

MEASURING COLUMN DENSITIES IN QUASAR OUTFLOWS: VLT OBSERVATIONS OF QSO 2359–1241¹

NAHUM ARAV,^{2,3} MAXWELL MOE,² ELISA COSTANTINI,⁴ KIRK T. KORISTA,⁵ CHRIS BENN,⁶ AND SARA ELLISON⁷

Received 2008 January 2; accepted 2008 March 17

ABSTRACT

We present high-resolution spectroscopic VLT observations of the outflow seen in QSO 2359–1241. These data contain absorption troughs from five resonance Fe II lines with a resolution of $\sim 7 \text{ km s}^{-1}$ and a signal-to-noise ratio per resolution element of order 100. We use this unprecedented high-quality data set to investigate the physical distribution of the material in front of the source and by that to determine the column densities of the absorbed troughs. We find that the apparent optical depth model gives a very poor fit to the data and greatly underestimates the column density measurements. Power-law distributions and partial covering models give much better fits, with some advantage to power-law models, while both models yield similar column density estimates. The better fit of the power-law model solves a long-standing problem plaguing the partial covering model when applied to large distance scale outflow: how to obtain a velocity-dependent covering factor for an outflow situated at distances thousands of time greater than the size of the AGN emission source. This problem does not affect power-law models. Therefore, based on the better fit and plausibility of the physical model, we conclude that in QSO 2359–1241, the outflow covers the full extent of the emission source but in a nonhomogeneous way.

Subject headings: galaxies: individual (QSO J2359–1241) — line: formation — quasars: absorption lines

Online material: extended figure set

1. INTRODUCTION

In recent years, the potential impact of quasar outflows on their environment has become widely recognized. (e.g., Blandford & Begelman 2004; Scannapieco & Oh 2004; Vernaleo & Reynolds 2006). Observationally, these outflows are detected as absorption troughs in quasar spectra that are blueshifted with respect to the systemic redshift of their emission-line counterparts. The absorption troughs are mainly associated with UV resonance lines of various ionic species (e.g., Mg II $\lambda\lambda 2796.35, 2803.53$; C IV $\lambda\lambda 1548.20, 1550.77$; Si IV $\lambda\lambda 1393.75, 1402.77$; and N V $\lambda\lambda 1238.82, 1242.80$). A large spread in maximum velocity is detected for different quasar outflows, from several hundred km s^{-1} to more than $30,000 \text{ km s}^{-1}$. Outflow troughs are known as intrinsic absorbers to distinguish them from intervening and associated absorbers, which are also detected in quasar spectra (Hamann et al. 1997; Barlow et al. 1997). The lines we study here are intrinsic narrow lines (see Arav et al. 2002 for full discussion) and are similar to the ones analyzed by de Kool et al. (2001).

Reliable measurements of the absorption column densities in the troughs are crucial for determining almost every physical aspect of the outflows: the ionization equilibrium and abundances, number density, distance, mass flux, and kinetic luminosity. Column density measurements depend on the absorption model one adopts for the observed troughs in the spectrum.

In the case of interstellar medium (ISM) and intergalactic medium (IGM) absorbers (the latter also known as intervening absorbers), the apparent optical depth method is an excellent absorption model. This method postulates that the absorber covers the full size of the emission source homogeneously and that the relationship between the residual intensity in the trough (I) and the optical depth of the line (τ) is given by $\tau_{\text{ap}} \equiv -\ln I$. The derived τ is then converted to column density using a standard formula (e.g., eq. [1] in Arav et al. 2001a).

However, our group (Arav 1997; Arav et al. 1999a, 1999b, 2002, 2001b, 2003; de Kool et al. 2001; Scott et al. 2004; Gabel et al. 2005a) and others (Barlow et al. 1997; Telfer et al. 1998; Churchill et al. 1999; Ganguly et al. 1999) have shown that the apparent optical depth method is not a good approximation for outflow troughs (see also § 3). Most of this evidence came from data sets that showed fully resolved troughs from unblended doublets (e.g., the N V, Si IV, C IV, and Mg II lines mentioned above). For all these doublets, the oscillator strength of the blue (the shorter wavelength) transition is twice that of the red transition. Therefore, the expected optical depth ratio is 2:1 in favor of the blue transition. (It is not exactly 2:1, since the wavelengths differ by a small amount.) For the apparent optical depth method to hold, we must have $\tau_{\text{ap}}(\text{blue}) = 2\tau_{\text{ap}}(\text{red})$, which requires $I_B = I_R^2$, where I_B and I_R are the residual intensities of the blue and red absorption troughs, respectively. Using high signal-to-noise ratio (S/N) observations of fully resolved and unblended outflow doublets, it became clear that in most cases I_B is significantly smaller than I_R^2 , demonstrating that the τ_{ap} model is not applicable to active galactic nucleus (AGN) outflows.

To avoid this contradiction for cases of unblended doublets, most works use a pure partial covering model to determine the real optical depth and hence the actual column density. This model assumes that only a fraction C of the emission source is covered by the absorber and then solves for a combination of C and τ that will fit the data of both doublet troughs while maintaining the intrinsic 1:2 optical depth ratio (see § 3 of Arav et al. [2005] for the full formalism, including velocity dependence). The weakness

¹ Based on observations made with ESO Telescopes at the Paranal Observatories under program ID 078.B-0433(A).

² CASA, University of Colorado, 389 UCB, Boulder, CO 80309-0389; arav@colorado.edu.

³ Department of Physics, Virginia Tech, Blacksburg, Va 24061; arav@vt.edu.

⁴ SRON National Institute for Space Research Sorbonnelaan 2, 3584 CA Utrecht, Netherlands.

⁵ Western Michigan University, Department of Physics, Kalamazoo, MI 49008-5252.

⁶ Isaac Newton Group, Apartado 321, E-38700 Santa Cruz de La Palma, Spain.

⁷ Department of Physics and Astronomy, University of Victoria, Victoria, BC V8P 1A1, Canada.

in this method is that we solve for two unknowns (C and τ) given two residual intensity equations. As long as the ratio of the residual intensities is in the permitted physical range (see § 3 of Arav et al. 2005), such a procedure will always yield a solution. However, this is also the case for other two-parameter fitting algorithms, which yield different estimates of column densities. A main alternative to the pure partial covering model is the so-called inhomogeneous absorber model (de Kool et al. 2002b, hereafter dKKA02; see § 3.2.2 here). This model successfully matched the data of two quasar outflows (dKKA02), while proving less applicable for the outflow troughs of Mrk 279 (Arav et al. 2005; but see § 4).

Being able to distinguish between these two (and other) absorption models is important not only for column density measurements, but also for obtaining a more accurate geometrical picture of the outflow, including its internal structure. Since doublet observations are inherently inconclusive in distinguishing between these methods (or any other two-parameter models), to resolve this issue we need data sets that show outflow troughs from more than two lines of the same ion. These troughs must be unblended, be fully resolved, and have a high signal-to-noise ratio (S/N).

In this paper we present unprecedented high-quality observations of a quasar outflow, which covers troughs from five Fe II resonance lines associated with Fe II UV multiplets 1, 2, and 3 (see Table 3). We use the strong constraints available from these troughs to test models of absorption material distribution and to determine the ionic column densities in all the troughs seen in this outflow. The object we targeted is the extensively studied broad absorption line (BAL) quasar QSO 2359–1241 (Arav et al. 2001a; Brotherton et al. 2001). Having detailed knowledge of the existing Keck HIRES data allowed us to tailor the new observations toward a definitive test for models of absorption material distribution. In future papers we will use these measured ionic column densities to determine the ionization equilibrium of the outflow, its number density, distance, and ultimately the mass flux and kinetic luminosity associated with the outflow.

The plan of the paper is as follows. In § 2 we describe the observations, the data reduction, and the phenomenology of the outflow; in § 3 we test the goodness of the fit for the five Fe II resonance lines obtained from different models of absorption material distribution; in § 4 we discuss our results, and in § 5 we summarize them.

2. VLT OBSERVATIONS OF THE QSO 2359–1241 OUTFLOW

2.1. Data Acquisition and Reduction

We observed QSO 2359–1241 (radio source NVSS J235953–124148, $z = 0.868$; Arav et al. 2001a) using the Very Large Telescope (VLT) operated by the European Southern Observatory (ESO) in 2006 Fall (see Table 1 below). The observations were done in service mode (i.e., done by the VLT staff at their scheduling, without active participation of the science team while the observations were taken) using the UV-Visual Echelle Spectrograph (UVES). All observations used a slit width of $1.0''$. In order to cover the full wavelength range, we used two UVES settings. The observations taken in 2006 September used the following setting: wavelength centers blue = 346 nm (grating/filter CD1/HER5) and red = 580 nm (grating/filter CD3/SHP700), and the observations taken in 2006 October used wavelength centers $B = 437$ nm (grating/filter CD2/HER5) and $R = 860$ nm (grating/filter CD4/OG590). Both settings were read out in 2×2 pixel-binned, high-gain mode. The useful spectral coverage

TABLE 1
VLT UVES OBSERVATIONS OF QSO 2359–1241

Date	UT Start Time	Integration (s)	Mean Air Mass
2006 Sep 30.....	01:05	2850	1.40
2006 Sep 30.....	01:57	2850	1.19
2006 Sep 30.....	02:47	2850	1.08
2006 Sep 30.....	03:38	2850	1.03
2006 Oct 11.....	00:07	2850	1.49
2006 Oct 11.....	00:57	2850	1.25
2006 Oct 11.....	01:46	2850	1.11
2006 Oct 11.....	00:02	2850	1.50

was 3200–9000 Å, compared to 4320–7450 Å for the Keck HIRES data. More important, the signal-to-noise ratio of the VLT data over the spectral range covering the five Fe II resonance lines was roughly 3 times higher where both data sets have similar spectral resolution. This large increase in signal-to-noise ratio allowed us to perform the analysis we present in this paper, which was not possible with the Keck HIRES data. The extraction of the spectra was done using the UVES pipeline, which is based on ECHELLE routines in the data reduction package MIDAS. A detailed description of this process can be found in Ballester et al. (2000).

The UVES data reduction pipeline gave excellent data products overall. However, one issue that produced considerable systematic errors is spectral undulations introduced by the pipeline. These undulations mimic narrow emission line structure in the data, with FWHM of about 10 Å and maximum flux of 20% above the adjacent continuum. It is clear that these are introduced by the reduction pipeline, since we have overlapping spectral coverage over part of the affected range (4795–4950 Å) and undulations appear in only one of the settings. The origins of these emission-line structures probably come from the bad columns located at the bottom of the lower red CCD chip (C. Ledoux [UVES instrument scientist] 2007, private communication). Fortunately, the observed absorption troughs in that spectral region are considerably narrower (about 3 Å full width) than the FWHM of the undulations. Therefore, we are able to model the undulations out of the data to a large extent (see § 2.2). Careful examination of the continuum fits in the overlapping spectral regions leads us to estimate a 2%–3% maximum systematic error associated with these data reduction undulations. We note that the undulations affect less than 5% of the entire spectrum range, but unfortunately this includes the region of the important Fe II UV 1 multiplet, which includes two of the Fe II resonance lines under study here (2586.65 and 2600.17 Å).

2.2. Data Co-adding and Normalization

Each of the four exposures for a given setting and wavelength region was heliocentric-corrected to vacuum wavelengths given by the fits header file. We then used a single multiplier value for each exposure such that all spectra had the same mean value. Rejection of cosmic rays was done by replacing any flux value 3.5σ or more above the mean of the other 3 values and replacing it by that mean value. Finally, the four different spectra were co-added via the weighted average at each data point based on the given error for the four different exposures. The FWHM of the spectral resolution was determined to be 8.5 ± 0.2 km s⁻¹ based on the weighted average of fitting 1 Gaussian profiles to the normalized flux of several atmospheric absorption lines. In view of this resolution, we chose to bin the data to 7.0 km s⁻¹, which

TABLE 2
TRANSITIONS INCLUDED IN THE QSO 2359–1241 LINE LIST

$\log(gf)$	E_{low} (cm^{-1})	Ions
> -2.5	All	He I, Mg II, Al II, Al III, Si II, Ca II
> -2.0	$E_{\text{low}} < 1000$	Fe II
> -1.5	$1000 < E_{\text{low}} < 4000$	Fe II
> -1.0	$E_{\text{low}} > 4000$	Fe II
> -1.0	All	Mg I, Cr II, Mn II, Ni II
> -0.8	All	H I

slightly oversamples the actual resolution, before normalization and analysis.

The algorithm used to obtain a model for the unabsorbed continuum entails initially rejecting expected absorption lines a priori from the line lists (see Table 2). An initial cubic fit was applied over every spectral point (pixel) spanning 125 points on either side of a given spectral point. Data points 8σ or more below the fit value were then rejected as absorption, and the cubic fit was reapplied. This process of refitting and rejecting was reiterated at 6, 4, and then 3 σ . The chosen continuum at this spectral point was the value of the last cubic fit for that point. This process was repeated for all data points in the spectrum. The resultant continuum fit was smoothed with a boxcar average of five data points. To normalize the data, we then divided the observed flux values by this model continuum. For problem areas such as the data reduction undulations, atmospheric absorption regions, or wide/multiple absorption troughs, the parameters above were slightly altered to accommodate these areas, but the basic principle of fitting, rejecting, and refitting remained the same. The only exception to this normalization process was at the spectral vicinity of the Mg II broad emission line (BEL), where a pair of χ^2 fitted Gaussians was used to divide out the Mg II emission before normalization.

2.3. Identification of Spectral Features

The atomic transitions data used to identify spectral features, normalize the data (see above), and fit the absorption troughs (see § 3) came from Morton (2003) for Fe II transitions with lower energies (E_{low}) below 1000 cm^{-1} and from Kurucz & Bell (1995) for all other transitions. Our assembled line list contains all transitions in the spectral region of the VLT data for QSO 2359–1241 with constraints $E_{\text{low}} < 10,000 \text{ cm}^{-1}$ (except for H I and He I, which can go up to $160,000 \text{ cm}^{-1}$), and the constraints on their gf values (where g is the degeneracy and f is the oscillator strength of the transition) are shown in Table 2.

In order to identify the spectral features in the spectrum, we first measured the redshift of the main outflow component for the Fe II $\lambda 2587$ resonance line ($z = 0.8594$) and then identified absorption features of other lines from all ions at that redshift. Table 3 gives the spectroscopic information for the five Fe II resonance lines that are the focus of the analysis presented here. A list of all the lines from these multiplets (including those from excited levels) is given in Table 1 of de Kool et al. (2001). The oscillator values (f) are from Morton (2003).

2.4. Description of Outflow Absorption Features

In the HIRES paper (Arav et al. 2001a), we have shown that the absorption features in this object originate from an outflow connected with the QSO, i.e., intrinsic in nature as opposed to intervening or associated absorption. In that paper we also labeled

TABLE 3
Fe II RESONANCE LINES ANALYZED IN THIS PAPER

Multiplet	λ Vacuum	Relative f^a
1.....	2600.17	0.81
	2586.65	0.23
2.....	2382.77	1.00
	2374.46	0.10
3.....	2344.21	0.35

^a Oscillator strength relative to $f(2382.77) = 0.32$.

the outflow components, and since these were unchanged between the two epochs, we use the same labeling system here. Figure 1 shows the absorption components seen in the Fe II $\lambda 2587$ (from the Fe II UV 1 multiplet) resonance line. Most of our analysis in this paper concentrates on the widest and highest velocity component, namely, *e*. This component is seen in all the outflow absorption features detected in this object. In Figure Set 2 we show the normalized flux for the spectral range of the Fe II UV 1 multiplet, where component *e* is labeled for all the lines from this multiplet. Most of the troughs arise from low excitation levels associated with the Fe II UV 1 multiplet. The online version of this figure shows the entire spectrum (spread over 46 pages), where the absorption features of all the detected ions are identified in a similar way.

3. MODELING THE ABSORPTION TROUGHS

As stated in § 1, in order to extract quantitative information about the outflow, we need to measure the column density associated with the absorption troughs. In this section we test four models for the distribution of absorption material in front of the emission source, in order to find which one gives the best fit for the VLT data of QSO 2359–1241.

3.1. Failure of the Apparent Optical Depth Model

The apparent optical depth method implicitly assumes that the absorption material covers the source completely and homogeneously. That is, all light rays that arrive at the observer pass through material with the same optical depth at a given wavelength. Under these assumptions, the residual intensity of a given line is $I_1 = e^{-\tau_{\text{ap}}}$, and the optical depth ratio for two $E = 0$ lines from the same multiplet is $R_{21} = f_2 \lambda_2 / f_1 \lambda_1$, where f is the oscillator strength and λ is the wavelength of the transition. Combining the two relations, we derive the expected residual intensity of line 2 given the observed residual intensity of line 1: $I_2 = e^{-R_{21} \tau_{\text{ap}}} = I_1^{R_{21}}$.

In Figure 3 we compare the residual intensities of the outflow troughs arising from two resonance lines ($\lambda 2587$ and $\lambda 2600$) of the UV 1 Fe II multiplet, with the apparent optical prediction for $I(\lambda 2600)$ based on $I(\lambda 2587)$, using the above relationship ($R_{21} = 3.5$ for this case). The inadequacy of the τ_{ap} model in predicting $I(\lambda 2600)$ is clear, demonstrating that the τ_{ap} model assumptions are invalid for the Fe II resonance troughs. Similar results are obtained for other ions. We therefore need to explore more complicated models. Below, we describe three such models: partial covering, inhomogeneous absorber, and modified partial covering.

3.2. Two- and Three-Parameter Models

3.2.1. Partial Covering Model

Applying the apparent optical depth method is in essence using a one-parameter model (τ_{ap}) for fitting the absorption troughs

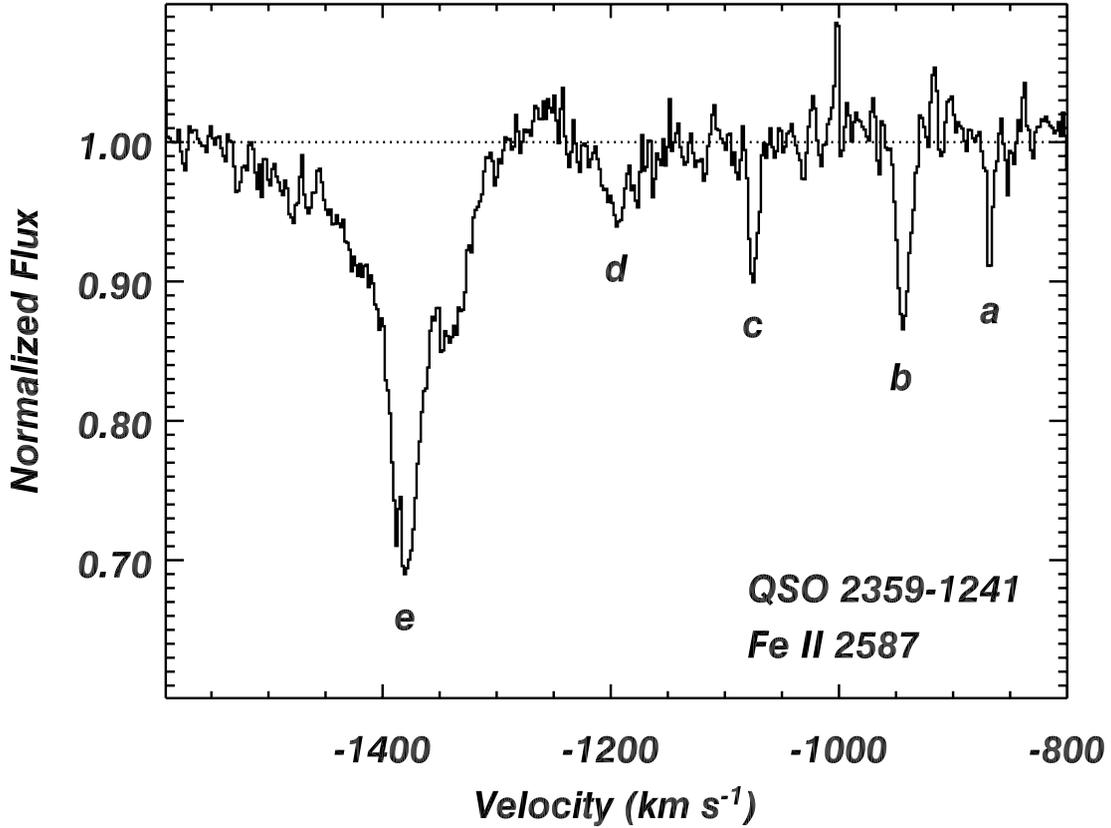


FIG. 1.—Morphology of the outflow in QSO 2359–1241. Plotted are the normalized VLT UVES data for the Fe II λ 2587 resonance line. The outflow velocity is measured compared to the systemic redshift ($z = 0.868$). The labeling of the components follows that of Arav et al. (2001a), and most of our analysis is done on component *e*.

$[\tau_{\text{ap}} \equiv -\ln I]$. Two-parameter models are the obvious next step. Most of the previous work, which addressed the inadequacy of the apparent optical depth method in modeling outflow troughs, dealt with data of unblended doublets (e.g., Barlow et al. 1997; Arav et al. 1999b, 2002; Ganguly et al. 1999; Scott et al. 2004). The preferred method for modeling these troughs is the partial covering model, which is a two-parameter model (Hamann et al. 1997; Barlow et al. 1997). To summarize this method, we assume that only a fraction C of the emission source is covered by the absorber and then solve for a combination of C and optical depth (τ) that will fit the data of both doublet troughs while maintaining the intrinsic 2:1 optical depth ratio (see § 3 of Arav et al. 2005 for the full formalism, including velocity dependence). Implicit in this model are the assumptions that the absorber covering the fraction C of the source has the same optical depth across this area and that the rest of the source is covered by material with $\tau = 0$ in that transition. A geometrical illustration of this assumption is shown in Figure 4. The partial covering model allows us to find a good fit for the doublet troughs and, under the model assumption, a consistent determination of the optical depth and hence the column density of the trough. This is an important improvement over the apparent optical depth method, which cannot yield consistent optical depth estimates for the two troughs.

As stated in § 1, the weakness of the partial covering model as applied to doublet troughs is that we solve for two unknowns (C and τ) given the two residual intensity equations of the doublet troughs. As long as the ratio of the residual intensities is in the permitted physical range (see § 3 of Arav et al. 2005), such a procedure will always yield a solution. Therefore, a good fit is a necessary but not sufficient condition to validate the underlying model. In addition, other two-parameter models may fit the data equally

well, where the most studied alternatives are inhomogeneous absorption models.

3.2.2. Inhomogeneous Absorption Models

A more general way to modify the apparent optical depth assumptions is to assume that the absorption material does not cover the source homogeneously. Since there is an infinite number of ways of doing so, we need a quantitative model that can be directly compared with the partial covering results. The first attempt to study generic inhomogeneous distributions of absorbing material in AGN outflows was made by de Kool et al. (2002a; dKKA02). A formalism was developed to simulate the effects of an inhomogeneous absorber on the emerging spectrum, and several examples of fitting existing data with this model were presented. Arav et al. (2005) built on this formalism and produced a two-parameter inhomogeneous model that can be directly compared with the partial covering results. For full details we refer the reader to these two papers, and here we give a brief description of the idea based on § 2 of Arav et al. (2005). Our starting point is equation (6) from dKKA02:

$$F(\lambda) = \int \int S(x, y, \lambda) e^{-\tau(x, y, \lambda)} dx dy, \quad (1)$$

where $S(x, y, \lambda)$ is the surface brightness distribution of the background source and $\tau(x, y, \lambda)$ is the line-of-sight optical depth at wavelength λ in front of a specific (x, y) location, as defined by dKKA02, equations (3)–(5).

In its most general form, equation (1) is not practical for modeling spectra. Both $S(x, y, \lambda)$ and $\tau(x, y, \lambda)$ are unconstrained, two-variable functions for each given λ . We need a simpler and testable

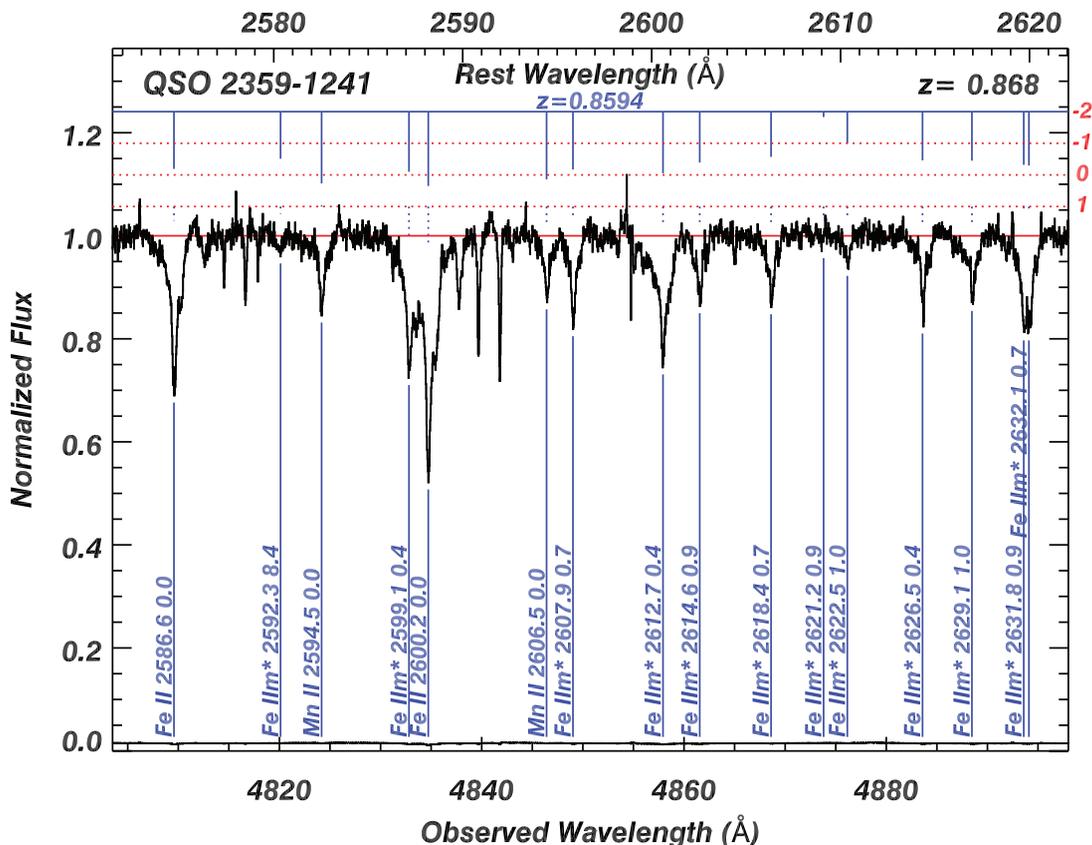


FIG. SET 2.—Small portion of the normalized VLT UVES spectrum of QSO 2359–1241, showing the outflow features associated with the Fe II UV 1 multiplet. (A similar presentation of the entire spectrum is available online.) The data plotted on Fig. 1 are shown here between 4806 and 4819 Å (observed). Blue vertical lines denote the expected position of component e for various transitions, based on the $z = 0.8594$ redshift of this component. Most of these transitions are associated with Fe II UV 1 multiplet (two transitions from Mn II and one from a higher Fe II multiplet are also marked). The size of the identification line (measured from the top solid line) gives the $\log(gf)$ value of the transition (where f is the oscillator strength and g is the degeneracy of the transition), which can be read from the values associated with the dotted lines (right axis). Each transition is labeled by its ion, wavelength, and energy level ($\text{cm}^{-1}/1000$). The three unlabeled troughs around 4840 Å and similarly around 4815 Å are components b , c , and d of Fe II $\lambda 2600$ and Fe II $\lambda 2587$, respectively. The error spectrum is also plotted (just above the x -axis). The short dotted vertical lines just above the continuum show the exact position of the trough around the continuum level. [See the electronic edition of the Journal for Figs. 2.1–2.47.]

model. To this end, we introduce the following simplifying assumptions (whose physical validity and plausibility are discussed in Arav et al. 2005).

1. $S(x, y, \lambda) = S(\lambda)$, which can be set to $S(\lambda) = 1$ when dealing with normalized data.

2. $\tau(x, y, \lambda) = \tau(x, \lambda)$, a simplifying assumption entailing little loss of generality. We can think of $e^{-\tau(x, \lambda)}$ as the integrated attenuated flux along a y -strip at a specific (x, λ) .

3. $\tau(x, \lambda) = \tau_{\max}(\lambda)x^a$, which yields a two-parameter model that can be directly compared with the partial covering model. We note that a Gaussian characterization also yields a two-parameter model. However, these models were investigated by Arav et al. (2005; see their discussion of Fig. 3) and were shown to yield results very similar to power-law models. We therefore use the power-law model as the sole representative of two-parameter inhomogeneous absorption models.

With the last simplification, we have a power-law inhomogeneous model in which each of the two parameters (τ_{\max} and a) can be velocity dependent, similar to C and τ for the partial covering model. In general, these two models have quite different physical interpretations, as discussed in § 4.

3.2.3. Modified Partial Covering Model

The pure partial covering model is unphysical, since there is a sharp edge (essentially a step function) between a region with

finite optical depth and a region with $\tau = 0$. Physically, there needs to be a transition zone between these two regions. This motivated us to add a parameter to the partial covering model to control the gradient between these two regions. The phenomenological formalism we chose is

$$\tau = \tau_0(1 - 1/[e^{b(x-x_0)}]), \quad (2)$$

where x_0 and τ_0 are similar to C and τ for the pure partial covering model and b controls the gradient.

In Figure 4 we show the optical depth distribution across the source for all four models. The examples we show are tailored to the best fits of each model to the Fe II $E = 0$ lines (see § 3.3)

3.2.4. Conclusive Testing of the Absorption Models

In order to test the partial covering and the inhomogeneous absorption models, we need data sets that will overconstrain them, that is, by fitting more than two equations with the same two unknowns (or three, in the case of the modified partial covering model). To do so, we need to observe more than two troughs from the same ion.

Our VLT UVES observations of QSO 2359–1241 yield the best such data set to date. In the outflow of this object, we detect absorption troughs associated with more than 30 lines of the Fe II multiplets UV 1, UV 2, and UV 3 (see Fig. Set 2 for the spectral region of the Fe II UV 1 multiplet). These lines include

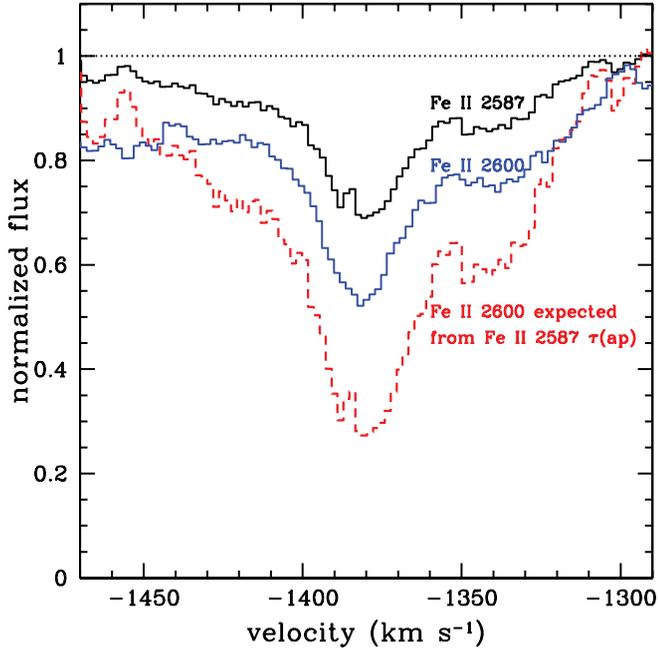


FIG. 3.— Failure of the apparent optical depth model (τ_{ap}). Normalized residual intensity for Fe II 2587 and Fe II 2600 Å are shown for the main component of the outflow (solid histograms). The dashed red histogram shows the expected depth of the Fe II 2600 Å based on the residual intensity of the Fe II 2587 Å line and using the τ_{ap} model (see text for more details). It is evident that the expected depth is a poor fit to the observed Fe II 2600 trough, demonstrating the inadequacy of the τ_{ap} assumption for this outflow. Similar results are obtained for the majority of quasar outflows.

five transitions from the $E = 0$ energy level, which in principle can give us five residual intensity equations to fit the two and three free parameters of the above models. To realize this possibility, we chose this target, S/N, and spectral resolution such that we would acquire an unprecedented data set for these purposes. The troughs are fully resolved, have $S/N \geq 80$ per resolution element, and show only minor blending with other troughs compared to other objects showing intrinsic Fe II* troughs (e.g., Wampler et al. 1995; de Kool et al. 2001, 2002a).

3.3. χ^2 Fitting

In order to get a quantitative measure of how well each model fits the data, we use χ^2 minimization. Since the trough is well resolved, we binned the data to match the measured resolution. The resultant data products in the case of the Fe II resonance lines are residual intensities for five $E = 0$ lines covering the span -1475 to -1305 km s^{-1} . Each velocity bin is treated as an independent measurement, and we use a least-squares χ^2 minimization routine to find the best parameters of each model for that specific bin. The process is repeated for all velocity bins. Two of the troughs suffer a moderate amount of blending with other lines, and we omit the contaminated region of the individual trough while performing the fit. These contaminations occur for the Fe II 2600 Å trough between -1475 and -1420 km s^{-1} and for the Fe II 2344 Å trough between -1340 and -1305 km s^{-1} . As a result, we always fit at least four data points with each model, five for the -1420 to -1340 range, which includes the deepest part of the trough. Since the maximum number of parameters we fit is three, the models are always overconstrained.

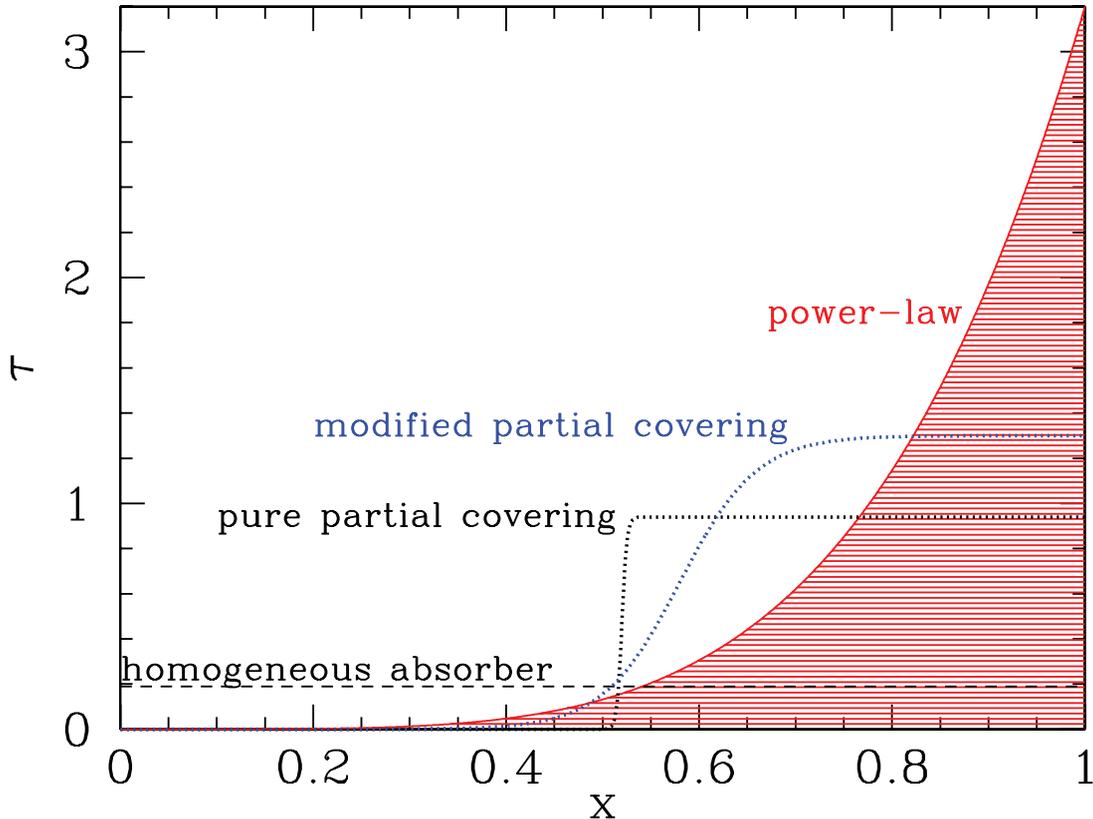


FIG. 4.— Four distributions of absorbing material that we investigate in this paper. The constant emission source coincides with the x -axis, and the observer is far above the plot. For the power-law case we shaded the region of absorbing material, and for all distributions the absorbing material is below their curves. For each distribution the parameters are representative of the deepest part in the e component of the outflow, where a χ^2 fit for all five resonance lines was performed at that velocity, and the best-fit model for each functional form is shown; τ is that of the Fe II 2587 Å line. For the different distributions, the optical depth in front of an x -location is given by (1) homogeneous absorber (apparent optical depth case): $\tau(x) = 0.19$; (2) partial covering: $\tau(x) = 0$ for $0 < x < x_0$, and $\tau(x) = \tau_{\text{max}}$ for $x_0 < x < 1$, where $\tau_{\text{max}} = 0.94$ and $x_0 = 0.52$; (3) power law: $\tau(x) = \tau_{\text{max}} x^a$, where $\tau_{\text{max}} = 3.2$ and $a = 4.6$; (4) modified partial covering: $\tau = \tau_0 [1 - 1/(e^{b(x-x_0)})]$, where $\tau_0 = 1.3$, $b = 25$, and $x_0 = 0.58$.

To compare the goodness of the fit for each model, we calculate the reduced χ^2 defined in the following way (see Press et al. 1989, pp. 498–546):

$$\chi_{\text{red}}^2 = \frac{\sum_{i,j} (I_{i,j} - M_{i,j})^2 / \sigma_{i,j}^2}{N_{\text{tot}} - N_v N_p}, \quad (3)$$

where the sum over i is taken over the number of fitted data points in a given velocity bin (four or five in our case), the sum over j is taken over the 25 velocity bins we fit, $I_{i,j}$ is the residual intensity of trough i at velocity bin j , $M_{i,j}$ is the modeled residual intensity of trough i at velocity bin j , $\sigma_{i,j}$ is the associated error for each data point, N_{tot} is the total number of velocity bins used in the fit, N_v is the number of velocity bins (25 in this case), and N_p is the number of free parameters for each velocity bin (1 for the apparent τ model, 2 for the partial covering and power-law models, and 3 for the modified partial covering model). We note that since the minimization procedure omits the blended regions of two troughs, $N_{\text{tot}} = 112$, not 125.

Figure 5 shows the best χ^2 fit for each model, where for the sake of clarity we show only three of the five Fe II $E = 0$ lines detected in the spectrum. The fits were done using all five lines, and those that are shown span the full range in oscillator strength. From both visual inspection and the χ_{red}^2 value, it is clear that the τ_{ap} model gives a poor fit to the data. This is simply a more quantitative way of arriving at the conclusion shown in Figure 3. The important results shown in Figure 5 are the following:

1. Both the pure partial covering and the power-law distribution models give similar fits (with a small advantage to the power-law model).
2. The modified partial covering model does not yield an improvement in χ_{red}^2 value, although the fit itself is somewhat better. This surprising result is discussed in § 3.4.

Before we compare the various fits and discuss how physical they are, we need to address the absolute χ_{red}^2 values. Normally, we would expect the χ_{red}^2 value of a “good” fit to be ~ 1 . For our two leading models, we find $\chi_{\text{red}}^2(\text{PL}) = 3.8$ and $\chi_{\text{red}}^2(\text{PC}) = 4.5$ (where PL stands for power law and PC for pure partial covering). Nominally, these values exclude these models at a very high probability. For a model with 60 degrees of freedom, $\chi_{\text{red}}^2 = 3.8$ has probability less than 10^{-10} of being the correct model. However, the main issue here is the systematic errors that are not addressed by the fit. First, as noted in § 2, the undulations superimposed on the spectrum due to data reduction issues can contribute a 2%–3% systematic error in the normalized intensity of the troughs. The troughs that are most affected by these undulations (2600 and 2587 Å) also have the highest S/N value per resolution element (upward of 80). As a result, this systematic error by itself can produce $\chi_{\text{red}}^2 \sim 4$, even if the model matches the data perfectly. Another important source of systematic error is caused by uncertainties in the oscillator values of the Fe II resonance lines. Relative uncertainties in these values can produce a considerable χ_{red}^2 on its own. We tested this hypothesis by artificially changing the oscillator strength for the weakest transition (2374 Å) and found considerable changes in the χ_{red}^2 values. For example, for the PL model the value dropped from 3.8 to 3.4 when the oscillator strength of Fe II 2374 Å was increased by 40%.

We therefore are not troubled by the fact that in this actual case we obtain $\chi_{\text{red}}^2 \sim 4$, and we focus our discussion on the differences seen between the different models. It will be left to future data sets, with better control of the systematic errors, to show whether these models give reasonable absolute values of χ_{red}^2 .

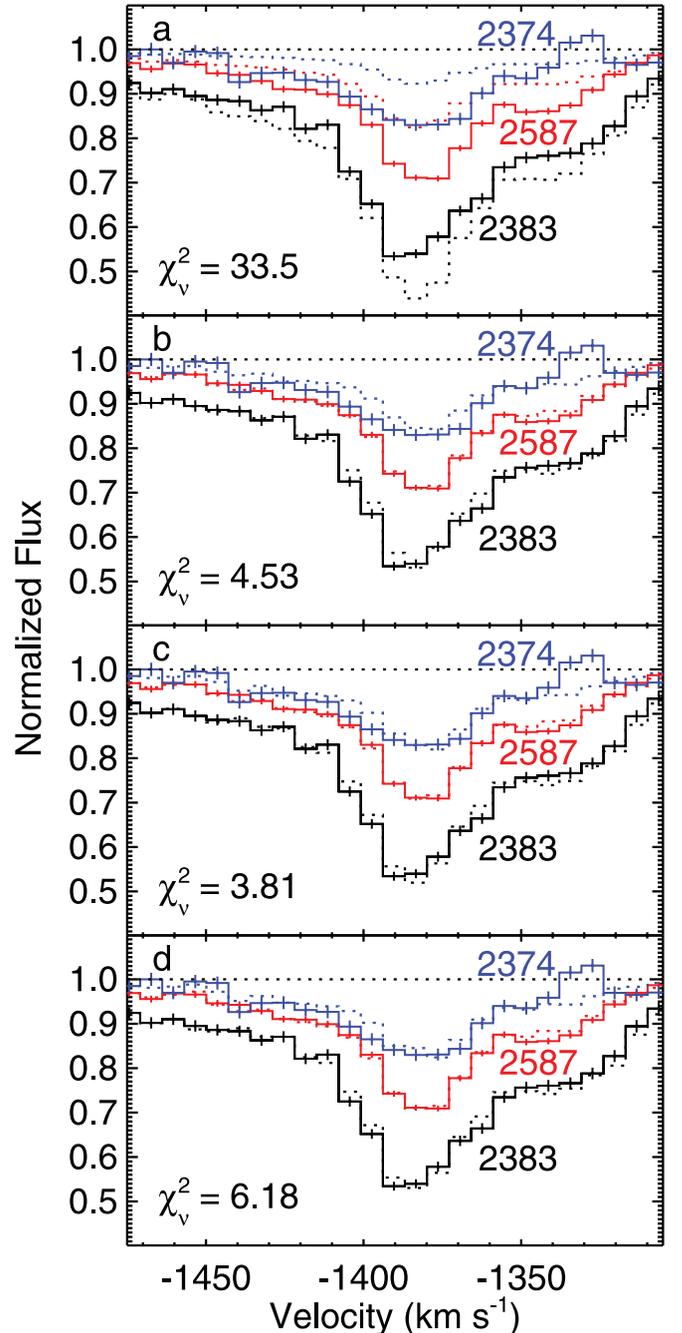


FIG. 5.—Best χ^2 fits for the Fe II $E = 0$ lines seen in the spectrum of QSO 2359–1241. Each panel shows the fit for a different model: (a) τ_{ap} model; (b) pure partial covering; (c) power-law model; (d) modified partial covering (see text for more details). The data plus error bars are shown in solid-line histograms, the fits are shown in dotted-line histograms, and the wavelength (Å) of each transition is labeled. Best-fit reduced χ^2 values are printed at the lower left corner for each model.

3.4. Comparison between the Different Models

As noted above, it has become clear over the past 10 years that the simplest possible model (τ_{ap}) is inadequate for modeling outflow troughs. Our analysis here only adds a well-constrained quantitative case to this effect. The next step is to compare the two-parameter models, namely, the pure partial covering model and the power-law model. These models represent rather different physical pictures with regard to the distribution of absorbing material in front of the emission source. We discuss their pros and cons in § 4, but as a first step, it is useful to make a simple

comparison of the fits to the same data by models with the same number of free parameters.

We start by assuming that the main reason for the high absolute values of χ_{red}^2 for these models is the systematic errors discussed above. By that, we accept both models as being reasonably good fits for the data. The $\chi_{\text{red}}^2 = 3.8$ of the power-law model is better than the $\chi_{\text{red}}^2 = 4.5$ of the pure partial covering model. However, assessing the significance of this difference in lieu of the large contribution from systematic error is more problematic, since it is difficult to predict what will happen to the relative differences in the χ_{red}^2 values once the systematic errors are better controlled. For the mean time, we cautiously accept the better fit of the power-law model. As we see below, evidence from other Fe II transitions and from other ions lends support to the better applicability of the power-law model as well.

It is also important to compare the velocity dependence of these models. For the partial covering model, we find the known result that the covering factor traces the shape of the strongest modeled resonance line $C(v) \approx 1 - I(v)$. This behavior is seen in most covering factor analyses of AGN outflows (e.g., Arav et al. 1999a; Scott et al. 2004). For the power-law model, we find that the lowest exponent is achieved at the point of lowest residual intensity. We also find that for velocities redward of the lowest residual intensity, the value of the exponent is modest, $4 < a < 9$, with errors of $\sim 20\%$. At higher velocity, the value of a climbs to roughly 20, with 50% error typical (1 σ errors based on the fitting process). For such high exponent values, the absorption behavior of the power-law model becomes increasingly similar to that of the pure covering factor model. The technical reason for the high value of a in the interval $-1450 \text{ km s}^{-1} < v < -1410 \text{ km s}^{-1}$ as compared to $-1370 \text{ km s}^{-1} < v < -1320$ can be seen in Figure 5. The intensity ratio of the strong to weak lines is considerably smaller in the $-1450 \text{ km s}^{-1} < v < -1410 \text{ km s}^{-1}$ interval compared to the $-1370 \text{ km s}^{-1} < v < -1320$ interval. Independent of a specific model, this behavior suggests a real difference in the physical behavior of the material in the two velocity intervals. Under the assumption of the power-law model, a large exponent value can be interpreted as material spread in smaller concentrated entities, but at this stage such an interpretation is tentative at best.

The modified covering factor model is an attempt to address the inherent unphysical nature of each of the above models. In the pure partial covering model, there is an unphysically sharp edge (essentially a step function) between a region with finite optical depth and a region with $\tau = 0$. Physically, there should be a transition zone between these two regions. The modified covering factor model solves this issue by adding a parameter to the partial covering model to control the gradient between these two regions. For the power-law model, the entire region where $\tau > 3$ yields virtually no flux. Since the data require steep power laws with high values of τ_{max} , in some cases most of the material comes from the $\tau > 3$ region (see discussion in Arav et al. 2005). It is disconcerting to have most of the material coming from regions that do not influence the spectrum much. The modified covering factor model allows for a gradual distribution of material, but caps τ_{max} at a physically reasonable value. However, for the QSO 2359–1241 data set τ_{max} , the weakest Fe II resonance line ($\lambda 2374$) has $\tau_{\text{max}} < 3$ across the entire velocity range, except for two data points, where $\tau_{\text{max}} < 3.5$. Therefore, there is no strong saturation at any point in front of the absorber (not even at $X = 1$, where $\tau = \tau_{\text{max}}$).

This may be the reason why the modified covering factor model gives a poorer fit to the data ($\chi_{\text{red}}^2 = 6.18$). That is, adding another degree of freedom was not necessary, since the data could be rea-

sonably fitted with a nonsaturated power law. Alternatively, if most of the χ_{red}^2 comes from systematic errors (the normalization of the effective continuum and/or oscillator strength issues discussed in § 3.2), we expect a higher χ_{red}^2 when the number of free parameters is increased. This is because while the χ^2 does not change much, the number of degrees of freedom drops considerably. In this actual case, the denominator of equation (3) changes from 65 to 40, while moving from a 2 to a 3 parameter fit, predicting an increase of 62% in χ_{red}^2 for the same χ^2 value. This is almost exactly the excess in χ_{red}^2 for this fit compared with that of the power-law model. It is therefore possible that the statistical fit of the modified covering factor model will improve enough to make it a good candidate for modeling the absorption material distribution, once the systematic errors are under better control.

3.5. Differences in Inferred Column Densities

Figure 6 shows the extracted column density of the Fe II $E = 0$ level for component e of the outflow [$N_{\text{Fe II}(E=0)}$] for all four methods. The column densities shown are the average value across the emission source at each velocity bin. For example, for the pure partial covering, we multiply the extracted column density by the covering factor [$C(v)$], since the model assumes that only $C(v)$ of the emission source is covered by the absorbing material.

Several important conclusions can be drawn from Figure 6. First, the column density derived from the τ_{ap} model is roughly a factor of 3 lower than those derived by the other methods. Not only is the τ_{ap} model inapplicable to AGN outflow troughs, it also severely underestimates the column density. Second, the inferred column density for the pure covering factor and power-law methods differ by only 25%. Although these models may suggest a rather different physical picture of the absorbing material, this quantitative similarity shows that the ionization equilibrium results based on their modeled column density are quite similar. Furthermore, this last conclusion is strengthened by the observation that not only is the total derived column density similar, but so is the distribution over velocity.

3.6. Results for Excited Fe II and He I

As we have shown in § 3.3, the power-law model yields the best fit to the Fe II resonance line data. Similar results are obtained for the five He I lines that arise from the metastable level 2^3S , as well as for the low excitation levels of the Fe II ground state. In addition to the five resonance lines (i.e., $E = 0$) that we studied so far, UV multiplets 1, 2, and 3 of Fe II contain 31 lines from four low excitation levels (with energies between 385 and 977 cm^{-1} ; see Table 1 of de Kool et al. 2001). Troughs from these lines will provide the diagnostics for finding the number density of the outflow and its ionization equilibrium (K. T. Korista et al. 2008, in preparation). In the context of this paper, these lines supply further constraints regarding which of our two main models (partial covering or power law) better represents the distribution of the absorbing material in front of the source.

In order to test which model works better and to extract the crucial column densities of these energy levels, we again used χ^2 minimization. Since all these levels arise from the same ion, their distribution in front of the source should be similar to that of the resonance lines; the only freedom we allow is the ratio of level populations, as this depends on the density and temperature of the absorber. Therefore, our procedure is to use the same $C(v)$ and $a(v)$ that we found for the partial covering or power law, respectively, and vary $\tau(v)$ and $\tau_{\text{max}}(v)$ in order to find an optimal

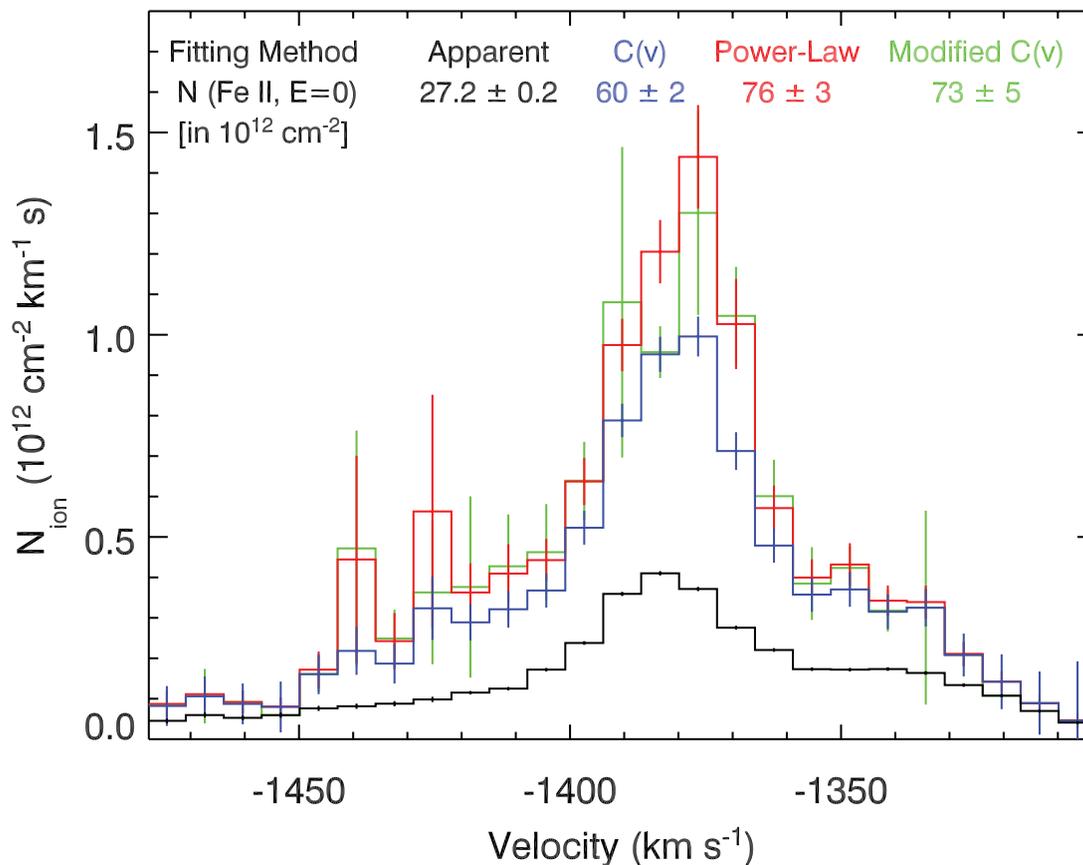


FIG. 6.—Column density of the Fe II $E = 0$ level for component e of the outflow. Shown are the best-fit determination using the four models discussed in the text. The labels at the top give the integrated column density for each of these methods and identifies each of the curves with the corresponding model.

solution. This is a robust and powerful test, since the absorption troughs for each level are a completely independent set of measurements from those of the resonance lines and from each other. In addition, we also observe five lines from the H I metastable level at $159,856 \text{ cm}^{-1}$. Although this material does not have to be distributed identically to the Fe II distribution, this is still the simplest and most constraining assumption. Therefore we use the exact same methodology on the He I* lines.

The results are shown in Table 4. It is evident that for each of the five independent Fe II levels and for the He I* lines, the χ_{red}^2 of the power-law model is lower than that of the partial covering model. This supplies strong support to the result already obtained for the resonance Fe II trough: a two-parameter power-law model fit the data better than a two-parameter partial covering model. As we discuss in § 4, the power-law distribution is also a

better physical model for the absorber in QSO 2359–1241. For these reasons, we adopt the power-law results for all the troughs observed in the object. The few exceptions are blended lines, for which we can only give lower limits.

4. DISCUSSION

Understanding the nature of quasar outflows and their influence on the nuclear and galactic environments begins with deciphering their geometry on both small and large scales. Since all of the observational information is embedded within the absorption troughs, the first step is to determine the distribution of absorbing material that creates the troughs. Advances in our understanding of this issue will shed light on the microphysics of the outflows and, more importantly, will yield reliable measurements of the ionic column densities. These measurements in turn are crucial to determining almost every physical aspect of the outflows: the ionization equilibrium and abundances, gas density, distance, mass flux, and kinetic luminosity.

Until the late 1990s, all the work published on quasar outflows (e.g., Korista et al. 1992; Turnshek et al. 1996; Hamann 1998) used the apparent optical depth (τ_{ap}) model to extract column densities. This method works very well for ISM and IGM absorbers and corresponds to a simple geometrical interpretation for the absorbing medium, a homogeneous screen in front of the light source. The well-demonstrated failure of the τ_{ap} model when applied to AGN outflow troughs (see §§ 1 and 3) motivates us to look more carefully at the geometry involved. Relatively secure is the following picture: For quasars, the size of the continuum source is $\sim 10^{16} \text{ cm}$, and the size of the broad emission line region is

TABLE 4
 χ_{red}^2 FITS FOR EXCITED Fe II AND He I LEVELS

Energy Level (cm ⁻¹) ^a	Lines	χ^2 (PC)	χ^2 (PL)
Fe II (0)	5	4.5	3.8
Fe II (385)	7	5.8	4.8
Fe II (668)	8	5.5	4.9
Fe II (863)	6	3.0	2.9
Fe II (977)	5	1.4	1.3
Fe II (7955)	4	2.4	2.3
He I (159856)	5	5.0	4.3

^a Number of lines used in the fit.

$\sim 10^{18}$ cm (Kaspi et al. 2005). Along our line of sight, there is material flowing toward us with radial velocities ranging from a few hundred to more than 30,000 km s⁻¹. This material is, cosmologically speaking, in the vicinity of the quasar, within 0.1–10,000 pc from the central source: 0.1 pc in BALQSO 1603+3002 (Arav et al. 1999a), 25 pc in NGC 3783 (Gabel et al. 2005b), 1000 pc in BALQSO 1044+3656 (de Kool et al. 2001), and 28 kpc in quasar 3C 191 (Hamann et al. 2001). The question we try to address here is: what is the distribution of the absorbing material across the quasar emission region?

Most of the work that tried to advance beyond the τ_{ap} model has used the pure covering factor method (see §§ 1 and 3). This method has three advantages. (1) It lends itself to a simple geometrical picture in which the absorbing material covers only part of the source. (2) It yields perfect fits to doublet troughs, which comprise the vast majority of high-resolution data needed to advance beyond the τ_{ap} model. However, as mentioned in § 3.2, the perfect fit arises mainly from fitting two equations with two unknowns. (3) Changes in the covering factor as a function of velocity were interpreted as accelerated motion with a small nonradial component (the dynamical-geometrical model of Arav et al. 1999b).

However, the pure covering factor model has some significant weaknesses as well. (1) It has been known for a long time that troughs of different ions show different covering factors at the same velocity (e.g., BALQSO 1603+3002; Arav et al. 1999a). Explaining such behavior involves more complicated variants of the model; chief among them is abandoning the notion of a homogeneous material distribution over the covered surface. We see a strong case of the above problem in QSO 2359–1241, in which the best-fit covering factor solution for the excited levels of Fe II is different (smaller) than the covering factor for the $E = 0$ level. (2) While the issue above can be somewhat addressed using more complicated variants of the model, the following issue poses a more serious problem. We now have several cases for which the distance of the outflow from the source was obtained by determining the number density of the outflow and combining it with the ionization equilibrium solution (knowledge of the ionization parameter). Some outflows, including the one studied here, are found at distances of ~ 1000 pc from the source and at the same time show clear evidence that the τ_{ap} model is inapplicable to their troughs (e.g., BALQSO 1044+3656; de Kool et al. 2001). At distances up to ~ 10 times the size of the emission source, the dynamical-geometrical model for the changes in the covering factor as a function of velocity can work well. However, it is very difficult to envision a scenario in which this model can work when the distance of the outflow from the emission source is 10^3 – 10^4 times the size of the emission source. We note that the covering factor model does not suffer from these difficulties for outflows at much closer distances to the nucleus (e.g., BALQSO 1603+3002 and NGC 3783, mentioned above).

Two main efforts were undertaken to address this serious flaw. One effort concentrated on finding a way to interpret the number density using a model that reduces the inferred distance by 2 or 3 orders of magnitude (Everett et al. 2002). This was done by invoking a shielded multiphase gas that is made of a continuous low-density wind with embedded high-density clouds. The model also requires that the clouds are dusty and have a strong differential dust depletion of iron compared to magnesium. While the Everett et al. (2002) model solves the covering factor problem by reducing the inferred distance substantially, this achievement comes at the heavy price of invoking several special conditions (an ionization shield, multiphase gas, and dusty clouds with strong differential dust depletion).

The second effort studied inhomogeneous distribution of absorbing material across the emission source as an alternative to the pure partial covering model. The strong feature of these models is that they do not require a finely tuned covering factor at large distances and at the same time do not invoke special physical conditions to do so. Before we discuss the physical picture of these models, we need to address the history of these models, which at first glance looks contradictory. De Kool et al. (2002a, dKKA02) showed that such models (similar to the ones used here) can yield adequate fits for two quasar outflows. In contrast, Arav et al. (2005) found that the same models could not fit the outflow troughs of the Seyfert galaxy Mrk 279. However, in the Mrk 279 data analyzed by Arav et al. (2005), more than half the flux at the spectral location of the outflow troughs is contributed by the BEL of the same transition. Gabel et al. (2005a) determined that for Mrk 279 the outflow covers the entire continuum source, but only a portion of the BEL. This picture looks physically plausible if the distance of the outflow is cospatial or just outside the BEL region. Supporting evidence on this distance scale comes from the fact that the troughs in Mrk 279 vary on timescales of ~ 1 yr. Arav et al. (2005) tried to find a simple inhomogeneous power-law model that could fit the data without differentiating between the BEL and continuum flux distributions and concluded that this is not feasible. It is not known whether more complicated inhomogeneous models that differentiate between the BEL and continuum flux sources could give a good fit to the Mrk 279 data. In contrast, for QSO 2359–1241 the contribution of the BELs at the spectral location of the Fe II troughs is less than 10%, and the distance scale is ~ 1000 pc (based on the inferred number density from Fe II* troughs combined with a determination of the ionization parameter N . (Arav et al. 2008, in preparation). The combination of these conditions permits us to ignore the contribution of the separate BELs and yields a better fit for the power-law model than the pure covering factor model does.

Finally, we need to address the fundamental question: What is the physical distribution of absorbing material in quasar outflows? In this work we have shown that a power-law distribution of material is a better model both statistically and physically than the pure covering factor model. A power-law model is of course a highly simplified version of what the real distribution might be. A hint about the actual situation can be inferred from resolved images of other astrophysical outflows. Planetary nebula, super nova remnants, and other resolved outflows tend to show a fragmented or weblike structure of the emitting material. We speculate that quasar outflows show similar structures on scales smaller than the size of the emission source. This is a straightforward way of explaining the breakdown of the homogeneous absorption screen hypothesis. It is our hope that future dynamical models will be able to yield physically motivated predictions with regard to the distribution of outflowing absorbing material. These predictions could then be tested against the QSO 2359–1241 and similar data sets.

5. SUMMARY

1. We presented 6.3 hr of VLT UVES high-resolution ($R \sim 40,000$) spectroscopic observations of QSO 2359–1241 and identified all the absorption features associated with the outflow emanating from this object.
2. The unprecedented high signal-to-noise ratio data from five unblended troughs of Fe II resonance lines yielded tight constraints on outflow trough formation models.
3. We find that power-law distribution models for absorption material in front of the emission source give a better fit to the Fe II data than covering factor models.

4. This finding alleviates the problem of obtaining a velocity-dependent partial covering factor at distances 10^3 – 10^4 times larger than the size of the emission source.

We thank the referee for numerous valuable suggestions. We acknowledge support from NSF grant AST 0507772 and from NASA LTSA grant NAG5-12867.

REFERENCES

- Arav, N. 1997, in ASP Conf. Ser. 128, *Mass Ejection from Active Galactic Nuclei*, ed. N. Arav, I. Shlosman, & R. J. Weymann (San Francisco: ASP), 208
- Arav, N., Becker, R. H., Laurent-Muehleisen, S. A., Gregg, M. D., White, R. L., Brotherton, M. S., & de Kool, M. 1999a, *ApJ*, 524, 566
- Arav, N., Brotherton, M. S., Becker, R. H., Gregg, M. D., White, R. L., Price, T., & Hack, W. 2001a, *ApJ*, 546, 140
- Arav, N., Kaastra, J., Kriss, G. A., Korista, T. K., Gabel, J., & Proga, D. 2005, *ApJ*, 620, 665
- Arav, N., Kaastra, J., Steenbrugge, K., Brinkman, B., Edelson, R., Korista, K., & de Kool, M. 2003, *ApJ*, 590, 174
- Arav, N., Korista, T. K., & de Kool, M. 2002, *ApJ*, 566, 699
- Arav, N., Korista, T. K., de Kool, M., Junkkarinen, V. T., & Begelman, M. C. 1999b, *ApJ*, 516, 27
- Arav, N., et al. 2001b, *ApJ*, 561, 118
- Ballester, P., Modigliani, A., Boitquin, O., Cristiani, S., Hanuschik, R., Kaufer, A., & Wolf, S. 2000, *Messenger*, 101, 31
- Barlow, T. A., Hamann, F., & Sargent, W. L. W. 1997, in ASP Conf. Ser. 128, *Mass Ejection from Active Galactic Nuclei*, ed. N. Arav, I. Shlosman, & R. J. Weymann (San Francisco: ASP), 13
- Blandford, R. D., & Begelman, M. C. 2004, *MNRAS*, 349, 68
- Brotherton, M. S., Arav, N., Becker, R. H., Tran, H. D., Gregg, M. D., White, R. L., Laurent-Muehleisen, S. A., & Hack, W. 2001, *ApJ*, 546, 134
- Churchill, C. W., Schneider, D. P., Schmidt, M., & Gunn, J. E. 1999, *AJ*, 117, 2573
- de Kool, M., Arav, N., Becker, R. H., Laurent-Muehleisen, S. A., White, R. L., Price, T., & Gregg, M. D. 2001, *ApJ*, 548, 609
- de Kool, M., Becker, R. H., Gregg, M. D., Arav, N., White, R. L., & Korista, K. T. 2002a, *ApJ*, 567, 58
- de Kool, M., Korista, K. T., & Arav, N. 2002b, *ApJ*, 580, 54 (dKKA02)
- Everett, J., Konigl, A., & Arav, N. 2002, *ApJ*, 569, 671
- Gabel, J. R., et al. 2005a, *ApJ*, 623, 85
- . 2005b, *ApJ*, 631, 741
- Ganguly, R., Eracleous, M. C., Charlton, J. C., & Churchill, C. W. 1999, *AJ*, 117, 2594
- Hamann, F. 1998, *ApJ*, 500, 798
- Hamann, F., Barlow, T. A., Junkkarinen, V., & Burbidge, E. M. 1997, *ApJ*, 478, 80
- Hamann, F. W., Barlow, T. A., Chaffee, F. C., Foltz, C. B., & Weymann, R. J. 2001, *ApJ*, 550, 142
- Kaspi, S., Maoz, D., Netzer, H., Peterson, B. M., Vestergaard, M., & Jannuzi, B. T. 2005, *ApJ*, 629, 61
- Korista, K. T., et al. 1992, *ApJ*, 401, 529
- . 2008, *ApJ*, submitted
- Kurucz, R., & Bell, B. 1995, CD-ROM 23, *Atomic Line Data* (Cambridge: Smithsonian Astrophys. Obs.)
- Morton, D. C. 2003, *ApJS*, 149, 205
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1989, *Numerical Recipes in Pascal: The Art of Scientific Computing* (Cambridge: Cambridge Univ. Press)
- Scannapieco, E., & Oh, S. P. 2004, *ApJ*, 608, 62
- Scott, J. E., et al. 2004, *ApJS*, 152, 1
- Telfer, R. C., Kriss, G. A., Zheng, W., Davidson, A. F., & Green, R. F. 1998, *ApJ*, 509, 132
- Turnshek, D. A., et al. 1996, *ApJ*, 463, 110
- Vernaleo, J. C., & Reynolds, C. S. 2006, *ApJ*, 645, 83
- Wampler, E. J., Chugai, N. N., & Petitjean, P. 1995, *ApJ*, 443, 586