Fig. 5) is adequately modeled with a single

¹⁷ Note that the portion of the posterior predictive histograms lying on the unconstrained side of limits are not plotted in our PP-plots, though the full histogram is used to determine the correct color mapping on the constrained side of the limits, as exemplified in Figure 11.

Table 6. Selection Metrics for Models of Absorber 0.5A

Model ID ^a	UV par. ^b	N_{pars}^c	Selection Metrics		Note ^f
			$\log_{10} K^d$	PP p^e	
PPa-C,N		7	+0.10±0.02	0.312	
PPa-C,N		7	+0.06±0.05	0.319	
Pa-C		4	0.00±0.04	0.151	
Pa-C,N		5	-0.07±0.01	0.328	
Pa-C,N,Fe		6	-0.78±0.03	0.275	
Pa-C,Fe		5	-0.97±0.01	0.125	
PCa-C,N,Fe	✓✓	11	-1.07±0.05	0.322	
Pa-C	✓	5	-1.13±0.02	0.118	
PPa-C,N,Fe	✓✓	10	-1.22±0.01	0.255	
Ca-C,N		6	-1.35±0.04	0.297	
Pa-C,Fe	✓	6	-1.35±0.01	0.120	
Pa-C,N,Fe	✓✓	8	-1.54±0.03	0.241	
PPa-		5	-1.63±0.05	0.029	Rej
PPa-C,N,Fe	✓✓	10	-1.65±0.03	0.253	
PPa-		5	-1.86±0.02	0.032	Rej
PPa-		6	-1.93±0.03	0.035	Rej
Ca-C,Fe		6	-2.14±0.06	0.117	
Pa-C,N		5	-2.49±0.21	0.033	Rej
Ca-C,N,Fe	✓✓	9	-2.52±0.65	0.236	
Pa-		3	-2.73±0.02	0.003	Rej
Pa-N		4	-2.86±0.07	0.012	Rej
Pa-C,N		5	-3.36±0.55	0.009	Rej
Pa-	✓	4	-3.62±0.40	0.004	Rej
Pa-	✓✓	5	-3.68±0.46	0.004	Rej
Ca-C		4	-3.83±0.56	0.007	Rej

^a The model identifier; the left side follows the shorthand convention discussed in Section 3.3. The right side lists any elements for which we assigned a free (relative abundance) parameter.

^b One checkmark indicates α_{UV} was a free parameter in the model, two checkmarks means the same for $\beta_{\text{He II}}$, and models with both parameters fixed (at -1.4, 1, respectively) get no checkmark. Note that checkmarks do not convey the same meaning in analogous tables for the other absorbers.

^c The number of free parameters in the model.

^d Logarithmic Bayes factors, i.e., the ratio of each model's marginal likelihood $\log Z_i$ to that of the highest ranked model $\log Z_{\text{max}}$ (in this case, $P_{5\text{par}}$, with C, N varied): $\log K_i = \log(Z_i/Z_{\text{max}})$. Models meeting our selection criterion, $\log K > -1.5$, lie above the midline.

^e Posterior predictive p -values. Models with PP p -values <0.05 are rejected.

^f "Rej" indicates that the model is rejected.

gas phase, but this absorber also has the fewest metal detections. Its spectral profiles show robust low-ion absorption but no trace of high-ionization stages. Despite the simplicity of this absorber, it is poorly fit by a basic 3-parameter model, where the density, overall metallicity, and total hydrogen column density are the only parameters. The PP-plot in the top-right of Figure 10 shows this model, and we can see why it falls short, as discussed above. The PP p -value of 0.003 for this model (see Table 6) is less than our adopted $\alpha=0.05$ significance threshold for rejection. If we introduce some additional model flexibility, allowing the carbon and nitrogen relative abundances to vary, the PP-plot improves significantly (Figure 10, bottom-right). We compare this to the 8-parameter model shown in Figure 12, which produces almost identical posterior predictive distributions despite the introduction of

three further free parameters (the iron relative abundance plus two UV parameters for eight total parameters).

While neither the 5- nor 8-parameter model can be rejected on the basis of their (satisfactory) PP p -values (0.328 and 0.241 respectively), the 5-parameter model *seems* to be strongly preferred over the 8-parameter model by a Bayes factor of $10^{1.47}$ (i.e., $10^{-0.07}/10^{-1.54}$ the reciprocal of the ratio implicit in Table 6), though this is slightly less than our decision threshold of $10^{1.5}$ – see Section 3.3.3. The comparison is ill-advised in this case, however, because Bayes factors are ideal for comparing mutually exclusive models, whereas the 5-parameter model is really an instance of the 8-parameter model with three fixed parameters *arbitrarily* selected from a continuous parameter space in the fuller model.¹⁸ Instead, it is preferable to evaluate the 5-parameter model that results from marginalizing over the three extra parameters of the 8-parameter model, which is what we choose to do in extracting our inferences about this absorber. In this particular case, the choice between models would hardly impact the parametric inferences, since there are no significant inconsistencies between the two models with respect to the marginal distributions of their five common parameters. Note, however, that comparing between models with P+C vs. PIE-only ionization mechanisms is an appropriate use of the Bayes factor, since these are exclusive models. In fact, we find that the P+C 6-parameter model (which adds the gas temperature as a free parameter to the PIE-only 5-parameter model) is disfavored relative to the PIE-only model by a Bayes factor of $\approx 10^{-1.28}$, though this is not large enough to be decisive. Interestingly, the P+C model prefers cooler temperatures anyway; in this temperature regime, photoionization dominates over collisional ionization and the P+C model is very similar to the PIE-only model.

Looking further into the plethora of model variations listed in Table 6, several things stand out. First, all the models rejected by their PP p -value (the two rightmost columns) are likewise rejected on account of their Bayes factors (4th column, below the midline). Second, the viable models appear to include a mix of single-phase and two-phase models. This is somewhat misleading, because we intuitively assign a higher prior probability to the single-phase models, which would boost the posterior odds of the single-phase models relative to the two-phase ones. This model prior is not accounted for by Table 6. Hence, we listed the best single-phase model as the reference model (i.e., where $\log_{10} K=0$) even though there are a pair of two-phase models that are ranked slightly higher on the basis of their Bayes factors alone. The contrast between two-phase and single phase models is also misleading because the posterior distributions of the two-phase models for this absorber do not deviate significantly from that of the single-phase models (which themselves are very similar). For example, the two-phase model includes two density parameters, one of which has the same distribution as the single-phase density, while the other has a flat (uniform) distribution across the explorable parameter space implying that it is unconstrained by the data.

¹⁸ If we could make a strong argument that our particular choice of fixed parameters is in some way a special combination when reducing from eight to five parameters, then perhaps Bayes factors would be appropriate for comparing these two models.

Table 7. Selection Metrics for Models of Absorbers 0.4A & 0.4B

Model ID ^a	Ind. UV ^b	N_{pars}^c	Selection Metrics		Note ^f
			$\log_{10} K^d$	PP ^e	
PPPa-CCCb		27	0.00±0.41	0.625	
PCCa-CCCb		29	-0.31±0.40	0.541	
PPCa-CCCb		28	-0.55±0.40	0.611	
CCCa-CCCb		30	-1.31±0.78	0.447	
PPPa-CCCb	✓	29	-1.45±0.78	0.613	
PCCa-PCCb		28	-1.70±0.56	0.381	
PCCa-PPb		26	-1.85±0.37	0.468	
PCCa-CCCb	✓	31	-2.00±1.30	0.537	
PPCa-PCCb		27	-2.47±0.43	0.368	
PPPa-PCCb		26	-2.59±0.88	0.363	
PCCa-PPCb		27	-2.96±0.51	0.487	
PPCa-CCCb	✓	30	-3.40±1.01	0.591	
PCa-CCCb		26	-3.51±0.66	0.356	
PPa-PPb	✓	22	-3.52±0.39	0.307	
PCa-CCCb	✓	28	-3.62±0.28	0.349	
PCa-PCCb		25	-3.72±0.30	0.148	
PCa-PPb		21	-4.27±0.23	0.254	
PPPa-PPb		24	-4.43±0.66	0.317	
PCa-PPb		23	-4.68±0.16	0.213	
PPCa-PPCb		26	-5.06±0.78	0.274	
PCa-PPCb		24	-5.35±0.37	0.225	
PPa-PCb	✓	23	-5.37±0.42	0.219	
PPCa-PPb		25	-5.55±1.16	0.292	
PPa-PPb		20	-6.27±0.18	0.105	
PPPa-PPCb		25	-6.38±0.36	0.311	
PCa-PPb	✓	23	-8.17±0.50	0.318	
PPa-PCb		21	-8.31±0.39	0.081	
PCa-PCb	✓	24	-8.34±0.63	0.234	
PCa-PCb		22	-10.76±0.78	0.116	
Pa-Pb	✓	16	-21.50±0.10	0.000	Rej
Pa-Pb		14	-46.73±0.01	0.000	Rej

^a The model identifier follows the shorthand convention discussed in Section 3.3.

^b A checkmark indicates that each absorber (0.4A & 0.4B) was modeled with an independent set of UV parameters.

^{c-f} These tablenotes follow Table 6, except that here the highest ranked model is a “joint UV” PPPa-CCCb model.

A third interesting insight from Table 6 is that the more dramatic improvements in the models for Absorber 0.5A come from varying the relative abundance parameters for C and/or N in lieu of the UV radiation field parameters (indicated by checkmarks in the second column) rather than vice versa. This seemingly stands in contrast to the models for the other four absorbers (discussed below) which show much higher sensitivity to α_{UV} in particular. Upon closer inspection, however, we find that the fiducial value of $\alpha_{\text{UV}} = -1.4$ assumed in models where we fix the parameter happens, by chance, to be close to the optimal one. Forcing α_{UV} to be harder or softer (taking values of -1 or -2 respectively) in a 5-parameter model (with C/ α and N/ α free) results in much poorer fits, with PP p -values <0.1. Note that the 8-parameter model naturally arrives at this conclusion, with a peak in the marginal distribution for α_{UV} near the fiducial value. Thus, even this simple, single-phase absorber offers some constraint on the UV radiation field.

Table 8. Selection Metrics for Models of Absorbers 0.6A & 0.6B

Model ID ^a	Ind. UV ^b	N_{pars}^c	Selection Metrics		Note ^f
			$\log_{10} K^d$	PP ^e	
PCCa-CCCb		30	0.00±0.27	0.273	
CCCa-CCCb		31	-0.69±0.39	0.255	
PPPa-CCCb		28	-0.69±0.17	0.324	
PCCa-CCCb	✓	32	-0.71±0.74	0.270	
PPCa-CCCb		29	-0.98±0.34	0.310	
PPPa-CCCb	✓	30	-1.46±0.31	0.225	
PPa-CCCb		26	-1.94±0.45	0.178	
PPa-CCCb	✓	28	-2.75±0.41	0.155	
CCCa-CCCb	✓	33	-2.79±0.31	0.217	
PPCa-CCCb	✓	31	-2.98±1.02	0.253	
PCCa-PCCb		29	-3.17±0.46	0.140	
CCCa-PCCb		30	-3.22±0.55	0.146	
PPPa-PCCb		27	-3.68±0.49	0.145	
PPCa-PCCb		28	-4.36±0.38	0.143	
PCa-CCCb		27	-13.65±0.25	0.021	Rej
PPa-PCb	✓	24	-16.32±0.25	0.001	Rej
PCa-PCb	✓	25	-16.83±0.26	0.001	Rej
PPa-PCb		22	-20.98±0.11	0.000	Rej
PCa-PCb		23	-23.68±1.02	0.000	Rej
Pa-Pb	✓	17	-62.01±0.07	0.000	Rej
Pa-Pb		15	-73.13±0.06	0.000	Rej
CCCb		16	0.00±0.27	0.225	
PCCb	✓	15	-2.30±0.39	0.317	
PPCb	✓	14	-10.18±0.17	0.012	Rej
PPb	✓	13	-14.30±0.17	0.000	Rej
PCb	✓	12	-14.55±0.34	0.000	Rej
PPb	✓	11	-14.92±0.31	0.000	Rej
CCb	✓	13	-15.37±0.74	0.000	Rej
PPPa	✓	14	0.00±0.27	0.261	
PPCa	✓	15	-0.37±0.39	0.232	
PCCa	✓	16	-1.71±0.17	0.200	
PPa	✓	12	-3.11±0.34	0.005	Rej
PCa	✓	13	-4.76±0.31	0.003	Rej
CCCa	✓	17	-5.39±0.74	0.046	Rej
CCa	✓	14	-7.86±0.45	0.005	Rej

^a The model identifier follows the shorthand convention discussed in Section 3.3.

^b A checkmark indicates that each absorber (0.6A & 0.6B) was modeled with an independent set of UV parameters.

^{c-f} These tablenotes follow Table 6, except that here the highest ranked model is a “joint UV” PCCa-CCCb model.

4.3 Additional Phases

In contrast to Absorber 0.5A (described above), the remaining four absorbers require *multiphase* models with at least two distinct phases per absorber. Tables 7 and 8 summarize the models we have tried for these absorbers, including the PP p -values for each model.

4.3.1 Multiphase Models for PG1338+416 Absorbers 0.6A and 0.6B

We first discuss the absorption system we have studied in the PG1338+416 spectrum. Figure 13 tracks the improvement for 1-, 2-, and 3-phase models of Absorbers 0.6A and 0.6B (see Figure 6) through their respective PP-plots. Single-phase models with all parameters free (16 total—eight for each

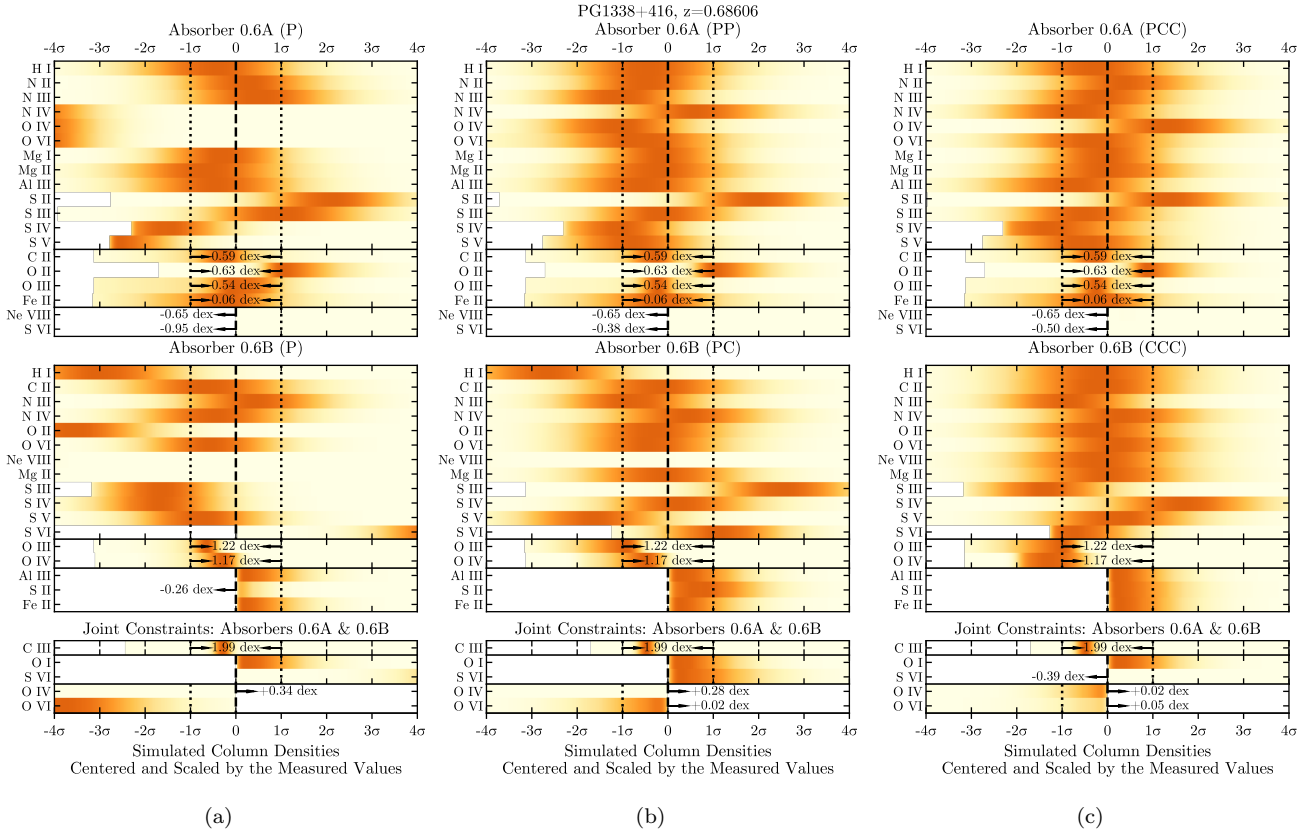


Figure 13. PP-plots for ions/limits detected in Absorber 0.6A (*top panels*), Absorber 0.6B (*middle panels*), and where the measurement could not be resolved into either absorber (*bottom panels*) are shown for selected models of increasing complexity (from left to right). These types of plots are explained in detail in Figure 10 and Section 4.1. The single-phase model (a) cannot adequately account for more than half of the species in Absorbers 0.6A and 0.6B. A two-phase model (b), which also assigns independent UV parameters to the two absorbers does better, yet still misses H I, S III, S V, and Ne VIII in Absorbers 0.6B. The addition of a third phase for Absorber 0.6B (c) results in a statistically acceptable model (PP *p*-values and Bayes factors for these models are listed in Table 8) without requiring independent UV parameters for each absorber. The letters ‘P’ and ‘C’ in the panel titles reference the number of phases used and their ionization mechanism. See Section 3.3 for a review of our convention for model specification.

absorber in the plotted PIE-only model) fail for both absorbers (Figure 13a). A single PIE-only phase dramatically underproduces the high-ions N IV, O IV, S V, and O VI in Absorber 0.6A, resulting in a PP *p*-value of 0.00023¹⁹, which falls well below our rejection threshold discussed in Section 3.3.3. A single-phase model for Absorber 0.6B performs even more poorly, only managing to fit a handful of species. To ensure that our conclusions are driven by the observations themselves and not modeling systematics (e.g., such as those arising from uncertainties in the Voigt profile and ionization model atomic parameters), we reran our single-phase models after artificially boosting the measured column density errors to a minimum of 0.1 dex (with the exception of $N_{\text{H I}}$). The slight improvements brought about by loosening these constraints fall far short of what is needed to salvage the single-phase models, implying that our conclusion about the need for multiphase models is robust.

To improve upon these single-phase ionization models, we explored adding a second and even a third model phase to each of Absorbers 0.6A & 0.6B. Each new phase introduces at

least two (three if using P+C) additional parameters to govern its density, total hydrogen column density, and (for P+C phases) temperature. In light of the observed kinematic alignment across various ionization stages (Section 3.1), which implies some similarity in the spatial distributions of the different gas phases, we assume that a common UV radiation field illuminates all the model phases associated with a given absorber. Thus additional phases share the UV parameters of the first model phase. Since both Absorbers 0.6A & 0.6B are distinct velocity components of a single absorption system at redshift $z=0.68606$, we explore two classes of models. One with unique pairs of UV parameters *per absorber* and another “joint UV” model, with a single pair of UV parameters shared across all absorbers in this PG1338+416 system.

With regard to the metallicity and relative abundance parameters, we considered two choices: 1) vary the metallicity parameters independently in every model phase, for every modeled absorber or 2) assign a single set of parameters to cover all the phases modeled for a given absorber, i.e., enforce a homogeneous gas composition in each absorber. The second choice is reasonable and realistic if the circumgalactic gas is well mixed. If the CGM is not well mixed, then situations such as a cool, metal-enriched clump of gas embedded in a

¹⁹ Computed from just the simulations and observations represented in the top panel of Figure 13a.

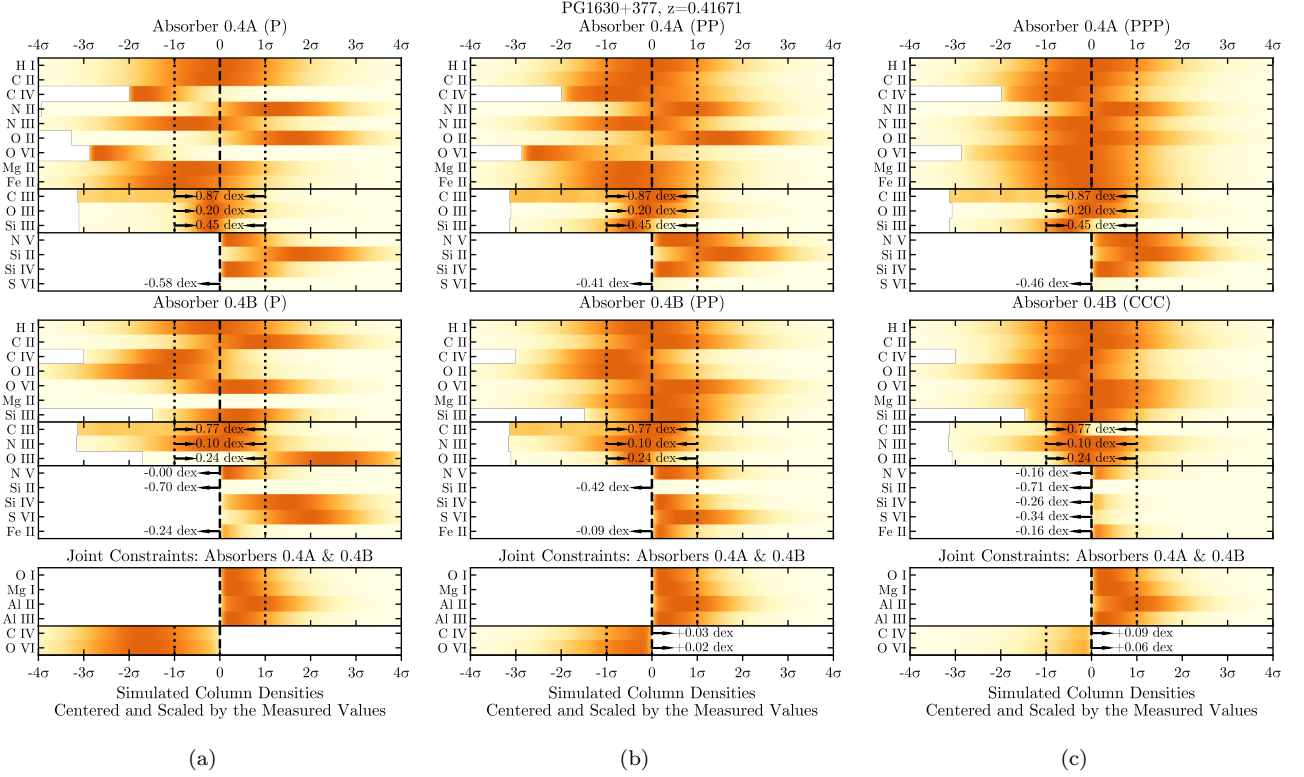


Figure 14. PP-plots for Absorbers 0.4A & 0.4B modeled with an increasing number of PIE-only phases in each panel (from left to right). The meaning of the top, middle, and bottom panels is analogous to those in Figure 13. Also as in Figure 13, the single-phase model clearly fails owing to $\gtrsim 2\sigma$ discrepancies between the simulated and observed O II, Si II, C IV, and O VI column densities in Absorber 0.4A and Mg II, O III, and S VI in Absorber 0.4B (a). In this case (as opposed to Figure 13), a two-phase model (b) provides posterior predictive distributions that are reasonably consistent with the measured columns, while a third phase offers only subtle improvements (c). We note that the mixed two-phase model (PIE-only plus P+C) shown here performs better for O VI than a PIE-only two-phase model for Absorber 0.4A (not shown), though a PIE-only three-phase model also suffices. See Figure 10 and Section 4.1 for details to aid in deciphering these PP-plots.

hot, lower-metallicity halo could occur. In practice, however, the metallicity of the high-ions is difficult to constrain and, at least for the absorbers considered here, allowing the metallicity to vary independently in each phase does not improve the model result. This conclusion is borne out by the Bayes factors, which effectively reject variable-metallicity models in favor of those having a homogeneous metallicity across all the gas phases. These simpler models achieve equally good fits with fewer parameters. Moreover, we find that variable-metallicity models tend to yield metallicity results that are quite similar to those from homogeneous (well-mixed) models anyway. Hence the models we present here all include a single set of shared gas composition parameters across all phases in a given absorber. Parameterized in this manner, each multiphase model only adds an extra two or three parameters per additional phase (per absorber) beyond those required by the full single-phase model.

Jointly modeling Absorbers 0.6A & 0.6B with two phases each noticeably improves the predictive consistency for both absorbers as shown in Figure 13b, though H I and Ne VIII still pose a problem for the model. The posterior distributions of the parameters associated with each of these two absorbers are not independent in this model on account of the joint

constraints²⁰ on the C III, O I, O IV, O VI, and S VI column densities. While the top panel of Figure 13b (Absorber 0.6A) suggests a reasonably good model based on its posterior predictive power (only the weak S II detection is off by $\gtrsim 1\sigma$), the model’s failure to reproduce H I, S III, S V, and Ne VIII for Absorber 0.6B (middle panel) accounts for its low Bayes factor ($10^{-20.98}$) and consequently its low ranking in Table 8. The key difference between 0.6A and 0.6B, as we discuss below, is the detection of Ne VIII in 0.6B. The Ne VIII is hard to explain without a third phase.

Adding a third model phase to cover the high-ions in Absorber 0.6B (Figure 13c) largely resolves the problems with the two-phase model. The model in Figure 13c also includes a third phase for Absorber 0.6A. The extra phase does not significantly improve this absorber’s PP plot (top panel), yet the model ranks higher in Table 8 because it satisfies the stringent pressure equilibrium prior, whereas the two-phase model cannot (see below, Table 9). The three-phase model presented here also differs from the two-phase one by sharing the same set of UV parameters for both absorbers (what we call “joint UV”) rather than assigning independent sets of free UV parameters to each (i.e., “independent UV”). We find

²⁰ These were derived from equivalent measurements independently of any VP measurements for the same species.

that our “joint UV” models all rank near or higher than our “independent UV” models.

4.3.2 Multiphase Models for PG1630+377 Absorbers 0.4A and 0.4B

The pair of absorbers in PG1630+377 at $z_{\text{abs}} = 0.41760$ (0.4A and 0.4B) similarly require multiphase models, as demonstrated by Figure 14. As was the case with the PG1338+416 Absorbers 0.6A and 0.6B, artificially boosting the column density errors does not resolve the failures of the single-phase models. Modeling these two absorbers with two phases each, however, *does* produce a nearly satisfactory PP-plot (Figure 14b), though the model struggles to produce enough O VI for Absorber 0.4A. As with the models for the $z_{\text{abs}} = 0.68606$ absorbers, a third phase improves not only the PP plot (Figure 14c), but also the prospects of thermal pressure equilibrium (see below, Table 9). The small number of high-ion detections in this absorber (just C IV and O VI, although O III may be almost as much of a “high ion” as C IV, see Figure 2) precludes strong constraints on a third phase in either velocity component, but Table 7 suggests that models fare better when including one.

These two absorbers (0.4A and 0.4B) underscore how the number and degree of ionization of detected species play a pivotal role in dictating the optimal number of model phases. The real significance in the number of model phases stems from the degree to which our inferences depend on them. In general, we find good agreement between two and three-phase models (for all our multiphase absorbers) where their parameters overlap. This observation, coupled with the large number of acceptable competing model variations listed in Tables 6, 7, and 8, leads us to merge competing models via a simple Bayesian model averaging procedure: the average model is obtained by drawing samples from the posterior of each of the competing models in proportion to (i.e., weighted by) their Bayes factors (listed in the same tables). Since, as we noted above, the “joint UV” models typically outperform the “independent UV” alternatives, we choose to include only the “joint UV” models in our Bayesian average. This simplistic approach glosses over several nuances related to Bayesian model averaging (see, for example, Hoeting et al. 1999), which are of minor importance to us given the large overlap in the posteriors of the competing models. Where significant inferential discrepancies do exist between classes of models, we acknowledge them below.

4.4 Ion Distribution Among the Phases

Despite assumptions that are often made in the literature,²¹ many ions can arise in a range of physical conditions, and more than one phase in a multiphase model can contribute

some appreciable fraction of a specific ion’s total column density. Here we investigate how these fractions vary from ion to ion, and we consider some insights about the ions we have studied. Figure 15 presents the constraints from our models on the fractions of ions that arise in a low-ionization phase vs. a high-ionization phase for two of our absorbers. In both panels, the solid blue curves depict the relative probability that the model phase having the highest gas density contains a given fraction (specified as a logarithm along the horizontal axis) of the total modeled column density. The dot-dashed orange curves and dashed brown curves plot the same for the intermediate and highest-density phases, respectively. Said another way, the blue, orange, and brown relative probability densities (i.e., p_{blue} , p_{orange} , and p_{brown}) reference the same posterior sample wherever $10^{x_{\text{blue}}} + 10^{x_{\text{orange}}} + 10^{x_{\text{brown}}} = 1$ is true for any set of x -axis values x_{blue} , x_{orange} , and x_{brown} corresponding to the labeled curves of matching colors. The total fraction is always unity. The colored curves are *relative* probabilities since they have been normalized by their peak probability to maximize visibility. The relative height of or area under the various colored curves is otherwise meaningless. The curves in Figure 15 depict the fraction of each ion that arises in a given phase; e.g., if a curve peaks at an x -axis value ≈ 0 , then $\approx 100\%$ of the ion column density originates in that phase.

In addition to confirming our expectation that the highest density phase produces the dominant fraction of low-ions while the majority of high-ions are located in the lowest density phase(s), Figure 15 also reveals two important results.

The first is that the low- and high-ionization gas phases are not always neatly separable. For example, $\approx 20\%$ of the neutral hydrogen in the PPP model for Absorber 0.6A (Figure 15a) is produced by the intermediate density phase, while the lowest density phase only contains $\approx 80\%$ of the total. In fact, this model predicts that there are ten other species that have $\gtrsim 20\%$ contributions from two distinct phases (i.e., C II, N II/IV, O II/III/IV, and S II/III/IV). The fraction of neutral hydrogen in a given phase though, is particularly important. While H I is often assumed to be entirely associated with a separable cold phase, this example shows that this does not have to be true. The fraction of a given species’ column density that originates in a non-dominant phase (e.g., the fraction of C III produced by the lowest ionization phase in Figure 15a) is sensitive to the ionization mechanism. Contrast the PIE-only model for Absorber 0.6A referred to above (Figure 15a) with a P+C model for the same absorber (Figure 15b). Whereas a substantial fraction of the PPP model’s low ion columns come from an intermediate ionization phase, the CCC model attributes essentially all the low ions to just the lowest density phase. The production of the intermediate phase is evidently shifted more toward the higher-ions. This makes sense, because the P+C models typically probe higher temperatures than the PIE-only ones. Finally, we also find some absorbers, such as Absorber 0.6B (Figure 15c), where the overlap between phases is largely negligible for most ions. Even in this case, however, the delineation between low-ions and mid-ions is not always clear cut. Taken as a whole, all the panels in Figure 15 demonstrate that there is not a straightforward way to reliably determine the relative contribution from multiple phases *a priori* for the majority of ions, so models that account for all ionization stages are quite valuable.

²¹ Commonly in the literature, simplifying assumptions are made such as assigning all of the H I to the lowest ionization phase (unless the H I absorption profile shows evidence of a broad component, e.g., Tripp et al. 2001; Richter et al. 2006). However, more nuanced efforts to partition the H I have appeared in more recent papers. For example, Stern et al. (2016) have proposed a model in which the H I is distributed across multiple phases.




Figure 15. Fractional, model-predicted column densities of selected ions for Absorbers 0.6A (PPP model) (*left*), 0.6A (CCC model) (*middle*), and 0.6B (*right*). The curves represent the (unnormalized) probability that a given fraction of the total model-predicted log column density resides in a particular phase (labeled in the legend). The fractional contributions of all phases in a given panel sum to unity (zero on the log scale).

The second observation is that intermediate-ions such as C III, N III, and O III can be predominantly produced in the same phase as “higher”-ions N IV, O IV, and S V (Figure 15c), in the same phase as low-ions C II, N II, and O II (Figure 15b), or in their own unique phase (Figure 15a). These same figures indicate that the sets of “low” and “high” ions might be similarly indeterminate. In particular, if the model has three phases, then the “high” ions may or may not include N IV, O IV, and S IV/v. These observations suggest that the second ionization phase in our models (the first being low-ion dominated) is distinct from both the “cold” and hot phases described by a two-phase halo paradigm. Furthermore, these observations and those of the preceding paragraph support our hypothesis in the introduction, namely, that analyses based on, e.g., fitting C II and C IV²² with a single-phase ionization model likely suffer from serious systematic errors, and in

general, caution is warranted when only limited constraints on absorber ionization are available.

Since our models involve summing column densities across phases at varying temperatures, it is important to consider whether the superposition of Voigt profiles implied by the pairs of column densities and temperatures ($N_{\text{H } i}^i$ and T^i for the i th phase) is consistent with the spectra. We carried out this consistency check for the models presented in Figures 13c and 14c and found that modeled spectra do not exceed the observed absorption profiles. In most cases, the modeled absorption lines have noticeably smaller linewidths, which we interpret as evidence for some degree of turbulent or non-thermal line broadening that is beyond the scope of the models in this paper. We did not analyze these linewidth differences further here, though we are currently working to develop an integrated Voigt profile and ionization fraction model that can encompass the extra linewidth information in a rigorous manner.

²² The C IV ion, which is commonly detected in ground-based QSO observations, would be categorized similarly to N IV and O IV in Figure 15 on account of their similar ionization potentials and photoionization rates (see Figure 2).

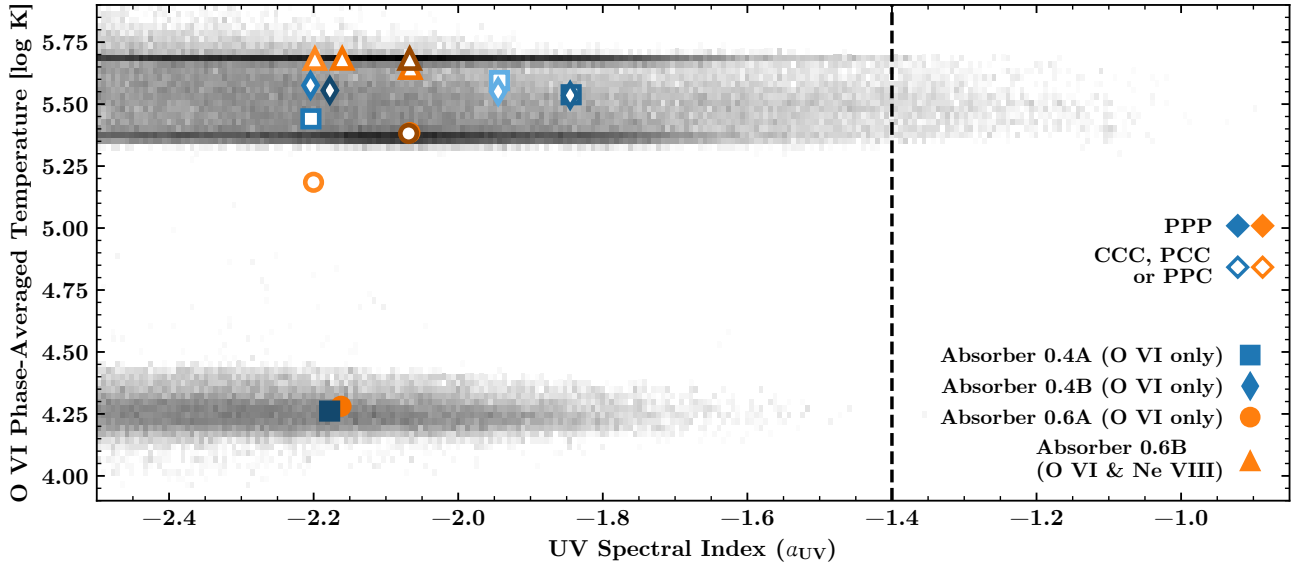


Figure 16. An illustration of the ionization regimes preferred for the high-ions by the best models. This demonstrates that most of the successful models, including all the ones constrained by a Ne VIII detection, produce O VI at high temperatures ($\log_{10}(T_{\text{O VI}}/\text{K}) \approx 5.55$ K). The vertical axis is the linear average of the gas temperature across all model phases, weighted by the fraction of (modeled) O VI in each phase. The horizontal axis tracks the UV spectral index parameter, with lower values indicating softer UV radiation. The dotted vertical line is placed at $a_{\text{UV}} = -1.4$, the slope of the [Haardt & Madau \(2012\)](#) UVB at these redshifts ($z \approx 0.5$). The log-scaled, gray 2-D histogram bins samples from the Bayesian weighted average model (see Section 4.3), while the plotted points denote mean values for each of the models with Bayes factors $> 10^{-1.5}$ in Tables 7 and 8. The middle-right legend indicates our use of closed symbols to denote purely photoionized (PIE-only) models and open symbols to denote any partially collisionally ionized (P+C) models. The symbol shapes and colors refer to particular absorbers and absorption systems, respectively. Finally, the color shading reflects the relative probability of the various models (in each color group), with darker symbols denoting the most probable models according to their Bayes factors.

4.5 Highly Ionized Gas

The question of whether photoionization can realistically produce O VI and even Ne VIII in CGM/IGM absorbers has been a recalcitrant question in the literature, and the answer is still being debated (e.g., [Tripp et al. 2001, 2008, 2011](#); [Savage et al. 2005, 2010, 2014](#); [Muzahid et al. 2012](#); [Meiring et al. 2013](#); [Werk et al. 2016](#); [Hussain et al. 2017](#); [Pachat et al. 2017](#); [Rosenwasser et al. 2018](#); [Burchett et al. 2019](#)). While the ionization state in our PIE-only models depends solely on the UV radiation field and gas density, the P+C models allow for tuning the gas heating independently of the incident radiation. When tuned to produce lower ionization states, e.g., C II and Si II, the additional flexibility afforded by the P+C ionization mechanism becomes redundant, since photoionization on its own can readily produce these low-ions. However, as we consider increasingly more ionized species, e.g., C IV, N IV, O IV, and higher ions, the number of ionizing photons with sufficient energy probably decreases rapidly, and an additional source of gas heating and ionization may be required. PIE-only models must compensate for the low density of high-energy photons (see Figure 2) either by increasing the ionization parameter (either through lowering the gas density or via increasing the intensity of the radiation field) or by varying the shape of the ionizing UV spectrum (for example, harder ionizing UV slopes will generally produce a larger O VI/C IV ratio than softer slopes).

Figure 16 shows how the ionization mechanism and radiation field are intertwined in our best multiphase models. The log-scaled, gray 2-D histogram shows the joint probability of all the multiphase models in this paper across $T_{\text{O VI}}$

a_{UV} space, where $T_{\text{O VI}}$ is the average phase temperature weighted by the fraction of O VI in each modeled phase. The histogram reveals a well-defined photoionization branch ($T_{\text{O VI}} \approx 10^{4-4.5}$) as well as collisionally ionized, warm-hot ($T_{\text{O VI}} \gtrsim 10^{5.25}$) structures. The blue and orange symbols mark point estimates for individual models.

Three patterns in Figure 16 are particularly notable. First, all the models with favorable Bayes factors (*colored symbols*) lie along the collisional ionization branch, except for one photoionization model each of Absorbers 0.4A and 0.4B. Thus, we conclude that O VI, which we detect in all four multiphase absorbers, can be readily explained by a collisional ionization origin. The two satisfactory models with photoionized O VI, however, indicate that the detection of O VI is not itself a sufficient cause to discriminate between ionization mechanisms.

Second, the two absorbers that do not produce any favorably ranked photoionization models exhibit a higher ratio of high- to low-ions than the other two multiphase models. A brief examination of Figures 4 and 6 corroborates this statement. The low-ions (e.g., C II, Mg II) are significantly weaker and the mid/high-ions (e.g., C III/IV, O III/OVI) much stronger in Absorber 0.4B (no satisfactory photoionization model) than in Absorber 0.4A (which has a good photoionization model). A similarly striking pattern emerges for Absorber 0.6B, which has a much stronger high-ion to low-ion ratio than Absorber 0.6A, as also evidenced by a strong Ne VIII line detected in the former absorber.

Our stringent pressure equilibrium prior may play an outsized role in precluding photoionization models for Absorber 0.4B. Indeed, Table 7 lists a model (PCCa-PP Pb) with

an excellent PP p -value and a Bayes factor not far outside the rejection threshold, but with zero probability of thermal pressure equilibrium (Table 9). For the Ne VIII absorber (0.6B) however, a photoionization model fails badly (see the Bayes factor for the PPPb model in Table 8) because strong Ne VIII columns are simply impossible to produce in a photoionized gas phase that obeys our modest cloud size prior, even with a multiphase model. In addition to Ne VIII, Absorber 0.6B also exhibits very strong high-ion absorption in numerous other species for which we have coverage (i.e., N V, O IV/V, Ne VIII, and S IV/V). Even ignoring O VI and Ne VIII, the other high-ions in this absorber, which exist in an intermediate ionization phase (see Figure 15c), exhibit similarly derived temperatures that on their own are indicative of collisional ionization (e.g., $T_{\text{O IV}} \approx T_{\text{S IV}} \approx 10^{5.0}$ K).

Third, the two collisional ionization bands that appear in the gray histogram in Figure 16, i.e., at $T_{\text{O VI}} \approx 10^{5.35}$ K and $10^{5.7}$ K, should be interpreted cautiously. Although they undoubtedly represent collisionally ionized O VI (and other high-ions), their narrow thermal precision (i.e., dispersion along the vertical axis) may be an artifact of the discrete nature of our models. In the likely scenario that the absorber temperatures actually follow a continuous multiphase structure, the phase-averaged temperatures in Figure 16 could mildly underestimate the distribution of true (hot) gas temperatures. Another concern is the lack of constraints at the high energy edge of the model. It is possible, for example, that the quasar sightlines probe some additional phase of hotter gas (e.g., the putative hot, virialized halo), and detection of this phase would require species such as O VII at shorter wavelengths than those covered by COS. Under this scenario it is possible that a comprehensive model constrained at all wavelengths could split the O VI and Ne VIII columns between the collisionally ionized branch in Figure 16 and some higher temperature branch, thereby forcing an (upward) adjustment to the $T_{\text{Ne VIII}} \approx 10^{5.7}$ K temperature in Figure 16.

5 PHYSICAL CONDITIONS

5.1 Pressure Gradients

An important utility of our multiphase models is their ability to track thermal pressure ($P_{\text{therm}}/k = nT$) across individual gas phases. If the absorbers are in hydrostatic equilibrium within a hot and virialized halo, then the thermal pressures of individual phases should be nearly equal modulo the assumptions discussed in Section 3.4 (e.g., assuming non-thermal pressure support is negligible). Indeed, as illustrated in Figure 17, we can find models that provide good solutions with pressure balance between the phases for all four multiphase absorbers, though we also find models with satisfactory Bayes factors where the individual phases are not in thermal pressure equilibrium. These latter models could indicate that more complex models of the CGM may be required, possibly including time-dependent physics, or this could indicate that non-thermal pressure support (e.g., from turbulence, cosmic rays, and/or magnetic fields) is important.²³ In this section, we

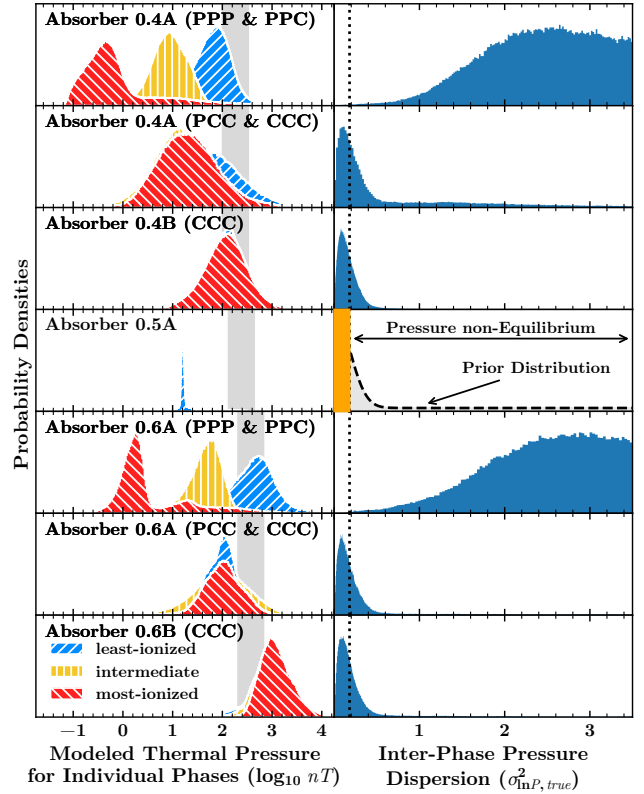


Figure 17. *Left panels:* Bayesian model averaged thermal pressures (nT) in individual ionization phases of each absorber. The blue, yellow, and red hatched histograms represent individual phases from those having the lowest degree of ionization (highest density) to the highest (lowest density), respectively. For comparison, vertical gray bands represent the range of hydrostatic halo pressures predicted by Maller & Bullock (2004) between $0.2-1R_{\text{vir}}$ for a $10^{11.5}M_{\odot}$ halo (higher halo masses would push the gray bands toward higher pressures). *Right panels:* the distribution of $\sigma_{\ln P, \text{true}}^2$, which we use to interpret the probability of thermal pressure equilibrium as discussed in the text (see also Section 3.4). The middle panel on the right side illustrates our prior on $\sigma_{\ln P, \text{true}}^2$.

present a detailed analysis of the thermal gas pressures required by the models that takes advantage of our Bayesian framework.

In Section 3.4 we laid out a framework for evaluating 1) whether thermal pressure equilibrium (TPE) is a possible solution for our absorbers, and 2) how likely any such TPE solutions are, i.e., $P(\text{TPE})$. Our approach introduces a hyperparameter to model the thermal pressure dispersion of the absorbing medium at any scale (e.g., a single gas cloud or whole galaxy halo), assuming the pressure is log-normally distributed throughout the medium. This hyperparameter, $\sigma_{\ln P, \text{true}}^2$, governs the prior probability of the modeled pressure dispersion, i.e., the sample variance $s_{\ln P}^2$ of the set of thermal pressures produced by a multiphase model for a given absorber. To ensure our posterior inferences are properly calibrated, we place a hyperprior on $\sigma_{\ln P, \text{true}}^2$ that is designed to give equal prior volume/weight to equilibrium and non-

²³ To be clear, our analysis of *thermal pressure* equilibrium vs non-equilibrium probabilities is based on models that assume *ionization* equilibrium, although we briefly speculate about the poten-

tial relevance of exploring non-equilibrium models. Any subsequent discussion of models that a

equilibrium solutions alike, and to give greater preference to lower pressure dispersions. This prior is illustrated in the middle-right panel of Figure 17 along with a vertical dotted line marking the threshold between TPE and non-TPE models (i.e., α_{TPE} from Section 3.4). Posterior samples in the narrow region left of the dotted line are consistent with our definition for TPE, while samples rightward are inconsistent with TPE.

Figure 17 clearly shows that TPE solutions exist for each of the multiphase absorbers, although we also find excellent models for Absorbers 0.4A and 0.6A that do not achieve TPE. Interestingly, there appears to be a clear distinction between the pressure solutions from models with two or more PIE-only phases (PPP and PPC) compared to models with two or more P+C phases. The key difference here is that while the P+C models can assume any range of temperatures between 10^4 K and $10^{6.4}$ K, the PIE-only phases are generally limited to $T \lesssim 10^{4.3}$ K (e.g., see Figure 16). Consequently, those models with no more than one PIE-only phase have enough flexibility to balance their thermal pressures against the range of inter-phase densities required to fit the measured data. As summarized in Table 9, for three of the four multiphase absorbers, we find favorably ranked models (based on their Bayes factors listed in Tables 7 and 8) where TPE is not only possible but probable (these models indicate $>50\%$ probability of TPE, i.e., the cumulative posterior probability that $\sigma_{\ln P, true}^2 < \alpha_{TPE}$ is greater than 0.5, which is readily apparent upon inspection of the right side of Figure 17.)

Our evidence for TPE in what are most likely CGM absorbers (based on their N_{HI} and metallicity, see below) contrasts with the prior conclusions of Werk et al. (2014), who analyzed absorption in the CGM of the COS-Halos L^* galaxies (see also, e.g., McQuinn & Werk 2018). Their work compared gas densities inferred from single-phase, photoionization models with the predictions of analytic halo models of a single and two-phase hydrostatic equilibrium CGM, finding no agreement with either one. A close examination of Figure 17 and Table 9, however, reveals that our *predominantly photoionized*, multiphase models actually show significantly different thermal pressures in the various phases, consistent with the results of Werk et al. (2014). Moreover, it is notable that the PPP and PPC (non-TPE) models for Absorbers 0.4A and 0.6A have favorable Bayes factors, because the shape of our pressure prior (middle panel on the right in Figure 17) severely disfavors non-TPE solutions. Were these models to only weakly disfavor pressure equilibrium, then we might expect a non-negligible fraction of their posterior volume for $\sigma_{\ln P, true}^2$ to lie within the threshold for TPE. Instead, they altogether reject such a scenario. In fact, even the poor-performing models (for all absorbers) in Table 9, reveal a bifurcation in their TPE probabilities, with those that are predominantly photoionized consistently having little or no success satisfying the TPE prior on $\sigma_{\ln P, true}^2$ compared to those with at least two partially collisionally ionized, P+C phases. The reason is simply that these models cannot produce a large enough spread in temperature between the relatively cool, photoionized gas phases to balance their inter-phase differences in gas density.

The left side of Figure 17 displays the (Bayesian model averaged) posterior distribution of logarithmic thermal pressures for each absorber/phase combination. Recall that these pressures scale with the intensity of the UV radiation field

Table 9. Thermal Pressure Equilibrium Probability for Multiphase Models

Absorbers 0.4A and 0.4B				Absorbers 0.6A and 0.6B			
Model ID ^{a,c}	Joint	P(TPE) ^d		Model ID ^{a,c}	Joint	P(TPE) ^d	
		UV ^b	0.4A 0.4B			UV ^b	0.6A 0.6B
PPPa-CCCb		0.0	70.9	PCCa-CCCb		59.8	64.3
PCCa-CCCb		43.8	67.3	CCCa-CCCb		69.9	64.5
PPCa-CCCb		0.4	68.8	PPPa-CCCb		0.0	64.5
CCCa-CCCb		29.9	68.6	PCCa-CCCb	✓	62.3	62.9
PPPa-CCCb	✓	0.0	66.2	PPCa-CCCb		0.2	59.1
PCCa-PCCb		45.4	52.9	PPPa-CCCb	✓	0.0	62.6
PCCa-PPb		60.2	0.0	PPa-CCCb		0.0	58.9
PCCa-CCCb	✓	40.0	69.4	PPa-CCCb	✓	0.0	65.1
PPCa-PCCb		0.3	65.6	CCCa-CCCb	✓	74.1	64.0
PPPa-PCCb		0.0	62.1	PPCa-CCCb	✓	0.0	62.2
PCCa-PPCb		58.3	0.0	PCCa-PCCb		52.5	4.2
PPCa-CCCb	✓	1.1	63.4	CCCa-PCCb		69.1	4.3
PCa-CCCb		18.6	69.9	PPPa-PCCb		0.0	4.8
PPa-PPb	✓	0.1	0.0	PPCa-PCCb		0.4	6.7
PCa-CCCb	✓	22.4	66.7	PCa-CCCb		15.7	65.2
PCa-PCCb		21.7	60.4	PPa-PCb	✓	0.0	65.1
PCa-PPb		34.7	0.0	PCa-PCb	✓	2.3	67.2
PPPa-PPb		0.0	0.0	PPa-PCb		0.0	54.5
PCa-PPb		24.9	0.1	PCa-PCb		23.7	40.4
PPCa-PPCb		0.3	0.0				
PCa-PPCb		32.8	0.0				
PPa-PCb	✓	0.0	5.4				
PPCa-PPb		0.0	0.0				
PPa-PPb		0.0	0.0				
PPPa-PPCb		0.0	0.0				
PCa-PPb	✓	14.0	0.0				
PPa-PCb		0.1	0.1				
PCa-PCb	✓	13.0	3.8				
PCa-PCb		25.0	0.3				

^{a,b} These tablenotes are identical to those in Tables 7 and 8.

^c This table includes only multiphase models from Table 7 (for Absorbers 0.4A and 0.4B) and Table 8 (for Absorbers 0.6A and 0.6B), in the same order as listed therein (i.e., best to worst, from top to bottom). The midlines, which delineate between favored and disfavored models, are likewise located between the same models as in Tables 7 and 8.

^d Percent probability of thermal pressure equilibrium, $P(\sigma_{\ln P, true} < 0.21)$, for each multiphase absorber (see Section 3.4).

owing to the degeneracy between the gas density and the UV intensity (Section 3.4.3). A consequence of our interpreting the Haardt & Madau (2012) UVB as a lower limit on the radiation field intensity (as discussed in Section 3.4.2) is that the absolute pressures we compute are likewise lower limits. The *relative* pressures of individual phases (modeled for a given absorber), however, are well characterized, so they should scale upward in lock step if one postulates an elevated (above the Haardt & Madau (2012) baseline) UVB intensity.

We compare our modeled thermal pressures in Figure 17 to the two-phase analytic halo model of Maller & Bullock (2004), which is the same model used as a point of reference by Werk et al. (2014). Strikingly, we find that the *highest* pressure phase in each of our multiphase absorbers is *less* than or approximately equal to the minimum hydrostatic equilibrium pressure predicted by Maller & Bullock (2004)

for a $10^{11.5}M_{\odot}$ halo. If the TPE models are correct, then this could indicate that our absorbers are located in sub- L^* mass halos. For the non-TPE models, this represents the high end of the pressures in individual phases. The higher temperature phases require even lower pressures. We chose to compare against a $10^{11.5}M_{\odot}$ halo since that is roughly the threshold in halo mass above which halos are expected to remain virialized to the present epoch (e.g., Kereš et al. 2005, 2009). For example, the hydrodynamical simulations of Fielding et al. (2017) for a few idealized halos suggest that non-thermal pressure support can become important within several Gyr of virialization for halos with $M_{\text{Halo}} < 10^{11.5}M_{\odot}$. Such halo sizes correspond to relatively low-mass present day galaxies, with stellar masses of $M_{\star} = 10^{9-10}M_{\odot}$ corresponding roughly to $M_{\text{Halo}} = 10^{11.0-11.5}M_{\odot}$.²⁴

A speculative explanation for the low thermal pressures suggested by some of our models may be that these absorbers exist in low mass, non-virialized halos or some other context where thermal pressure equilibrium is not attained. For example, a disk cloud that is ejected into a halo might be over-pressurized for some time as it moves out into the halo. In such a halo, it might be possible for cool gas ejected from the disk to dissipate as it expands into the low-pressure CGM. As the gas clouds expand, their density would drop, causing the ionization parameter to rise and leading to a greater degree of gas ionization until it melds with the ambient halo.

Either low overall halo pressures or significant support from non-thermal pressures could cause the non-TPE results shown in Figure 17. Ji et al. (2019), for example, find that cosmic rays added to the FIRE-2 simulations can lower the thermal pressure for the low-ion dominated gas phase by more than an order of magnitude relative to the thermal pressure without cosmic rays. According to their simulations, however, we should expect to find a thermal pressure gradient pointing in the opposite direction of what we find for the PPP and PPC models of Absorbers 0.4A and 0.6A.

5.2 Cloud Sizes

The ratio of the total hydrogen column density to the hydrogen number density provides a scale for the physical extent of a given modeled gas phase, which we can use to estimate absorber sizes (Equation 11). In Figure 18, we present such “cloud” size estimates, broken down by phase and absorber, for our Bayesian model averages (see Section 4.3). Note that all our models include a prior discouraging cloud sizes larger than 10 kpc (but with a long tail, see Section 3.4.5). Also shown for comparison is the powerlaw size vs. density relation from the multiphase CGM modeling by Stern et al. (2016). Their multiphase approach differs fundamentally from ours in that their ionic column densities are produced by gas following a powerlaw density distribution, while our approach treats the measured column densities as composites of ionization phases at discrete, constant densities.²⁵




Figure 18. Cloud scale lengths $l \geq N_{\text{H}}/n_{\text{H}}$ of individual model phases vs. the hydrogen gas density normalized by the cosmic mean baryon number density (taken here as $\bar{n}_b = 1.85 \times 10^{-7}(1+z)^3$). These are extracted from the Bayesian weighted average model (see Section 4.3). The top axis provides a conversion to hydrogen number density (i.e., the parameter used in our models) at redshift $z=0.51850$; one can shift the units of the top axis by -0.09 or $+0.14$ to convert to redshift $z=0.41760$ or $z=0.68606$, respectively. Shapes indicate phases pertaining to the averaged-models of specific absorbers, and filled/semi-filled/open symbols denote low/intermediate/high-ion dominated phases within each absorber. For comparison, the gray bands cover the range of densities predicted for a $10^{11.5}M_{\odot}$ halo by the two-phase analytic halo model of Maller & Bullock (2004). The isolated, leftmost band shows the predicted hot-phase density between $0.2-1R_{\text{vir}}$ while the wider, rightmost bands show the predicted cool-phase density at the same radii for temperatures $T=10^4$ K (*thin band*) and $T=10^{3.7-4.3}$ K (*wide band*).

Examining individual model phases, we find that the cloud size correlates with the degree of ionization, e.g., the phase having the largest scale length is also the one that produces the dominant fraction of high-ions. The intermediate (in ionization) phases of the Bayesian weighted average models for Absorbers 0.4A, 0.4B, 0.6A, and 0.6B are also intermediate in cloud size, accounting for approximately 6%, 47%, 5%, and 3%, of the respective absorber’s total modeled path length, assuming a volume filling factor of unity in every phase, (i.e., $f_v = 1$ in Equation 11). For dense, low-ion dominated phases these numbers shrink to approximately 0.1%, 0.5%, 2.3%, and 0.03%, respectively. The squares of these 1-D distances roughly approximate the relative probability of observing each of the two or three phases in arbitrary lines-of-sight. For example, if our absorbers are structured in a layered spherical shell geometry with cold gas embedded within successively hotter and volumetrically larger gas phases, then our models would imply an O VI cross section that is $\gtrsim 300$ and $\gtrsim 3000$ times larger than the associated (i.e., kinematically aligned) O III and Mg II cross sections, respectively. While higher ionization stages do appear more frequently in blind redshift surveys (see, e.g., Fig. 10 in Tripp et al. 2008) their enhanced

²⁴ Based on the halo-to-stellar mass ratios of Guo et al. 2010.

²⁵ Other key differences include our parameterization of the UVB and relative abundances, and our focus on modeling individual velocity components. Additionally, they adopted 0.2 dex uncertainties for all their measurements, while most of our measurements are at least three times more restrictive.

detection frequency relative to that of the lower ionization stages is much less than the multiple orders of magnitude discrepancy anticipated by our simple geometrical argument. Various theories have been proposed to solve this conundrum. For example, the lower ionization phases could be dispersed throughout the larger high-ion phase as a fine “mist” (e.g., [McCourt et al. 2018](#); [Liang & Remming 2019](#); [Gronke & Oh 2020b](#)), which would effectively increase the lower-ions’ cross-section. Indeed, [Werk et al. \(2019\)](#) find empirical evidence of small (≤ 10 pc) clumps of low-ions embedded in much larger (≥ 1 kpc) coherent mid-/high-ion structures. Alternatively, cool gas could form as a shocked, thin shell surrounding a larger hot bubble (e.g., [Thompson et al. 2016](#)). It is also possible that cold clouds form in situ in complex, multiphase outflows from galaxies (e.g., [Gronke & Oh 2020a](#)). The overall sizes of our absorbers are small enough to allow any of these scenarios. This contrasts with other studies in the literature (e.g., [Hussain et al. 2017](#)), which infer significantly larger cloud sizes that would be in tension with some of the theoretical studies above. Our smaller sizes partly result from our imposed prior penalty on large cloud sizes. This does not lessen our conclusion, however, that viable models of highly ionized clouds can be constructed that do not require tremendously large clouds.

5.3 The Ionizing Radiation Field

Models of the extragalactic ionizing flux background are fundamental to our understanding of the IGM/CGM and other topics, such as reionization, yet observational constraints on the shape and intensity of the UV background are scarce. Not only are models of the *mean* extragalactic radiation field (e.g., [Haardt & Madau 1996, 2001, 2012](#); [Faucher-Giguère et al. 2009](#)) intrinsically uncertain, but in the vicinity of AGN and star forming galaxies, their uncertainty is compounded by the potential for local additions to the background ionizing flux (see [Upton-Sanderbeck et al. 2018](#)). The column density ratios of many ions commonly detected in CGM/IGM absorbers are highly sensitive to the UV radiation field (see, for example, Figure 9a). Consequently, line-rich QSO absorption systems with excellent quality spectra present an important tool for constraining *local* variations in the ionizing UV radiation field.

To extract information about the UV radiation field from our absorbers, we parameterize the ionizing UV spectrum in our models as a modified powerlaw, with a variable slope and He II Lyman Limit absorption “trough” (see Section 3.4 for details). As described in Section 4.3, we model the ionizing flux in two ways, 1) with independent UV parameters for each absorber (“independent UV”), and 2) with a shared set of UV parameters covering all absorbers at the same systemic redshift (“joint UV”). Both approaches can produce models with acceptable Bayes factors (see Tables 7 and 8), but we find that “joint UV” models generally score better than the “independent UV” versions of the same models.

Figure 19a plots marginal posterior densities for α_{UV} and $\beta_{He II}$ for each absorption system. These suggest that the four multiphase absorbers (Absorbers 0.4A, 0.4B, 0.6A, and 0.6B) favor a softer UV radiation field than the fiducial one ([Haardt & Madau 2012](#)). Specifically, the most probable value of α_{UV} for the absorption system at $z=0.41760$ is ≈ -2.3 , and for the two absorbers at $z=0.68606$ it peaks at ≈ -2.1 . While we

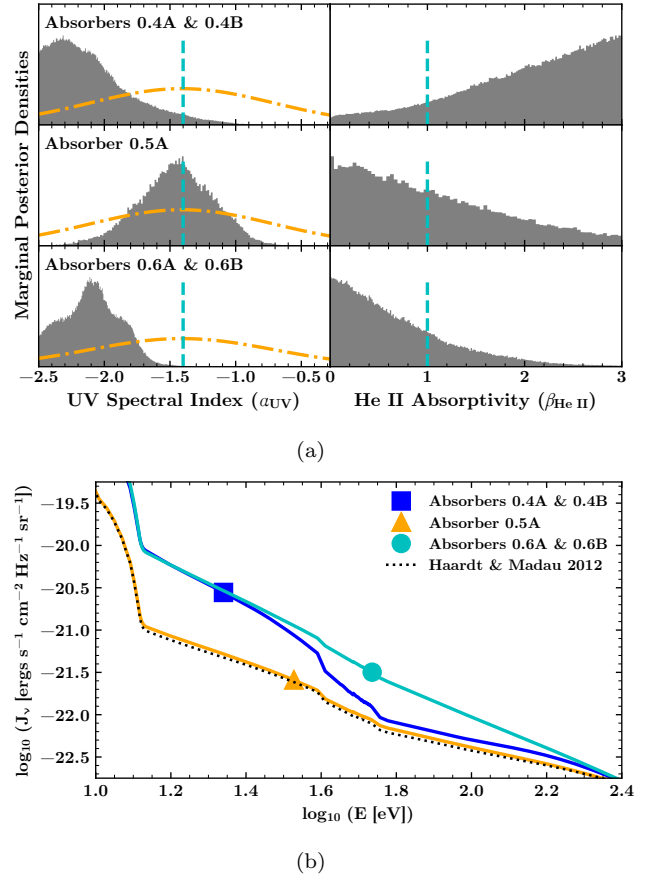


Figure 19. (Top) Full probability densities for the UV parameters inferred from our Bayesian model-averaged posteriors. The vertical dashed cyan lines mark the fiducial values, and the orange dash-dotted curves illustrate the prior density for α_{UV} . Note that the distributions plotted for Absorber 0.5A were derived from the subset of acceptable models that include α_{UV} and/or $\beta_{He II}$ as free parameters (see Table 6). (Bottom) Realizations of the most probable spectra derived from the posteriors shown in the top panel. The black dotted line illustrates the fiducial [Haardt & Madau \(2012\)](#) spectrum at redshift $z=0.68606$. For plotting purposes, we normalized all the UV spectral energy distributions by $\Gamma_{HI}^{HM12}(z=0.68606)$ (details in Section 3.4) to minimize redshift-dependent intensity variations.

understand that the ionizing UV spectral index of ([Haardt & Madau 2012](#)) ($\alpha_{UV}^{HM12} \approx -1.4$) is itself derived from uncertain models, it is unlikely that the cosmic UV background is as soft as our multiphase absorbers indicate. One possibility is that a nearby AGN with a comparably soft ionizing flux spectrum might boost the local radiation field ionizing flux density and tilt its spectral slope. Indeed, the distribution of spectral indices found by [Telfer et al. \(2002\)](#) in their subset of radio-loud AGN, with a mean of $\alpha_{UV} = -1.96 \pm 0.12$, supports this possibility. Other studies of composite ultraviolet quasar and AGN spectra, e.g., [Scott et al. \(2004\)](#) and [Stevens et al. \(2014\)](#), are also consistent with local values of ≈ -2.3 to -2.1 for individual cases. Alternatively, the ionizing flux field for these absorbers could have an important contribution from nearby galaxies. [Upton-Sanderbeck et al. \(2018\)](#) describe several potential galactic and circumgalactic sources of ionizing radiation associated with galaxies in addition to the omnipresent UV background. Radiation from hot

stars, for example, boosts the longer UV wavelengths more than the shortest ones, which, when combined with the UVB would help to soften its spectral index.

The second parameter governing the shape of the UV radiation field in our models, $\beta_{\text{He II}}$, provides more clues as to the origin of the extra UV flux. $\beta_{\text{He II}}$ adjusts the depth of the He II trough at 4 Ryd, which is sensitive to the amount of helium along the path from the radiation source to the absorber. While the effective helium opacity becomes significant for distant sources, radiation from a local AGN is unlikely to encounter sufficient helium to produce a significant effect. In contrast, stellar atmospheres generally imprint a strong He edge on stellar spectra. Thus, if galaxies (stellar sources) are leaking ionizing radiation, we generally expect a deeper He trough in the ionizing UV flux than the He feature of the diffuse UV background alone. Using these qualitative differences between AGN and stellar sources of ionizing radiation, the posterior distributions of $\beta_{\text{He II}}$ from models of Absorbers 0.4A and 0.4B (a strong He II trough) may indicate that the extra radiation has a stellar origin, while the models of Absorbers 0.6A and 0.6B (little/no He II trough) are consistent with ionizing flux contributions coming from an AGN.

Figure 19b offers another perspective on the inferred UV parameters. The cyan and blue lines reproduce the most probable UV ionizing flux SED for the multiphase absorbers. Their shapes (α_{UV}) look very similar, except that the absorption trough at 4 Ryd is strong in Absorbers 0.4A and 0.4B and almost non-existent in Absorbers 0.6A and 0.6B. Both, however imply an extra component of ionizing flux above the fiducial Haardt & Madau (2012) estimate (dotted black line). In contrast, the shape of the single-phase absorber (Absorber 0.5A) broadly aligns with Haardt & Madau (2012), based on those models for this absorber where the UV radiation field is allowed to vary.

5.4 Metallicity and Relative Abundance Ratios

Our ionization models all include parameters governing the absorbers' metallicity and deviation from solar C/ α , N/ α , and α /Fe ratios. Figure 20 plots point estimates of the marginal posterior densities of these parameters. Recall that all of the non- α elements we use to constrain our models are assigned independent abundance parameters. We use $[\alpha/\text{H}]$ as an indicator of the overall absorber metallicity. We chose to keep the metallicity and enrichment parameters constant across all phases within each absorber as discussed in Section 4.3. In three-phase models the metallicity of the highest ionization phase is minimally constrained owing to the limited number of high-ion detections, so the resulting metallicities are primarily driven by the two lowest ionization phases.

Unsurprisingly, all five of the absorbers are metal-enriched with α -metallicities greater than a quarter solar, a fact that we attribute to the selection of line-rich absorbers for this study. The metallicities are consistent with the distribution of metallicities measured in previously studied absorber samples with similar neutral hydrogen column densities (Lehner et al. 2013; Wotta et al. 2016; Prochaska et al. 2017). We note that a super-solar metallicity is indicated by the thermal pressure equilibrium solution for Absorber 0.4A. The metallicity for this model really follows a bimodal distribution (hence the large range extending to lower Z shown in Figure 20) with a

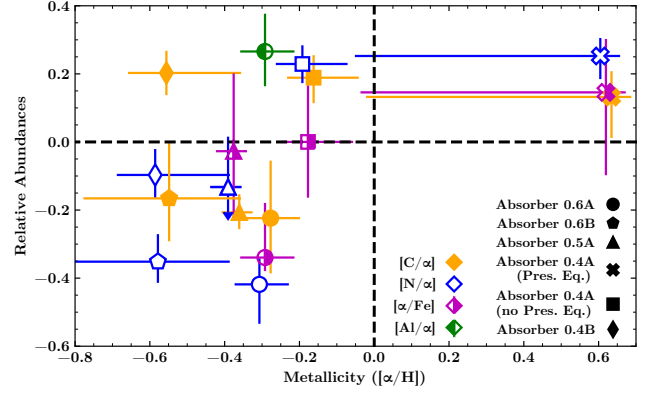


Figure 20. Modeled relative abundance ratios for C, N, Al, and Fe, with respect to the α -element abundances, vs. the modeled overall metallicity $[\alpha/\text{H}]$. The symbols represent point estimates of the Bayesian model-averaged (see Section 4.3) parameter distributions (i.e., marginal posterior densities). α/Fe is not shown for Absorbers 0.4B and 0.6B, where it is poorly constrained. Likewise, Absorber 0.6A is the only absorber where Al/α could be constrained. For clarity, we applied a -0.015 dex offset on the horizontal (metallicity) axis to the blue points ($[\text{N}/\alpha]$) and a +0.015 dex offset to the metallicities of the orange points ($[\text{C}/\alpha]$).

weaker secondary peak only slightly more enriched than the non-TPE solution for Absorber 0.4A (the square symbols). While we cannot rule out such a high metallicity arising in particular instances, the models present acceptable solutions that are less super-solar.

Our modeled metallicities also lie at the high end of the gas phase O/H metallicity gradient found in the nebular emission of higher-mass ($M_{\star} > 10^{10} M_{\odot}$) local galaxies (e.g., Belfiore et al. 2017). This correlation could suggest that the absorbers formed from enriched gas pushed out (e.g., via stellar/supernova winds or AGN outbursts) of the central regions of such galaxies into the CGM. That picture becomes less clear, however, when our relative abundance constraints are considered. Coincident with supersolar O/H, Belfiore et al. (2017) also found supersolar N/O ratios, in contrast with our modeling that requires subsolar N/O in all but one case. One possibility is that at the redshifts of our absorbers ($z=0.41760, 0.68606$, corresponding to look-back times of ~ 5 Gyr) the secondary production of nitrogen through the CNO-cycle has not yet had time to enrich the gas. Yet, the gas phase N/O ratio shows no evidence of evolving since $z \approx 1-2$ in the limited number of galaxies analyzed by Kojima et al. (2017) and in a stack of galaxies by Masters et al. (2014), for example.

Iron is another element with a delayed nucleosynthesis relative to the α -elements that provides an interesting point of comparison with three of our absorbers. Lower metallicity halo stars and globular clusters in the Milky Way offer well-known examples of enhanced α/Fe (e.g., McWilliam 1997) at low metallicities from the generation of stars formed before Type Ia supernovae had time to boost the iron abundances. Near solar metallicities, however, the α/Fe ratio is observed to decrease to solar levels, which is more consistent with the near-solar α/Fe ratios we find for Absorbers 0.4A and 0.5A. Surprisingly, our models prefer a sub-solar α/Fe ratio for Absorber 0.6A, i.e., it requires an excess of iron relative to what

is expected. In some models of galactic outflows, the hot outflowing metals are found to separate from the bulk of the ISM mass (Mac Low & Ferrara 1999), so it is possible that galaxy halos contain pockets of metal-enriched gas with abundance patterns connected to particular metal sources. A larger sample of CGM abundance ratios will be crucial to further clarify the possible progenitors of these CGM absorbers.

Absorber 0.6A also exhibits a robust Al/α overabundance relative to solar. The introduction of an Al/α parameter was necessary to obtain a satisfactory model for this absorber. Although uncommon, some enhanced Al/α abundances in QSO absorbers have been noted in the literature. For example, Ganguly et al. (1998) reported $[\text{Al}/\alpha] \approx 0.3$ at $-0.55 \leq [Z/Z_\odot] \leq -0.32$ for an absorber with similar $N(\text{H I})$ as for the systems in this paper but at $z=1.94$, and Richter et al. (2005) found $[\text{Al}/\alpha] \approx 0.4$ for a low-metallicity sub-DLA at $z=2.187$. Several other examples of enhanced aluminum abundances, predominantly in high redshift, low-metallicity, and/or high- $N(\text{H I})$ QSO absorbers, are noted in Prochaska & Wolfe (2002) and Glidden et al. (2016). These papers offer several speculations regarding the origin of non-solar aluminum relative abundances. The possibility of producing Al in novae stands out as the most consistent with our ejecta interpretation of the other observed relative abundance trends above, though these references also discuss the possibility that inferred non-solar aluminum relative abundances might originate in part owing to incorrect atomic physics (e.g., ionization and recombination rates).

Circumgalactic dust has the potential to confuse relative abundance studies since it differentially depletes metals from the gas phase. While evidence for circumgalactic dust comes from several different lines of inquiry (e.g., Ménard et al. 2010; Peek et al. 2015; De Cia et al. 2016; Frye et al. 2019), it is often neglected in analyses of optically thin CGM absorbers where exposure to the UV background is thought to suppress the amount of dust in the gas. In support of this assumption, Fumagalli et al. (2016) found that the physical conditions of a large sample of Lyman Limit Systems (LLSs) were minimally impacted by a model incorporating a dust abundance scaling based on Jenkins (2009), though it should be noted that most of these absorption systems are located at redshifts $z \geq 2.5$ and thus could have lower metallicities and lower dust contents. While we did not include a dust-specific relative abundance parameterization in our models, we have examined the abundances of species that are expected to be less depleted by dust (e.g., O and S) versus those that have a greater propensity to deplete (e.g., Si, Mg, and Fe), and we did not find any compelling indications of dust depletion patterns (see, e.g., Jenkins 2009). Indeed, the enhanced Al/α ratio discussed above is also in the wrong direction for a dusty absorber, since aluminum is expected to deplete strongly onto dust grains (Jenkins 2004).

6 SUMMARY

In this work, we developed a Bayesian multiphase ionization modeling scheme with the goal of accounting for all the observed ionization stages simultaneously at the maximum kinematic resolution possible. To explore the models' dependency on weakly constrained variables, we introduced several new or seldom encountered parameters governing 1) the number

of model phases per absorber, 2) the shape of the ionizing radiation field incident on the absorbing gas, and 3) relative abundance ratios of non α -elements in the same gas. We treated each velocity component (referred to throughout as independent “absorbers”) in an absorption system as being produced by an unresolved composition of gas from multiple different physical states, i.e., “phases”. To robustly navigate the enlarged parameter space in search of optimal solutions to our multiphase models, we employed an MCMC sampling procedure and several carefully chosen priors on the model parameters. These include a novel hyperprior on the thermal pressure, which we use to interpret the probability of thermal pressure equilibrium between the individual phases of our multiphase absorbers.

To showcase our ionization modeling approach, we selected absorption line measurements and multiphase ionization models for five strong H I absorbers ($\log N_{\text{H I}} > 15$) identified in three line-rich absorption systems at intermediate redshifts ($0.4 < z < 0.7$) in quasar spectra observed as part of the CASBaH program. The good signal-to-noise, full FUV-to-optical coverage of the background quasar spectra, and the optimal redshifts of the absorbers allowed us to detect several species that are difficult or impossible to observe at low ($z \lesssim 0.4$) and high ($z \gtrsim 1.5$) redshifts. The unusual number and precision of column density measurements gleaned from these absorbers were essential for constraining our complex ionization models.

Applying multiphase ionization equilibrium models to these absorbers, we find that only 1/5 of the absorbers is consistent with a single ionization phase. The other 4/5 absorbers require *at least two* distinct ionization phases to satisfactorily model all observed ionization stages simultaneously. In fact, a rigorous model comparison using Bayes factors strongly favors three-phase models for all four multiphase absorbers. An examination of the individual model phases reveals that for many of the species we detected, including commonly observed ions such as H I, C II, C III, Si III, and Si IV, multiple model phases may contribute a non-negligible fraction toward their total ionic column densities, making it impossible to categorize many of these ions as belonging exclusively to a particular gas phase. In some models, a significant amount (as much as $\approx 20\%$ or more) of the neutral hydrogen, for example, may originate outside the coolest, densest modeled gas phase that contains $> 95\%$ of the Mg II and Fe II. Interestingly, we also find that singly- and doubly-ionized species often *do not* originate in the same gas phase; ions such as C III, N III, and O III may share more in common (regarding their production) with “hotter” ions such as C IV, N IV, and O IV than with low-ions like C II, N II, and O II. We also examined whether photoionization-only and fixed-temperature photo-plus collisional ionization models (i.e., collisional ionization models where the incident ionizing radiation field is not ignored) are capable of satisfactorily reproducing the measured column densities of all detected species simultaneously. O VI, which we detected in all our multiphase absorbers, is readily accounted for by collisional ionization models. We also found photoionization models that could produce O VI in two of the absorbers, though not for the two remaining absorbers where we detect strong high-ion column densities, including Ne VIII.

We discussed several physical properties of our IGM/CGM absorbers, including thermal gas pressures, cloud sizes, deviations from solar metallicity and relative abundances (for C,

N, Al, and Fe), and implications for the UV ionizing radiation field impinging on these systems. For all four multiphase absorbers we found model configurations that can be interpreted as having multiphase pressures consistent with thermal pressure equilibrium, though for two absorbers we also found successful models with inter-phase thermal pressures severely out of equilibrium with each other. We observed that the key factor for building thermal pressure equilibrium models is the dynamic range of the modeled temperature. Photoionization models simply cannot achieve the higher temperatures needed to balance lower gas densities of more highly ionized gas phases. In addition to the question of equilibrium, we also find that all our multiphase models exhibit low overall thermal pressures, and the set of non-pressure equilibrium models fall well below reasonable estimations of the halo pressure based on a two-phase halo analytical model for a $10^{11.5}M_{\odot}$ halo. This suggests that the absorbers either reside in relatively low-mass (sub- L^*) halos (we blindly selected our absorbers with respect to their environments) or they have significant non-thermal pressure support.

Our models also successfully produce absorbers with realistic cloud scale lengths of $\lesssim 10 - 30$ kpc after setting a prior to restrict excessively extended solutions. The highest density gas phase in each absorber is consistent with sub-kiloparsec scale lengths. We also inferred the shape of the ionizing radiation field from our models, finding evidence that suggests the presence of a local source of radiation near four of our absorbers, while the other one is consistent with photoionization by the UV background ionizing flux. Finally, we derive metallicities greater than 25% solar in all five absorbers and find evidence of significant deviations from solar relative abundances, particularly for N/α , in all five. These case studies serve as a preview of our larger, ongoing CASBaH investigation of physical conditions in IGM/CGM absorbers that will utilize the multiphase ionization modeling scheme developed herein.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge Martin Weinberg for many helpful and stimulating discussions about statistical methodologies and the circumgalactic medium. This research is based on observations with the NASA/ESA Hubble Space Telescope, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. Financial support for HST programs HST-GO-11741 and HST-GO-13846 was provided through grants from the Space Telescope Science Institute under NASA contract NAS5-26555. N. Katz acknowledges support from NASA grant 80NSSC18K1016. Some data presented herein were obtained at the W. M. Keck Observatory, which is operated as a scientific partnership between the California Institute of Technology, the University of California, and the National Aeronautics and Space Administration. The Keck Observatory was made possible by the generous financial support of the W. M. Keck Foundation. The authors 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. The large grid of ionization models required for

this work were precomputed in parallel using the Linux package “GNU Parallel” (Tange 2011) on resources provided by the University of Massachusetts’ Green High Performance Computing Cluster (GHPCC). Financial support for computational costs was provided in part by a University of Massachusetts Amherst Graduate School Dissertation Research Grant.

DATA AVAILABILITY

The data underlying this article are presented within the article except for several large interpolation tables, which were generated from publicly available software (CLOUDY, <https://nublado.org>), and the original UV and optical quasar spectra, which are publicly available from the Mikulski Archive for Space Telescopes <https://archive.stsci.edu/>, and from the Keck Observatory Archive <https://www2.keck.hawaii.edu/koa/public/koa.php>. The various softwares used to analyze the spectra and compute Bayesian multiphase ionization models were developed by us and are not available for download at present.

REFERENCES

- Akerman, C. J., Carigi, L., Nissen, P. E., Pettini, M., & Asplund, M. 2004, *A&A*, 414, 931
- Akima, H., 1970, *Journal of the ACM*, 17, 4, 589
- Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, *ARA&A*, 47, 481
- Battisti, A. J., Meiring, J. D., Tripp, T. M. et al. 2012, *ApJ*, 744, 93
- Belfiore, F., Maiolino, R., Tremonti, C., et al. 2017, *MNRAS*, 469, 151
- Berg, D. A., Skillman, E. D., Marble, A. R. et al. 2012, *ApJ*, 754, 98
- Bowen, D. V., Chelouche, D., Jenkins, E. B., Tripp, T. M., Pettini, M., York, D. G., & Frye, B. L. 2016, *ApJ*, 826, 50
- Brooks, S. P. & Gelman, A. 1998, *J. Comput. Graph. Statist.*, 7, 434
- Buie, II, E., Fumagalli, M., Scannapieco, E. 2020, *ApJ*, 890, 33
- Burchett, J. N., Tripp, T. M., Werk, J. K., et al. 2013, *ApJ*, 779, L17
- Burchett, J. N., Tripp, T. M., Prochaska, J. X., et al. 2015, *ApJ*, 815, 91
- Burchett, J. N., Tripp, T. M., Bordoloi, R., et al. 2016, *ApJ*, 832, 124
- Burchett, J. N., Tripp, T. M., Wang, Q. D., et al. 2018, *MNRAS*, 475, 2067
- Burchett, J. N., Tripp, T. M., Prochaska, J. X., et al. 2019, *ApJ*, 877, L20
- Charlton, J. C., Ding, J., Zonak, S. G., et al. 2003, *ApJ*, 589, 111
- Chen, H.-W., Gauthier, J.-R., Sharon, K., Johnson, S. D., Nair, P., & Liang, C. J. 2014, *MNRAS*, 438, 1435
- Chen, H.-W., Johnson, S. D., Zahedy, F. S., Rauch, M., & Mulchaey, J. S. 2017, *ApJ*, 842, L19
- Cooke, R. J., Pettini, M., & Steidel, C. C. 2017, *MNRAS*, 467, 802
- Crighton, N. H. M., Hennawi, J. F., Simcoe, R. A., et al. 2015, *MNRAS*, 446, 18
- De Cia, A., Ledoux, C., Mattsson, L., et al. 2016, *A&A*, 596, A97
- Ding, J., Charlton, J. C., Churchill, C. W., & Palma, C. 2003, *ApJ*, 590, 746
- Dixon, W. V., Sahnou, D. J., Barrett, P. E., et al. 2007, *PASP*, 119, 527

- Fabbian, D., Nissen, P. E., Asplund, M., Pettini, M., & Akerman, C. 2009, *A&A*, 500, 1143
- Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, *ApJ*, 703, 1416
- Ferland, G. J., Chatzikos, M., Guzmán, F., et al. 2017, *Rev. Mex. Astron. Astrofis.*, 53, 385
- Fielding, D., Quataert, E., McCourt, M., & Thompson, T. A. 2017, *MNRAS*, 466, 3810
- Fielding, D. B., Ostriker, E. C., Bryan, G. L., & Jermyn, A. S. 2020, preprint (arXiv:2003.08390)
- Fischer, W. J., et al. 2019, *Cosmic Origins Spectrograph Instrument Handbook*, Version 11.0 (Baltimore: STScI)
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
- Frye, B., Bowen, D. V., Tripp, T. M., et al. 2019, *ApJ*, 872, 129
- Fumagalli, M., O’Meara, J. M., & Prochaska, J. X. 2016, *MNRAS*, 455, 4100
- Fumagalli, M., Haardt, F., Theuns, T., et al. 2017, *MNRAS*, 467, 4802
- Ganguly, R., Churchill, C. W., & Charlton, J. C. 1998, *ApJ*, 498, L103
- Gelman, A. & Rubin, D. B. 1992, *Statist. Sci.*, 7, 457
- Gelman, A., Meng, X.-L. & Stern, H. 1996, *Statistica Sinica*, 6, 733
- Gelman, A., Carlin, J. B., Stern, H. S. & Rubin, D. B. 2004, *Bayesian Data Analysis*, Chapman & Hall/CRC, Boca Raton, FL
- Glidden, A., Cooper, T. J., Cooksey, K. L., Simcoe, R. A., & O’Meara, J. M. 2016, *ApJ*, 833, 270
- Gnat, O., Sternberg, A. 2007, *ApJS*, 168, 213
- Goodman, J., & Weare, J. 2010, *Commun. Appl. Math. Comput. Sci.*, 5, 65
- Green, J. C., Froning, C. S., Osterman, S., et al. 2012, *ApJ*, 744, 60
- Gronke, M. & Oh, S. P. 2020, *MNRAS*, 492, 1970
- Gronke, M. & Oh, S. P. 2020, *MNRAS*, 494, L27
- Guo, Q., White, S., Li, C., & Boylan-Kolchin, M. 2010, *MNRAS*, 404, 1111
- Haardt, F., & Madau, P. 1996, *ApJ*, 461, 20
- Haardt, F., & Madau, P. 2001, in Neumann D. M., Tran J. T. V., eds, *XXIst Moriond Astrophysics Meeting, Clusters of Galaxies and the High Redshift Universe Observed in X-rays*. Editions Frontieres, Paris
- Haardt, F., & Madau, P. 2012, *ApJ*, 746, 125
- Handley, W., Millea, M. 2019, *Entropy*, 21(3), 272
- Henley, D. B., Gritton, J. A., & Shelton, R. L. 2017, *ApJ*, 837, 82
- Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. 1999, *Stat. Sci.*, 14, 382
- Hussain, T., Khaire, V., Srianand, R., Muzahid, S., & Pathak, A. 2017, *MNRAS*, 466, 3133
- Jenkins, E. B. 1996, *ApJ*, 471, 292
- Jenkins, E. B. 2004, in McWilliam A., Rauch M., eds, *Origin and Evolution of the Elements*, 336, preprint (arXiv:astro-ph/0309651)
- Jenkins, E. B. 2009, *ApJ*, 700, 1299
- Jeffreys, H. 1961, *Theory of Probability* 2nd ed., Oxford Univ. Press
- Ji, S., Chan, T. K., Hummels, C. B. et al. 2019, preprint (arXiv:1909.00003)
- Kass, R. E., & Raftery, A. E. 1995, *Journal of the American Statistical Association*, 90, 773
- Keeney, B. A., Stocke, J. T., Rosenberg, J. L. et al. 2013, *ApJ*, 765, 27
- Kereš, D., Katz, N., Weinberg D. H., Davé, R., 2005, *MNRAS*, 363, 2
- Kereš, D., Katz, N., Fardal, M., Davé, R., Weinberg, D. H., 2009, *MNRAS*, 395, 160
- Khaire, V., & Srianand, R. 2015, *MNRAS*, 451, L30
- Khaire, V., & Srianand, R. 2019, *MNRAS*, 484, 4174
- Khaire, V., Walther, M., Hennawi, J. F. et al. 2019, *MNRAS*, 486, 769
- Kojima, T., Ouchi, M., Nakajima, K., et al. 2017, *PASJ*, 69, 44
- Kollmeier, J. A., Weinberg, D. H., Oppenheimer, B. D., et al. 2014, *ApJ*, 789, L32
- Kriss, G. A. 2011, *COS Instrument Science Report 2011-01(v1)*, 17 pages,
- Lehner, N., Howk, J. C., Tripp, T. M., et al. 2013, *ApJ*, 770, 138
- Lepage, G. P. 1978, *J. Comput. Phys.* 27, 192
- Liang, C. J., & Remming, I. 2019, *MNRAS*, 491, 5056
- Lochhaas, C., Bryan, G. L., Li, Y., Li, M., & Fielding, D. 2020, *MNRAS*, 493, 1461
- Lodders, K. 2003, *ApJ*, 591, 1220
- Mac Low, M.-M., & Ferrara, A. 1999, *ApJ*, 513, 142
- Maller, A. H., & Bullock, J. S. 2004, *MNRAS*, 355, 694
- Masters, D., McCarthy, P., Siana, B., et al. 2014, *ApJ*, 785, 153
- McCourt, M., Oh, S. P., O’Leary, R., & Madigan, A.-M. 2018, *MNRAS*, 473, 5407
- McQuinn, M., & Werk, J. K. 2018, *ApJ*, 852, 33
- McWilliam, A. 1997, *ARA&A*, 35, 503
- Meiring, J. D., Tripp, T. M., Prochaska, J. X., et al. 2011, *ApJ*, 732, 35
- Meiring, J. D., Tripp, T. M., Werk, J. K., et al. 2013, *ApJ*, 767, 49
- Ménard, B., Scranton, R., Fukugita, M., & Richards, G. 2010, *MNRAS*, 405, 1025
- Murray, N., Ménard, B., & Thompson, T. A. 2011, *ApJ*, 735, 66
- Muzahid, S., Srianand, R., Bergeron, J., & Petitjean, P. 2012, *MNRAS*, 421, 446
- Muzahid, S., Srianand, R., Arav, N., Savage, B. D., & Narayanan, A. 2013, *MNRAS*, 431, 2885
- Muzahid, S., Srianand, R., Charlton, J., & Eracleous, M. 2016, *MNRAS*, 457, 2665
- O’Meara, J. M., Lehner, N., Howk, J. C., et al. 2015, *AJ*, 150, 111
- Oppenheimer, B. D., & Schaye, J. 2013, *MNRAS*, 434, 1063
- Oppenheimer, B. D., Segers, M., Schaye, J., Richings, A. J., & Crain, R. A. 2018, *MNRAS*, 474, 4740
- Pachat, S., Narayanan, A., Khaire, V., et al. 2017, *MNRAS*, 471, 792
- Peek, J. E. G., Ménard, B., & Corrales, L. 2015, *ApJ*, 813, 7
- Penprase, B. E., Prochaska, J. X., Sargent, W. L. W., Toro-Martinez, I. & Beeler, D. J. 2010, *ApJ*, 721, 1
- Penton, S. V., Stocke, J. T., & Shull, J. M. 2004, *ApJS*, 152, 29
- Pettini, M., Zych, B. J., Steidel, C. C., & Chaffee, F. H. 2008, *MNRAS*, 385, 2011
- Prochaska, J. X., Burchett, J. N., Tripp, T. M., et al. 2019, *ApJS*, 243, 24
- Prochaska, J. X., Naumov, S. O., Carney, B. W., McWilliam, A., & Wolfe, A. M. 2000, *AJ*, 120, 2513
- Prochaska, J. X., Werk, J. K., Worseck, G., et al. 2017, *ApJ*, 837, 169
- Prochaska, J. X., & Wolfe, A. M. 2002, *ApJ*, 566, 68
- Qu, Z., & Bregman, J. N. 2016, *ApJ*, 832, 189
- Rafelski, M., et al. 2018, *COS Data Handbook*, Version 4.0, (Baltimore: STScI)
- Rauch, M., Haehnelt, M. G., & Steinmetz, M. 1997, *ApJ*, 481, 601
- Rauch, M., Sargent, W. L. W., & Barlow, T. A. 2001, *ApJ*, 554, 823
- Ribaldo, J., Lehner, N., Howk, J. C., et al. 2011, *ApJ*, 743, 207
- Richter, P., Ledoux, C., Petitjean, P., & Bergeron, J. 2005, *A&A*, 440, 819
- Richter, P., Savage, B. D., Sembach, K. R., & Tripp, T. M. 2006, *A&A*, 445, 827
- Rosenwasser, B., Muzahid, S., Charlton, J. C., et al. 2018, *MNRAS*, 476, 2258
- Rubin, K. H. R., Diamond-Stanic, A. M., Coil, A. L., Crighton, N. H. M., & Stewart, K. R. 2018, *ApJ*, 868, 142
- Rudie, G. C., Steidel, C. C., Pettini, M., Trainor, R. F., Strom, A.

L., Hummels, C. B., Reddy, N. A., & Shapley, A. E. 2019, *ApJ*, 885, 61

Savage, B. D., & Sembach, K. R. 1991, *ApJ*, 379, 245

Savage, B. D., Lehner, N., Wakker, B. P., Sembach, K. R., & Tripp, T. M. 2005, *ApJ*, 626, 776

Savage, B. D., Narayanan, A., Wakker, B. P., et al. 2010, *ApJ*, 719, 1526

Savage, B. D., Kim, T.-S., Wakker, B. P., et al. 2014, *ApJS*, 212, 8

Scott J., Kriss G., Brotherton M., Green R., Hutchings J., Shull J., Zheng W. 2004, in Richards G. T., Hall P. B., eds, ASP Conf. Ser. Vol. 311, A Composite Extreme Ultraviolet QSO Spectrum from FUSE. Astron. Soc. Pac., San Francisco, p. 31

Sembach, K. R., & Savage, B. D. 1992, *ApJS*, 83, 147

Sembach, K. R., Tripp, T. M., Savage, B. D., & Richter, P. 2004, *ApJS*, 155, 351

Shull, J. M., Stocke, J. T., & Penton, S. 1996, *AJ*, 111, 72

Shull, J. M., Moloney, J., Danforth, C. W., & Tilton, E. M. 2015, *ApJ*, 811, 3

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

Stevens, M. L., Shull, J. M., Danforth, C. W., & Tilton, E. M. 2014, *ApJ*, 794, 75

Tange, O., 2011, *The USENIX Magazine*, 36, 42

Telfer, R. C., Zheng, W., Kriss, G. A., & Davidsen, A. F. 2002, *ApJ*, 565, 773

Thompson, T. A., Quataert, E., Zhang, D., & Weinberg, D. H. 2016, *MNRAS*, 455, 1830

Tripp, T. 2013, preprint (arXiv:1303.0043)

Tripp, T. M., Savage, B. D., & Jenkins, E. B. 2000, *ApJ*, 534, L1

Tripp, T. M., Giroux, M. L., Stocke, J. T., Tumlinson, J., & Oegerle, W. R. 2001, *ApJ*, 563, 724

Tripp, T. M., Aracil, B., Bowen, D. V., & Jenkins, E. B. 2006, *ApJ*, 643, L77

Tripp, T. M., Sembach, K. R., Bowen, D. V., et al. 2008, *ApJS*, 177, 39

Tripp, T. M., Meiring, J. D., Prochaska, J. X., et al. 2011, *Science*, 334, 952

Tumlinson, J., Thom, C., Werk, J. K., et al. 2011, *Science*, 334, 948

Tumlinson, J., Thom, C., Werk, J. K., et al. 2013, *ApJ*, 777, 59

Tumlinson, J., Peebles, M. S., & Werk, J. K. 2017, *ARA&A*, 55, 389

Upton Sanderbeck, P. R., McQuinn, M., D’Aloisio, A., & Werk, J. K. 2018, *ApJ*, 869, 159

van Zee, L., & Haynes, M. P. 2006, *ApJ*, 636, 214

Verner, D. A., Barthel, P. D., & Tytler, D. 1994, *A&AS*, 108, 287

Verner, D. A., Ferland, G. J., Korista, K. T., & Yakolev, D. G., 1996a, *ApJ*, 465, 487

Verner, D. A., Verner, E. M., & Ferland, G. J. 1996b, *Atomic Data and Nuclear Data Tables*, 64, 1

Vila Costa, M. B., & Edmunds, M. G. 1993, *MNRAS*, 265, 199

Vogt, S. S., Allen, S. L., Bigelow, B. C., et al. 1994, *Proc. SPIE Instrum. in Astron. VIII*, Vol. 2198, 362

Weinberg, D. H., Miralda-Escudé, J., Hernquist, L., & Katz, N. 1997, *ApJ*, 490, 564

Weinberg, M. D. 2012, *Bayesian Analysis*, 7, 737

Weinberg, M. D. 2013, *MNRAS*, 434, 1736

Weinberg, M. D., Yoon, I., & Katz, N. 2013, preprint (arXiv:1301.3156)

Werk, J. K., Prochaska, J. X., Tumlinson, J., et al. 2014, *ApJ*, 792, 8

Werk, J. K., Prochaska, J. X., Cantalupo, S., et al. 2016, *ApJ*, 833, 54

Werk, J. K., Rubin, K. H. R., Bish, H. V., et al. 2019, *ApJ*, 887, 29

Wolfe, A. M., Gawiser, E. & Prochaska, J. X. 2005, *ARA&A*, 43, 861

Woodgate, B. E., Kimble, R. A., Bowers, C. W., et al. 1998, *PASP*, 110, 1183

Wotta, C. B., Lehner, N., Howk, J. C., O’Meara, J. M., Oppenheimer, B. D., & Cooksey, K. L. 2019, *ApJ*, 872, 81

Wotta, C. B., Lehner, N., Howk, J. C., O’Meara, J. M., & Prochaska, J. X. 2016, *ApJ*, 831, 95

APPENDIX A: BAYESIAN MODEL DETAILS

A1 The Likelihood Function

The likelihood function (Equation A6) for our multiphase model assumes that column densities measured from the spectra are random variates of a zero-truncated normal distribution (ZTN). What we refer to throughout this paper as the zero-truncated normal (ZTN) distribution is really a simplification of the truncated normal distribution, which is defined as:

$$p_{\text{Trunc. Normal}}(x, \mu, \sigma, a, b) = \frac{\phi\left(\frac{x-\mu}{\sigma}\right)}{\sigma \left(\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right) \right)} \quad (\text{A1})$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ reference the standard normal probability density and cumulative density functions (respectively), which we provide here for completeness:

$$\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} \quad (\text{A2})$$

and

$$\Phi(x) = \frac{1}{2} \left(1 + \operatorname{erf}\left(\frac{x}{\sqrt{2}}\right) \right). \quad (\text{A3})$$

The ZTN probability density corresponds to setting $a, b = 0, -\infty$ in Equation A1, which simplifies to:

$$f_{\text{ZTN}}(x, \mu, \sigma) = \frac{\phi\left(\frac{x-\mu}{\sigma}\right)}{\sigma \left(1 - \Phi\left(-\frac{\mu}{\sigma}\right) \right)} \quad (\text{A4})$$

and the accompanying cumulative density function is:

$$F_{\text{ZTN}}(x, \mu, \sigma) = \frac{\Phi\left(\frac{x-\mu}{\sigma}\right) - \Phi\left(-\frac{\mu}{\sigma}\right)}{1 - \Phi\left(-\frac{\mu}{\sigma}\right)} \quad (\text{A5})$$

Our adoption of the ZTN as our column density sampling distribution was motivated by simulations of absorption features of different strengths in simulated counts spectra at various signal-to-noise ratios. For each test case, we ran 1000 different realizations of the spectrum, each time fitting a Voigt profile (with fixed velocity and b -value constrained between 10-25 km s⁻¹) to the absorption feature (even if non-detected) using χ^2 minimization. The histogram of best fitting column densities matches best with a ZTN distribution, even when the mean of the ZTN was fixed to the “true” (simulated) column density value. Note that the assumption of ZTN distributed column densities may break down in more complex scenarios than those considered above, such as multiple adjacent absorption lines (whether from the same system or a contaminating one), or for Voigt-profile fitting measurements with shared (and hence correlated) parameters. A proper treatment of such covariates has not been considered in the literature, however, and would be beyond the scope of this paper. For the present, using the simpler (truncated normal)

likelihood distribution described above should suffice as an approximation.

The full likelihood distribution is the product of independent ZTNs, one for each species-absorber pair (denoted by subscripts *ion, c*). The width of an individual ZTN (for species *ion*) is set to the measured uncertainty in the column density measurements, i.e., $\sigma_{\text{ZTN}} = \sigma_{\text{ion},c}^{\text{meas}}$. The mean²⁶ of each ZTN sampling distribution is set by the model-predicted column density for a given set of model parameters, i.e., $\mu_{\text{ZTN}} = N_{\text{ion},c}^{\text{model}}$. Equation A6 gives the likelihood of the observed column densities as the product of ZTN probabilities. The specific form of the ZTN for each term in the product depends on whether the measured column densities belong to the set of *Detections*, *Upper Limits*, or *Lower Limits* (i.e., *{D}*, *{UL}*, or *{LL}* in Equation A6). Since the column density measurement (kinematic) resolution often differs for various species (for the reasons discussed above), we introduce the concept of *component groups* to handle measurements that span one or more modeled absorbers (i.e., velocity components). These component groups (CG) are exclusive, meaning they encompass adjacent, but non-overlapping portions of the absorption features. A CG may contain anywhere from one to all of the measured column density constraints used in the model. CGs are in essence bookkeeping tools, enabling us to compare a lower (kinematic) resolution measurement to the more finely resolved models, by keeping track of which modeled absorbers (i.e., higher resolution gridpoints of the model) are encompassed by the CG.

While the likelihood of detection is given by the ZTN probability density function, for upper limits the likelihood is the *integrated* tail probability $P(N_{\text{ZTN}} \leq N_{\text{ion},c}^{\text{limit}})$, which is obtained by evaluating the ZTN cumulative distribution function (CDF, given by Equation A5) at the measured limit $N_{\text{ion},c}^{\text{limit}}$. Lower limits are analogous, but correspond to upper tail probabilities, i.e., $P(N_{\text{ZTN}} \geq N_{\text{ion},c}^{\text{limit}})$. Thus the full likelihood distribution is given by (with reference to Equations A4 and A5):

$$L(\mathbf{N}|\boldsymbol{\theta}) = \prod_{\text{ion}} \prod_{c=1}^{N_{\text{CG}}} \begin{cases} f_{\text{ZTN}}(N_{\text{ion},c}^{\text{meas}}, N_{\text{ion},c}^{\text{model}}, \sigma_{\text{ion},c}^{\text{meas}}), & \text{if } \text{ion} \in \{D\}^c \\ F_{\text{ZTN}}(N_{\text{ion},c}^{\text{limit}}, N_{\text{ion},c}^{\text{model}}, \sigma_{\text{ion},c}^{\text{limit}}), & \text{if } \text{ion} \in \{UL\}^c \\ 1 - F_{\text{ZTN}}(N_{\text{ion},c}^{\text{limit}}, N_{\text{ion},c}^{\text{model}}, \sigma_{\text{ion},c}^{\text{limit}}), & \text{if } \text{ion} \in \{LL\}^c \end{cases} \quad (\text{A6})$$

$$\text{where } N_{\text{ion},c}^{\text{model}} = \sum_{i=1}^{N_{\text{abs}}(c)} \sum_{p=1}^{N_{\text{phases}}(\text{ion}^i)} N_{\text{ion}}^{\text{CLOUDY}}(\boldsymbol{\theta}_p^i) \left(\frac{N_{\text{H}}^{i,p}}{N_{\text{H}}^0} \right),$$

and

$$\boldsymbol{\theta}_p^i = \{\log n_{\text{H}}^{i,p}, \log T^{i,p}, \alpha_{\text{UV}}^i, \beta_{\text{He II}}^i, \log Z^i, [C/\alpha]^i, [N/\alpha]^i, [\alpha/\text{Fe}]^i\}.$$

Here, N_{ions} , N_{CG} , N_{abs} , and N_{phases} , refer, respectively, to the number of unique species used to constrain the model, the number of component groups used to represent the measurements for these species, the number of absorbers being jointly modeled (e.g., two absorbers at the same systemic redshift), and the number of phases used by the model for each absorber. The modeled column densities $N_{\text{ion},c}^{\text{model}}$ actually repre-

sent the sum over column densities predicted at a more granular level (i.e., by individual absorbers and phases). These more granular column densities $N_{\text{ion}}^{\text{CLOUDY}}(\boldsymbol{\theta}_p^i)$ are obtained by interpolating over our grid (Table 5) of CLOUDY ionization model evaluations, with free grid parameters $\boldsymbol{\theta}_p^i$. Note that $\log T^{i,p}$ is only included for the P+C model. Finally, the total hydrogen ratio adjusts the gas content of the absorber with a normalization for the zeropoint of the interpolation grid.

A2 Prior Probability Density Functions

Section 3 includes a thorough qualitative description of the priors incorporated in our ionization modeling. Here we present the underlying formulae, using bold symbols throughout to represent parameter sets. The full prior distribution can be written as:

$$P(\boldsymbol{\theta}) = \hat{P}(\boldsymbol{\theta}) P(\sigma_{\ln P, \text{true}}^2) \frac{P(s_{\ln P}^2 | \sigma_{\ln P, \text{true}}^2)}{P(s_{\ln P}^2 | \hat{P}(\boldsymbol{\theta}), P(\sigma_{\ln P, \text{true}}^2))} \quad (\text{A7})$$

where the first term $\hat{P}(\boldsymbol{\theta})$ represents the joint prior distribution over the primary model parameters and the remaining terms in Equation A7 describe our thermal pressure prior, including a term for the hyperprior $P(\sigma_{\ln P, \text{true}}^2)$. Equation 15 implicitly defines $P(s_{\ln P}^2 | \sigma_{\ln P, \text{true}}^2)$, the prior on the sample variance of the inter-phase pressures (a derived parameter of the model). Following Handley & Millea (2019), as explained in Section 3, we also need to divide by $P(s_{\ln P}^2 | \hat{P}(\boldsymbol{\theta}), P(\sigma_{\ln P, \text{true}}^2))$, that is, the prior induced on $s_{\ln P}^2 | \sigma_{\ln P, \text{true}}^2$ by the explicit priors $\hat{P}(\boldsymbol{\theta})$.

The prior on the primary model parameters is a joint distribution

$$\hat{P}(\boldsymbol{\theta}) = P(\alpha_{\text{UV}}, \beta_{\text{He II}}, \mathbf{X}/\alpha, \alpha/\mathbf{H}, \log T, \log n_{\text{H}}, \log N_{\text{H}}), \quad (\text{A8})$$

which in our application simplifies to

$$\hat{P}(\boldsymbol{\theta}) = P(\alpha_{\text{UV}}) P(\beta_{\text{He II}}) P(\alpha/\mathbf{H}) P(\mathbf{X}/\alpha) P(\log T) P(\log n_{\text{H}}) \\ \times P(\log N_{\text{H}} | \log \tilde{n}_{\text{H}} = \log n_{\text{H}} + \log f(\alpha_{\text{UV}}, \beta_{\text{He II}})) \quad (\text{A9})$$

though we can imagine inserting various additional conditional dependencies given sufficiently strong prior information, e.g., future prior constraints on relative abundance ratios might be conditional on the metallicity. Note that the conditional prior for N_{H} has been formulated as a prior on the cloudsize, which requires that it be cast in terms of the adjusted density $\log \tilde{n}_{\text{H}}$ using the logarithmic form of Equation 9. The specific expressions for each of the terms in Equation A9 are given by Equations A10–A16, some of which make reference to the standard probability distributions listed in Equations A1–A3 and A17–A18. Ideally we would model every discrete velocity component as a distinct absorber, but in a crowded absorption system it is not always possible to resolve the individual components in every observed species. In these instances, two or more absorbers must be modeled together with joint constraints on any unresolved species. In the following, the index *i* extends over the number N_{abs} of distinct absorbers modeled together. For the ionizing UV spectral shape parameters, the index *u* extends over the set of distinct UV parameter pairs (with cardinality N_{UV}). Note that $N_{\text{UV}} \leq N_{\text{abs}}$ since we allow for the possibility of sharing UV parameters between models for one or more absorbers (at the same systemic redshift). Finally, the index *p* extends over the

²⁶ To be precise, we use μ_{ZTN} to denote the 1st moment of the *non-truncated* normal distribution.

set of phases in the model for the i th absorber. Thus,

$$p(\alpha_{UV}) = \prod_u^{N_{UV}} p_{\text{Trunc. Normal}}(\alpha_{UV}^u, \alpha_{UV}^{\text{HM12}}(z), 0.7, -2.5, -0.28) \quad (\text{A10})$$

where $\alpha_{UV}^{\text{HM12}}(z)$ is the fiducial (Haardt & Madau 2012) ionizing UV spectral index measured at redshift z ;

$$p(\beta_{\text{He II}}) = \prod_u^{N_{UV}} p_{\text{Uniform}}(\beta_{\text{He II}}^u, 0, 3) \quad (\text{A11})$$

and

$$p(X/\alpha) = \prod_{\text{element} \in X}^{\text{card}(X)} \prod_i^{N_{\text{abs}}(\text{element})} p_{\text{Uniform}}([element/\alpha]_i, a_{ele}, b_{ele}) \quad (\text{A12})$$

where X includes the elements C, N, and Fe with $(a_{ele}, b_{ele}) = (-0.8, 0.8)$, $(-1.2, 0.4)$, $(-1.2, 0.4)$ for C/ α , N/ α , and Fe/ α , respectively. As explained in the main text, we make an exception if $element = \text{Al}$, replacing the uniform probability density with $p_{\text{Trunc. Normal}}([element/\alpha]_i, 0, 0.2, -0.5, 0.5)$;

$$p(\alpha/\mathbf{H}) = \prod_i^{N_{\text{abs}}} p_{\text{Uniform}}([\alpha/\mathbf{H}]_i, -3.0, 0.7) \quad (\text{A13})$$

where $[\alpha/\mathbf{H}]$ represents the $[Z/\mathbf{H}]$ metallicity parameter varied in our grid of CLOUDY models, which is better understood as $[\alpha/\mathbf{H}]$ after we allow the significant non- α elements to vary freely

$$p(\log T) = \prod_i^{N_{\text{abs}}} \prod_p^{N_{\text{phases}, P+C}^i} p_{\text{Uniform}}(\log T_{i,p}, 4.0, 6.4) \quad (\text{A14})$$

where p is restricted to those model phases (if any) where P+C is used (since $\log T$ is not a parameter of PIE-only phase models)

$$p(\log n_{\mathbf{H}}) = \prod_i^{N_{\text{abs}}} \prod_p^{N_{\text{phases}}^i} (N_{\text{phases}}^i + 1)! \cdot 1\{n_{\mathbf{H}}^{\min} \leq n_{\mathbf{H}}^{i,p} < n_{\mathbf{H}}^{i,p-1} \leq n_{\mathbf{H}}^{\max}\} \quad (\text{A15})$$

where the indicator function only accepts solutions with one-to-one ordering between the density and index of individual phases in each modeled absorber. Limits on $n_{\mathbf{H}}^{\min}$ and $n_{\mathbf{H}}^{\max}$ are imposed by the boundaries of the CLOUDY ionization model grid (see Table 5). The normalizing constant $(N_{\text{phases}}^i + 1)!$ is simplest to understand if one considers the line segment between $n_{\mathbf{H}}^{\min}$ and $n_{\mathbf{H}}^{\max}$ as being split into $N_{\text{phases}}^i + 1$ smaller segments by the densities $n_{\mathbf{H}}^i$, which correspond to the N_{phases}^i phases specified by the model for the i th absorber. These line segments can be arranged in $(N_{\text{phases}}^i + 1)!$ possible ways, only one of which gives non-zero probability in Equation A15.

$$p(\log N_{\mathbf{H}} | \log \tilde{n}_{\mathbf{H}}) = \prod_i^{N_{\text{abs}}} \prod_p^{N_{\text{phases}}^i} p_{\text{Gamma}}(\log N_{\mathbf{H}}^{i,p}, 0.2, \tilde{n}_{\mathbf{H}}^{i,p} \cdot 15 \text{ kpc}) \quad (\text{A16})$$

where we use the adjusted density given by Equation 9. The following two equations list the standard Uniform and Gamma probability distributions invoked above for completeness. Note that we define the Truncated Normal distribution, which is also invoked above, in Equation A1.

$$p_{\text{Uniform}}(x, a, b) = \frac{1}{b-a} \text{ for } a \leq x < b \quad (\text{A17})$$

and

$$p_{\text{Gamma}}(x, k, \theta) = \frac{x^{k-1} e^{-x/\theta}}{\theta^k \Gamma(k)} \text{ for } x, k, \theta > 0. \quad (\text{A18})$$

A3 Numerical Estimation of the Marginal Likelihood

Weinberg (2013) and Weinberg et al. (2013) describe a novel approach for estimating the marginal likelihood from samples of the posterior distribution. Noting that posterior volumes with the smallest density of samples dominate the uncertainty in the marginal likelihood integral, they show how the variance can be reduced by restricting the marginal likelihood estimation to a small volume centered on the posterior mode. The marginal likelihood can be written as

$$P(N|M) \int_{\Omega_s} dP(\theta|N) = \int_{\Omega_s} P(\theta|M) P(N|\theta, M) \quad (\text{A19})$$

where $\Omega_s \subset \Omega$ and Ω represent the full set of posterior samples. Ω_s is carefully chosen to minimize the variance in the numerical evaluation of the two integrals. The integral on the left-hand side is just the fraction Ω_s/Ω if evaluated via Monte Carlo sampling. The right-hand side can be evaluated using the tree-based tessellation VTA algorithm proposed by Weinberg (2012) or by naive Monte Carlo resampling of the volume defined by Ω_s as suggested by Weinberg et al. (2013).

We make two adjustments to the algorithm. First, we use a custom iterative procedure to locate a hyperrectangle that circumscribes the posterior mode. We identify samples with marginal posterior probabilities within a fraction p of the marginal posterior mode that form a solid hyperrectangle containing the mode. We iteratively adjust p until the number of posterior samples satisfying this condition equals the desired value, N_{Ω_s} . We use a value of $N_{\Omega_s} = 1000$, which offers the best tradeoff in minimizing the variance of both integrals in Equation A19. Second, we estimate Equation A19 using the vegas adaptive Monte Carlo algorithm (Lepage 1978) implemented in the VEGAS Python package²⁷. This algorithm relies on both importance sampling and stratified sampling strategies to estimate the integral. We obtain much more efficient sampling using this algorithm than with using naive Monte Carlo integration.

A4 MCMC Tests

We tested our Bayesian ionization modeling scheme on numerous mock absorbers to 1) estimate the accuracy and precision of our parameter inferences and 2) assess the sensitivity of Bayes factors for model selection. These properties depend on both the (unknown) physical parameters of the true absorbers and on the number and quality of the absorption line constraints. Performing global tests over these variables would be computationally challenging and superfluous to our purposes. Instead, we narrow our scope and test mock absorbers designed to mimic the constraints and inferred parameters of the observed (real) absorbers presented in this study.

We designed the first set of tests, i.e., for accuracy and precision, as follows. We begin by randomly drawing ten parameter vectors from the posterior of highest ranked model

²⁷ Version 3.4 <https://vegas.readthedocs.io/en/latest/index.html>

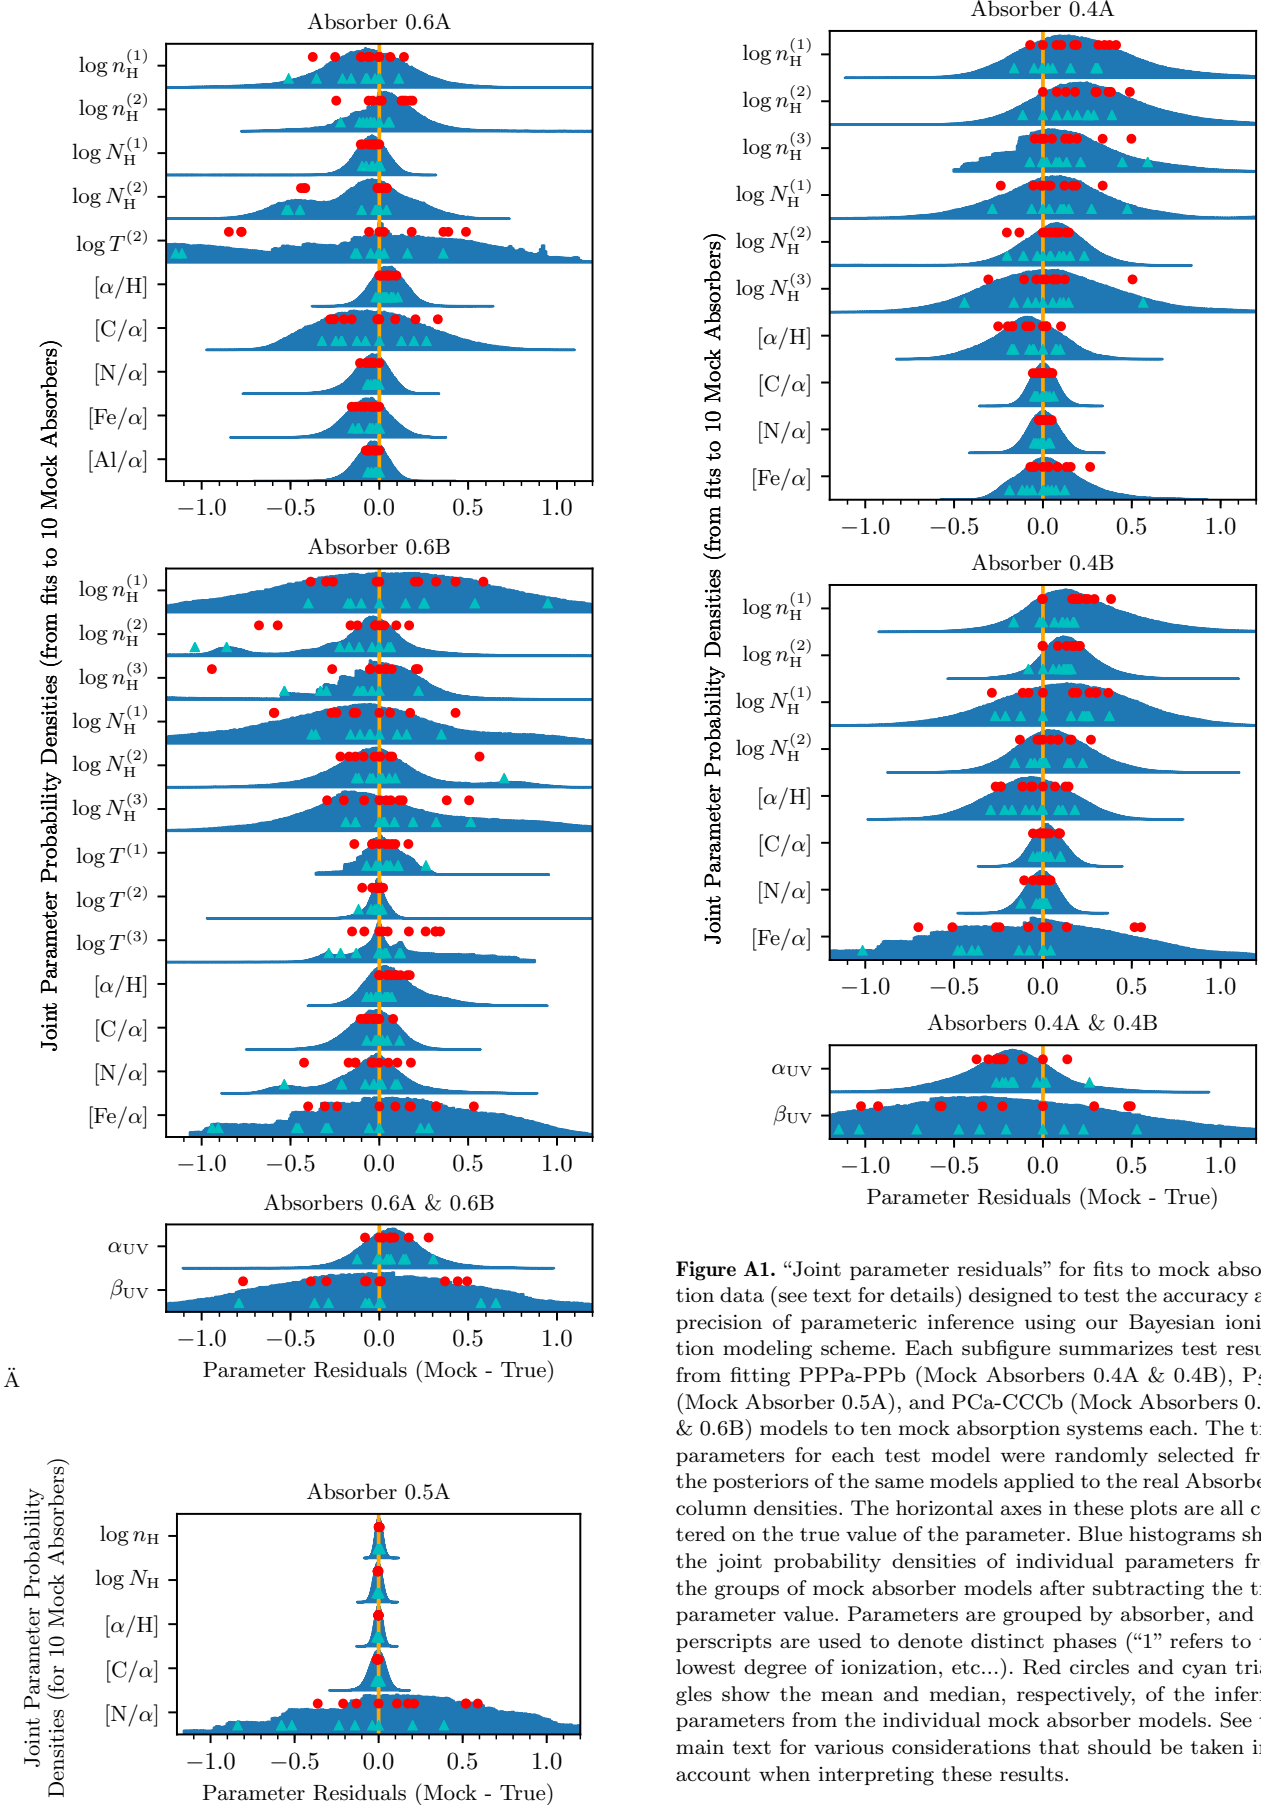


Figure A1. “Joint parameter residuals” for fits to mock absorption data (see text for details) designed to test the accuracy and precision of parametric inference using our Bayesian ionization modeling scheme. Each subfigure summarizes test results from fitting PPPa-PPb (Mock Absorbers 0.4A & 0.4B), P_{5par} (Mock Absorber 0.5A), and PCa-CCCb (Mock Absorbers 0.6A & 0.6B) models to ten mock absorption systems each. The true parameters for each test model were randomly selected from the posteriors of the same models applied to the real Absorbers’ column densities. The horizontal axes in these plots are all centered on the true value of the parameter. Blue histograms show the joint probability densities of individual parameters from the groups of mock absorber models after subtracting the true parameter value. Parameters are grouped by absorber, and superscripts are used to denote distinct phases (“1” refers to the lowest degree of ionization, etc...). Red circles and cyan triangles show the mean and median, respectively, of the inferred parameters from the individual mock absorber models. See the main text for various considerations that should be taken into account when interpreting these results.




Figure A2. A comparison of Bayes factors for fits to ten “simple” and ten “complex” mock absorption systems (see the main text for the distinction between these sets). The various models used to fit the mock absorbers are labeled along the vertical axis, with labels specifying the number and ionization type of phases (following the convention described in Section 3.3) for each component (“a” and “b”). The Bayes factor for each model-mock absorber pair is the ratio of its marginal likelihood to that of the true model. Colors and symbols redundantly refer to individual mock absorption systems. The solid vertical line represents alternative models with equal odds as the true model, while the dashed lines indicate a strong preference for or against the alternative model following the interpretation scale of Jeffreys (1961). The horizontal axis has been linearly scaled by a factor of four in the hatched region to minimize whitespace. A small vertical jitter is applied to enhance marker visibility.

(by Bayes factors) for each observed absorber and compute the corresponding model-predicted column densities. Each parameter vector and its associated column densities represents a distinct mock absorber. Drawing mock absorbers in this manner (i.e., from the posteriors of the models for real absorbers) ensures that the mock absorbers test a similar parameter space as that probed by the real ones, but with sufficient variation. We also want our mock column densities to have similar constraining power over the models as the real absorbers’ column densities. To achieve this, we 1) record column densities for the same set of species observed in the real absorbers while borrowing the same measurement classifications (detection or upper/lower limit)²⁸, and 2) apply the measurement uncertainties from the real data to the mock data.

After creating the mock absorbers, we then apply the same Bayesian ionization model that we used for the true absorbers to each mock absorber, after removing the pressure equilibrium and cloud size priors.

For each set of ten mock absorbers, we combined posterior samples of the individual models after first subtracting their known parameter vectors that we originally drew from the posteriors of real models; we refer to these as “joint parameter residuals.” Figure A1 shows histograms of these joint parameter residuals, along with the mean and median of the parameter residuals for individual mock absorbers. Using these figures, we can estimate the accuracy of the param-

eters²⁹ inferred by our models for the real absorption systems. For example, the modeled UV spectral index (α_{UV}) for Mock Absorbers 0.6A & 0.6B is overestimated by ≈ 0.08 dex on average, which can be interpreted to mean that the corresponding measurements for the real absorbers (See Figure 20) most likely overestimate the true (unknown) value by a similar amount. But Figure A1 also shows an $\approx 35\%$ chance that the true α_{UV} is underestimated (i.e., by integrating the blue histogram in the region >0), so the bias is relatively weak. Overall, the joint parameter residuals reveal relatively minor bias, in all cases $\lesssim 1\sigma$, but in some cases with rather large variances. Note, however, that the interpretation of the variance depends on the parameter. For example, the iron relative abundance parameter for Mock Absorber 0.6B has far too wide a variance to draw any useful insight from the model. This comes as no surprise, however, since iron in this absorber is only detected as an upper limit, and hence provides only a weak constraint on the model.

We also devised a second set of tests to evaluate the sensitivity of Bayes factors for selecting the true model from a set of models with various phase combinations, i.e., including both ionization mechanism (PIE-only vs. P+C) and number (of phases) per velocity component. Ideally, Bayes factors should show a preference for the true model regardless of its complexity, since Bayes factors naturally penalize both too simplistic and too overly complex models. In practice, the signal is not always strong enough to be conclusive, or it may be affected by insufficient data, bad priors, or imprecise and/or biased computation of the marginal likelihood. Hence,

²⁸ In some cases we had to extend upper and/or lower limits for individual mock absorbers so that the mock column densities would remain inside the limits.

²⁹ More precisely, we can evaluate the accuracy of *point estimates* of the parameter marginal posterior distributions.

we present these tests as an empirical calibration of Bayes factors specific to our application. For these tests we produced two sets of ten photoionized mock absorption systems, one generated from a complex model, the other from a simpler model. Specifically, the “complex” set of mock absorbers were drawn from the posterior of the PPPa-PPb model of the real Absorbers 0.4A & 0.4B in the same fashion as that described above for the accuracy and precision test (in fact the same mock absorbers were re-used). The ten “simple” mock absorbers were similarly produced, but drawn from the posterior of a lower complexity PPa-Pb model of the same real absorbers.³⁰ We then model the mock absorbers with a suite of ionization models with varying degrees of complexity, including the correct models from which they were derived. At every intersection of the mock absorbers and trial ionization models, we compute Bayes factors as the ratio of the marginal likelihood for the trial model over that of the true model.

Figure A2 presents logarithms of the Bayes factors from the second test relative to the true model. The logarithmic ratios are all essentially less than or equal to zero, meaning no model is strongly preferred over the true model. Further, we can conclusively rule out models with a lower degree of complexity (i.e., fewer phases) than the true model. This lends credence to our important conclusion regarding the inadequacy of single phase models (which implies the need for multiphase alternatives). Models with a slightly higher degree of complexity than the true model, however, are not clearly disfavored. One possible explanation for this ambiguity is that the extra phases in higher-complexity models might only contain a trivial amount of gas relative to the content of the other phases. The small amount of gas in these extra phases might produce an unnecessarily optimized fit, an improvement that is balanced by the penalty of adding parameters. Under this hypothesis, the higher complexity model minus the trivial phases would look similar to the true model. However, Occam’s razor would have us prefer the less complex model if the Bayes factor does not clearly favor the more complex one.

Additionally, we note that for these particular mock absorbers the Bayes factors indicate a weak, but consistent preference for purely photoionized models over ones incorporating collisional ionization. While this trends in the right direction (the true models are purely photoionized), it is not strong enough to be conclusive, likely owing to the paucity of high ion constraints in these mock absorbers (see Table 2 for a list of the species used). Note that this is not evidence for a universal preference for photoionization over collisional ionization, but only for the particular case of Absorbers 0.4A & 0.4B.

A5 Posterior Predictive p-values

Fundamentally, data (y^{sim}) simulated from a good model should look similar to the observed data (y^{obs}). This can be stated in terms of some test statistic $T(y)$ that summarizes each dataset: $T(y^{obs})$ should have a reasonable probability of being drawn from the distribution of $T(y^{sim})$. If $T(y^{obs})$ is

an outlier, we reject the generative model. p -values are commonly used to quantify the statistical significance of a measurement with respect to the null hypothesis or null distribution. The classical p -value is calculated over the distribution of y^{sim} for fixed (known) θ as:

$$p_C = \Pr(T(y^{sim}) \geq T(y^{obs}|\theta)). \quad (A20)$$

Thus, classical p -values summarize the discrepancy between the observed data and the values expected by a model with a particular value of θ .

In a Bayesian context, the observed data is compared to the posterior predictive distribution, which integrates over the posterior distribution of the (unknown) parameters. Thus, T is typically a function of θ as well as y , i.e., $T=T(y,\theta)$. The (Bayesian) posterior predictive p -value is thus defined as (Gelman et al. 1996, 2004):

$$p_B = \Pr(T(y^{sim}, \theta) \geq T(y^{obs}, \theta)|y) \quad (A21)$$

where the probability is taken over the posterior distribution of θ and the posterior predictive distribution of y^{sim} , i.e., the joint distribution $p(y^{sim}, \theta|y^{obs})$. In practice, one computes $T(y^{obs}, \theta)$ and $T(y^{sim}, \theta)$ for each set of parameters sampled from the posterior. The mock data, y^{sim} , are randomly drawn from the likelihood distribution, with one replicated dataset for each of the posterior samples.

The optimal test statistic for computing p_B depends on the model at hand. Here, we choose to sum the scaled residuals between the data and model column densities,

$$T(y, \theta) = \sum_i^{ions} \frac{|y_i - N_i^{model}(\theta)|}{\sigma_i^{obs}}. \quad (A22)$$

Here, y_i represents either the observed or simulated column density for the i th ion, N_i^{model} is the model-predicted column density corresponding to a particular set of parameters θ , and σ_i denotes the uncertainty in the measurement of the corresponding observed column density (even when $y_i=y_i^{sim}$). This statistic was chosen to complement the qualitative goodness-of-fit plots in Section 4, since discrepancies between the data and model are more intuitively understood when expressed as residuals. With this statistic, p_B should be interpreted as a two-tailed probability, with an ideal value at $p_B=0.5$. Large residuals drive $p_B \approx 0$ and indicate that the observed data are unlikely predictions of the model (i.e., potentially warranting a rejection of the model). Overly narrow residuals, which produce $p_B \approx 1$, are also unlikely to occur if the null hypothesis (the model being tested) is the correct one, so high numbers likely indicate some degree of overfitting.

Figure A3 illustrates the concept of scaled residuals. The likelihood distribution, $L(\mathbf{N}|\theta)$ (Equation A6), gives the probability of the data given the model, so we treat each observed column density as a random sample from a ZTN distribution (Equation A4). The ionization model sets the mean of this distribution, while its variance is fundamentally determined by the noise in the data, which we approximate as σ_i^{obs} . Figure A3 plots two sampling distributions for an arbitrary ion along with the corresponding simulated column densities. We are interested in the significance of any discrepancy rather than the absolute value of the discrepancy, hence the need for scaling by the standard deviation.

Extracting scaled residuals for censored data (upper and/or lower column density limits) is not as straightforward as with measurements for detected species. In these cases the

³⁰ We use the lower complexity PPa-Pb model simply to narrow the parameter selection space of the “simple” mock absorbers. The accuracy of the model with respect to the real data has no effect on the efficacy of the mock absorbers and tests derived from it.




Figure A3. A graphical illustration of scaled residuals. Model-predicted column densities μ_1 and μ_2 (blue and orange vertical lines, respectively), corresponding to two arbitrary parameter vectors θ_1 and θ_2 drawn from the (hypothetical) posterior, become the means of ZTN distributions (identically colored curves) each having the same standard deviation σ_{obs} . The distance r^{obs} (black arrows) between either μ_1 or μ_2 and the observed column density (vertical black dotted line) can be expressed as a column density (bottom axis) or as a scaled residual (top axis, $\mu^{\text{model}} \in \{\mu_1, \mu_2, \dots\}$). When computing posterior predictive p -values, a single simulated column density y_i^{sim} is drawn from each unique ZTN, which is linked to the unique posterior sample θ_i . We illustrate this concept for the two simulated column densities y_1^{sim} and y_2^{sim} , denoted by dashed vertical lines that are colored to correspond to the ZTN distributions from which they were drawn. A scaled residual is likewise computed for each of these simulated data points, with the distance r^{sim} shown by the respective colored arrows.

observed data residuals $|y_i^{\text{obs}} - N_i(\theta)|$ are themselves limits, a point illustrated in Figure A4. To make a fair comparison with model simulated column densities, any simulated columns that fall outside the observed detection limit are likewise treated as limits, and we ignore these when computing test statistics $T(y^{\text{obs}}, \theta)$ and $T(y^{\text{sim}}, \theta)$. In some cases, however, the posterior predictive distribution will generate y^{sim} values that would be considered detections in a real dataset given the signal-to-noise ratio of the observed spectrum. Since these provide a real constraint on the ionization modeling, we resort to a probabilistic analysis of the observed limits to provide a meaningful comparison to the simulated data. As noted above in Appendix A1, each column density limit is interpreted as the tail probability of the ZTN distribution implied by a given set of model parameters. Thus, we can realize an observed column density from the limits by sampling the relevant tail of the model-implied column density sampling distribution.³¹ We then use the realized column density to compute $T(y^{\text{obs}}, \theta)$ as though it were an actual detection. This approach does not bias the test statistic since the realized and simulated column densities are ultimately drawn from the same distribution (i.e., before applying the observational limit constraints). The darkest shaded regions in Figure A4 exemplify the tail regions from which upper limit column density realizations are drawn. Lower limits are similar but mirrored, while double limits specify sampling intervals rather than tails.

³¹ Note that without the tail constraints this would be identical to the sampling procedure used to generate simulated data.




Figure A4. An illustration of column density residuals for lower limits. In both cases, μ , the mean of the ZTN distribution (red curve), is the model-predicted column density corresponding to a particular set of parameters drawn from the posterior of a given model. We compute the posterior predictive probability distribution by integrating over all possible values of μ implied by the posterior. The red vertical dashed lines represent simulated column densities drawn from the ZTN for two illustrative cases, while the dotted black lines indicate the measured upper limit. As a necessary step to obtaining a p -value, we must determine if the simulated measurement produces a larger residual with respect to μ than the observed limit. Case 1 differs from Case 2 in the location of μ relative to the observed column density limit. In both cases the residual r^{sim} is straightforward to measure, while we can only measure an upper limit on r^{lim} . In some regions (white/not shaded) it is clear that either $r^{\text{sim}} < r^{\text{lim}}$ (Case 1) or $r^{\text{sim}} > r^{\text{lim}}$ (Case 2), but in the lightly shaded regions it is not clear which residual is larger. We overcome this problem by statistically generating realizations of the column density limits. That is, we draw samples from the posterior predictive distribution in the region allowed by the measured limit (shaded in dark grey). Note that when $y^{\text{sim}} < y^{\text{lim}}$ we treat both r^{lim} and r^{sim} as limits and ignore these points when computing the test statistics $T(y^{\text{lim}}, \theta)$ and $T(y^{\text{sim}}, \theta)$.

This paper has been typeset from a \LaTeX file prepared by the author.