Draft version November 28, 2011
Preprint typeset using LaTeX style emulateapj v. 11/10/09

THE FLAT TRANSMISSION SPECTRUM OF THE SUPER-EARTH GJ1214B FROM WIDE FIELD CAMERA 3 ON THE HUBBLE SPACE TELESCOPE

Zachory K. Berta[1], David Charbonneau[1], Jean-Michel Désert[1], Eliza Miller-Ricci Kempton[2,8], Peter R. McCullough[3,4], Christopher J. Burke[5], Jonathan J. Fortney[2], Jonathan Irwin[1], Philip Nutzman[2], Derek Homeier[6,7]

Draft version November 28, 2011

ABSTRACT

Capitalizing on the observational advantage offered by its tiny M dwarf host, we present HST/WFC3 grism measurements of the transmission spectrum of the super-Earth exoplanet GJ1214b. These are the first published WFC3 observations of a transiting exoplanet atmosphere. After correcting for a ramp-like instrumental systematic, we achieve nearly photon-limited precision in these observations, finding the transmission spectrum of GJ1214b to be flat between 1.1 and 1.7 μm. Inconsistent with a cloud-free solar composition atmosphere at 8.2σ, the measured achromatic transit depth most likely implies a large mean molecular weight for GJ1214b's outer envelope. A dense atmosphere rules out bulk compositions for GJ1214b that explain its large radius by the presence of a very low density gas layer surrounding the planet. High-altitude clouds can alternatively explain the flat transmission spectrum, but they would need to be optically thick up to 10 mbar or consist of particles with a range of sizes approaching 1 μm in diameter.

Subject headings: planetary systems: individual (GJ 1214b) — eclipses — techniques: spectroscopic

1. INTRODUCTION

With a radius of 2.7 $R_\oplus$ and a mass of 6.5 $M_\oplus$, the transiting planet GJ1214b (Charbonneau et al. 2009) is a member of the growing population of exoplanets whose masses and radii are known to be between those of Earth and Neptune (see Léger et al. 2009; Batalha et al. 2011; Lissauer et al. 2011; Winn et al. 2011). Among these exoplanets, most of which exhibit such shallow transits that they require ultra-precise space-based photometry simply to detect the existence of their transits, GJ1214b is unique. The diminutive 0.21 $R_\odot$ radius of its M dwarf stellar host means GJ1214b exhibits a large 1.4% transit depth, and the system's proximity (13 pc) means the star is bright enough in the near infrared ($H = 9.1$) that follow-up observations to study the planet's atmosphere are currently feasible. In this work, we exploit this observational advantage and present new measurements of the planet's atmosphere, which bear upon models for its interior composition and structure.

According to theoretical studies (Seager et al. 2007; Rogers & Seager 2010; Nettelmann et al. 2011), GJ1214b's 1.9 g cm^{-3} bulk density is high enough to require a larger ice or rock core fraction than the solar system ice giants but far too low to be explained with an entirely Earth-like composition. Rogers & Seager (2010) have proposed three general scenarios consistent with GJ1214b's large radius, where the planet could (i) have accreted and maintained a nebular H_2/He envelope atop an ice and rock core, (ii) consist of a rocky planet with an H_2-rich envelope that formed by recent outgassing, or (iii) contain a large fraction of water in its interior surrounded by a dense H_2-depleted, H_2O-rich atmosphere. Detailed thermal evolution calculations by Nettelmann et al. (2011) disfavor this last model on the basis that it would require unreasonably large bulk water-to-rock ratios, arguing for at least a partial H_2/He envelope, albeit one that might be heavily enriched in H_2O relative to the primordial nebula.

By measuring GJ1214b's transmission spectrum, we can empirically constrain the mean molecular weight of the planet's atmosphere, thus distinguishing among these possibilities. When the planet passes in front of its host M dwarf, a small fraction of the star's light passes through the upper layers of the planet's atmosphere before reaching us; the planet's transmission spectrum is then manifested in variations of the transit depth as a function of wavelength. The amplitude of the transit depth variations $\Delta D(\lambda)$ in the transmission spectrum scale as $n_H \times 2HR_p/R_\star^2$, where n_H is set by the opacities involved and can be 1-10 for strong absorption features, H is the atmospheric scale height, R_p is the planetary radius, and $R_\star$ is the stellar radius (e.g. Seager & Sasselov 2000; Brown 2001; Hubbard et al. 2001). Because the scale height H is inversely proportional to the mean molecular weight μ of the atmosphere, the amplitude of features seen in the planet's transmission spectrum places strong constraints on the possible values of μ and, in particular, the hydrogen/helium content of the atmosphere (Miller-Ricci et al. 2009).

zberta@cfa.harvard.edu
[1] Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138, USA
[2] Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA
[3] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
[4] Smithsonian Astrophysical Observatory, 60 Garden St., Cambridge, MA 02138, USA
[5] SETI Institute/NASA Ames Research Center, M/S 244-30, Moffett Field, CA 94035, USA
[6] Centre de Recherche Astrophysique de Lyon, UMR 5574, CNRS, Université de Lyon, École Normale Supérieure de Lyon, 46 Allée d'Italie, F-69364 Lyon Cedex 07, France
[7] Institut für Astrophysik, George-August-Universität, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany
[8] Sagan Fellow

Indeed, detailed radiative transfer simulations of GJ1214b's atmosphere (Miller-Ricci & Fortney 2010) show that a solar composition, H_2-dominated atmosphere ($\mu = 2.4$) would show depth variations of roughly 0.1% between 0.6 and 10 μm, while the features in an H_2O-dominated atmosphere ($\mu = 18$) would be an order of magnitude smaller. While the latter of these is likely too small to detect directly with current instruments, the former is at a level that has regularly been measured with the Hubble Space Telescope (HST) in the transmission spectra of hot Jupiters (e.g. Charbonneau et al. 2002; Pont et al. 2008; Sing et al. 2011).

Spectroscopic observations by Bean et al. (2010) with the Very Large Telescope found the transmission spectrum of GJ1214b to be featureless between 0.78-1.0 μm, down to an amplitude that would rule out cloud-free H_2-rich atmospheric models. Broadband Spizer Space Telescope photometric transit measurements at 3.6 and 4.5 μm by Désert et al. (2011a) showed a flat spectrum consistent with Bean et al. (2010), as did high-resolution spectroscopy with NIRSPEC between 2.0 and 2.4 μm by Crossfield et al. (2011). Intriguingly, the transit depth in K-band (2.2 μm) was measured from CFHT by Croll et al. (2011) to be 0.1% deeper than at other wavelengths, which would imply a H_2-rich atmosphere, in apparent contradiction to the other studies.

These seemingly incongruous observations could potentially be brought into agreement if GJ1214b's atmosphere were H_2-rich but significantly depleted in CH_4 (Crossfield et al. 2011; Miller-Ricci Kempton et al. 2011). In such a scenario, the molecular features that remain (predominantly H_2O) would fit the CFHT measurement, but be unseen by the NIRSPEC and Spitzer observations. Explaining the flat VLT spectrum in this context would then require a broadband haze to smooth the spectrum at shorter wavelengths (see Miller-Ricci Kempton et al. 2011). New observations by Bean et al. (2011) covering 0.6-0.85 μm and 2.0-2.3 μm were again consistent with a flat spectrum, but they still could not directly speak to this possibility of a methane-depleted, H_2-rich atmosphere with optically scattering hazes.

Here, we present a new transmission spectrum of GJ1214b spanning 1.1 to 1.7 μm, using the infrared slitless spectroscopy mode on the newly installed Wide Field Camera 3 (WFC3) aboard the Hubble Space Telescope (HST). Our WFC3 observations directly probe the predicted strong 1.15 and 1.4 μm water absorption features in GJ1214b's atmosphere (Miller-Ricci & Fortney 2010) and provide a stringent constraint on the H_2 content of GJ1214b's atmosphere that is robust to non-equilibrium methane abundances and hence a definite test of the CH_4-depleted hypothesis. The features probed by WFC3 are the same features that define the J and H band windows in the telluric spectrum, and cannot be observed from the ground.

Because this is the first published analysis of WFC3 observations of a transiting exoplanet, we include a detailed discussion of the performance of WFC3 in this observational regime and the systematic effects that are inherent to the instrument. Recent work on WFC3's predecessor NICMOS (Burke et al. 2010; Gibson et al. 2011a) has highlighted the importance of characterizing instrumental systematics when interpreting exoplanet results from HST observations.

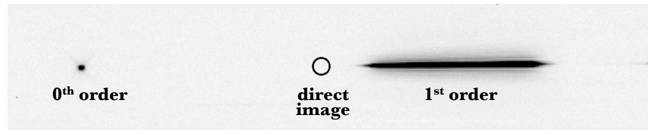

Figure 1. A 512x100 pixel cutout of a typical WFC3 G141 grism exposure of the star GJ1214. The $0^{\rm th}$ and $1^{\rm st}$ order spectra are labeled, and the start of the $2^{\rm nd}$ order spectrum is visible on the right. The location of the star in the direct images (not shown here) is marked with a circle.

This paper is organized as follows: we describe our observations in §2, our method for extracting spectrophotometric light curves from them in §3, and our analysis of these light curves in §4. We present the resulting transmission spectrum and discuss its implications for GJ1214b's composition in §5, and conclude in §6.

2. OBSERVATIONS

We observed three transits of GJ1214b on UT 2010 October 8, 2011 March 28, and 2011 July 23 with the G141 grism on WFC3's infrared channel (HST Proposal #GO-12251, P.I. = Z. Berta), obtaining simultaneous multiwavelength spectrophotometry of each transit between 1.1 and 1.7 μm. WFC3's IR channel consists of a 1024×1024 pixel Teledyne HgCdTe detector with a 1.7 μm cutoff that can be paired with any of 15 filters or 2 low-resolution grisms (Dressel et al. 2010). Each exposure is compiled from multiple non-destructive readouts and can consist of either the full array or a concentric, smaller subarray.

Each visit consisted of four 96 minute long HST orbits, each containing 45 minute gaps due to Earth occultations. Instrumental overheads between the occultations are dominated by serial downloads of the WFC3 image buffer, during which all science images are transferred to the telescope's solid state recorder. This buffer can hold only two 16-readout, full-frame IR exposures before requiring a download, which takes 6 minutes. Exposures cannot be started nor stopped during a buffer download, so parallel buffer downloads are impossible for short exposures.

Subject to these constraints and the possible readout sequences, we maximized the number of photons detected per orbit while avoiding saturation by gathering exposures using the 512×512 subarray with the RAPID NSAMP=7 readout sequence, for an effective integration time of 5.971 seconds per exposure. With this setting, four 12-exposure batches, separated by buffer downloads, were gathered per orbit resulting in an integration efficiency of 10%. Although the brightest pixel in the $1^{\rm st}$ order spectrum reaches 78% of saturation during this exposure time, the WFC3's multiple non-destructive readouts enable the flux within each pixel to be estimated before the onset of significant near-saturation nonlinearities.

A sub-region of a typical G141 grism image of GJ1214 is shown in Fig. 1. The 512×512 subarray allows both the $0^{\rm th}$ and $1^{\rm st}$ order spectra to be recorded, and the $1^{\rm st}$ order spectrum to fall entirely within a single amplifier quadrant of the detector. The $1^{\rm st}$ order spectrum spans 150 pixels with a dispersion in the x-direction of 4.65 nm/pixel and a spatial full-width half maximum in the y-direction of 1.7 pixels (0.2"). The $0^{\rm th}$ order spectrum is slightly dispersed by the grism's prism but is nearly a point source. Other stars are present in the

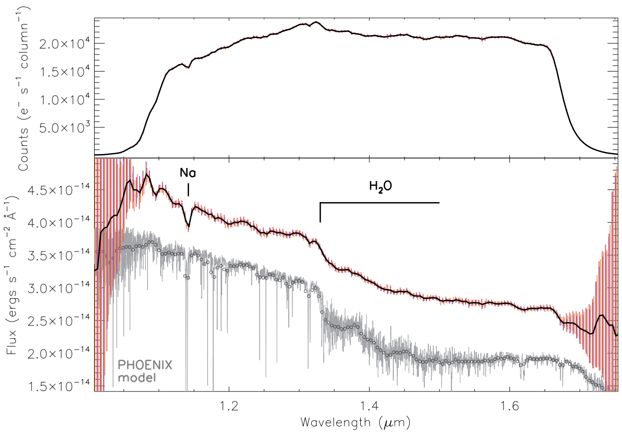

Figure 2. The mean out-of-transit extracted spectrum of GJ1214 (*black line*) from all three HST visits, shown before (top) and after (bottom) flux calibration. Individual extracted spectra from each visit are shown with their 1σ uncertainties (*color error bars*). For comparison, the integrated flux from the PHOENIX model atmosphere used to calculate the stellar limb darkening (see §4.2) is shown (*gray lines*) offset for clarity and binned to the WFC3 pixel scale (*gray circles*).

subarray's $68' \times 61'$ field of view, but are too faint to provide useful diagnostics of systematic trends that may exist in the data. For wavelength calibration, we gathered direct images in the F130N narrow-band filter; the direct images' position relative to the grism images is also shown in Fig. 1.

To avoid systematics from the detector flat-fields that have a quoted precision no better than 0.5% (Pirzkal et al. 2011), the telescope was not dithered during any of the observations. We note that a technique called "spatial scanning" has been proposed to decrease the overheads for bright targets with WFC3, where the telescope nods *during* an exposure to smear the light along the cross-dispersion direction, thus increasing the time to saturation (McCullough & MacKenty 2011). We did not use this mode of observation as it was not yet tested at the time our program was initiated.

3. DATA REDUCTION

The Python/PyRAF software package aXe was developed to extract spectra from slitless grism observations with WFC3 and other Hubble instruments (Kümmel et al. 2009), but it is optimized for extracting large numbers of spectra from full frame dithered grism images. To produce relative spectrophotometric measurements of our single bright source, we opted to create our own extraction pipeline that prioritizes precision in the time domain. We outline the extraction procedure below.

Through the extraction, we use calibrated 2-dimensional images, the "`flt`" outputs from WFC3's `calwf3` pipeline. For each exposure, `calwf3` performs the following steps: flag detector pixels with the appropriate data quality (DQ) warnings, estimate and remove bias drifts using the reference pixels, subtract dark current, determine count rates and identify cosmic rays by fitting a slope to the non-destructive reads, correct for photometric non-linearity (properly accounting for the signal accumulation before the initial "zeroth" read), and apply gain calibration. The resulting images are measured in $e^- s^{-1}$ and contain per pixel uncertainty estimates based on a detector model (Kim Quijano et al. 2009). We note that `calwf3` does *not* apply flat-field corrections when calibrating grism images; proper wavelength-dependent flat-fielding for slitless spectroscopy requires wavelength-calibrating individual sources and `calwf3` does not perform this task.

3.1. *Interpolating over Cosmic Rays*

`calwf3` identifies cosmic rays that appear partway through an exposure by looking for deviations from a linear accumulation of charge among the non-destructive readouts, but it can not identify cosmic rays that appear between the zeroth and first readout. We supplement `calwf3`'s cosmic ray identifications by also flagging any pixel in an individual exposure that is $> 6\sigma$ above the median of that pixel's value in all other exposures as a cosmic ray. Through all three visits (576 exposures), a total of 88 cosmic rays were identified within the extraction box for the 1$^{\rm st}$ order spectra.

For each exposure, we spatially interpolate over cosmic rays. Near the 1$^{\rm st}$ order spectrum, the pixel-to-pixel gradient of the point spread function (PSF) is typically much shallower along the dispersion direction than perpendicular to it, so we use only horizontally adjacent pixels when interpolating to avoid errors in modeling the sharp cross-dispersion falloff.

3.2. *Identifying Continuously Bad Pixels*

We also mask any pixels that are identified as "bad detector pixels" (DQ=4), "unstable response" (DQ=32), "bad or uncertain flat value" (DQ=512). We found that only these DQ flags affected the photometry in a pixel by more than 1σ. Other flags may have influenced the pixel photometry, but did so below the level of the photon noise. In the second visit, we also identified one column of the detector ($x = 625$ in physical pixels[9]) whose light curve exhibited a dramatically different systematic variation than did light curves from any of the other columns. This column was coincident with an unusually low-sensitivity feature in the flat-field, and we hypothesize that the flat-field is more uncertain in this column than in neighboring columns. We masked all pixels in that column as bad.

We opt not to interpolate over these continuously bad pixels. Because they remained flagged throughout the duration of each visit, we simply give these pixels zero weight when extracting 1D spectra from the images. This allows us to keep track of the actual number of photons recorded in each exposure so we can better assess our predicted photometric uncertainties.

3.3. *Background Estimation*

In addition to the target, WFC3 also detects light from the diffuse sky background, which comes predominantly from zodiacal light and Earth-shine, and must be subtracted. We draw conservative masks around all sources that are visible in each visit's median image, including GJ1214 and its electronics cross-talk artifact (see Viana & Baggett 2010). We exclude these pixels, as well as all

[9] For ease of comparison with future WFC3 analyses, throughout this paper we quote all pixel positions in physical units as interpreted by SAOImage DS9, where the bottom left pixel of a full-frame array would be $(x, y) = (1,1)$.

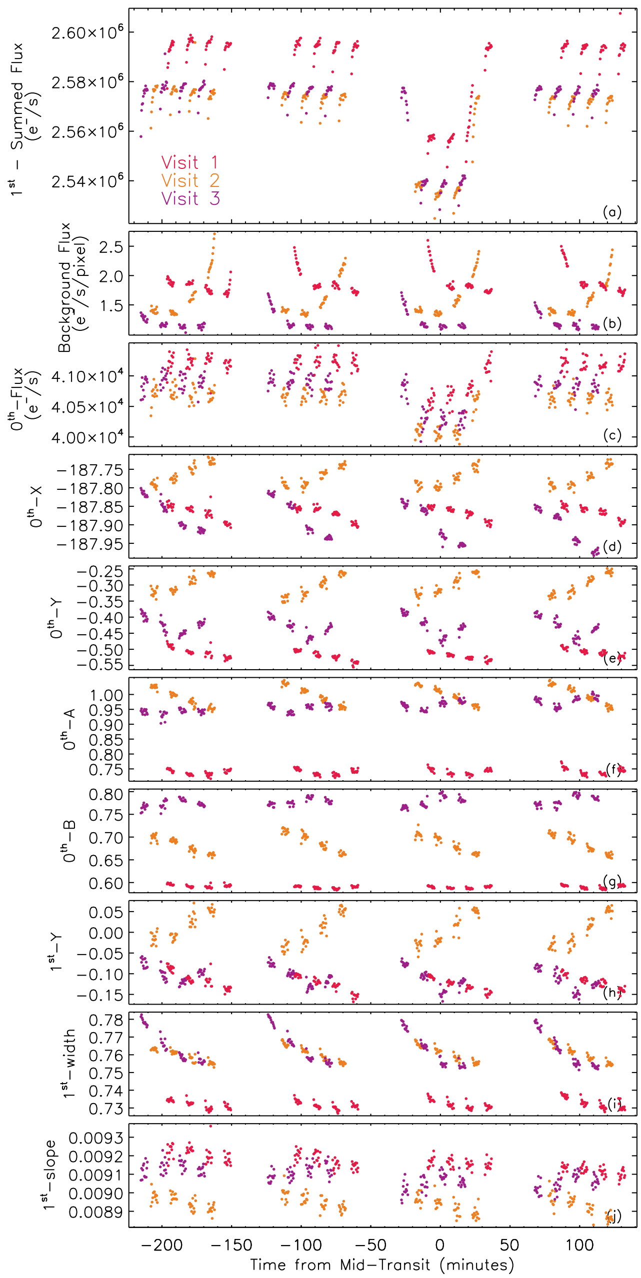

Figure 3. Extracted properties of the $0^{\rm th}$ and $1^{\rm st}$ order spectra as a function of time, including (a) the summed $1^{\rm st}$ order photometric light curves; (b) the estimated sky background level; (c-g) the total flux, x and y position (measured relative to the reference pixel), and Gaussian widths in the x and y directions of the $0^{\rm th}$ order image; and (h-j) the y offset, cross-dispersion width, and slope of the $1^{\rm st}$ order spectrum. All three visits are shown and are denoted by the color of the symbols. The 45 minute gaps in each time series are due to Earth occultations, the 6 minute gaps are due to the WFC3 buffer downloads.

pixels that have any DQ warning flagged. Then, to estimate the sky background in each exposure, we scale a master WFC3 grism sky image (Kümmel et al. 2011) to match the remaining 70-80% of the pixels in each exposure and subtract it. We find typical background levels of $1-3$ e^{-} s^{-1} pixel^{-1}, that vary smoothly within orbits and throughout visits as shown in Fig. 3 (panel b). As a test, we also estimated the background level from a simple mean of the unmasked pixels; the results were unchanged.

3.4. *Inter-pixel Capacitance*

The normal calibration pipeline does not correct for the inter-pixel capacitance (IPC) effect, which effectively couples the flux recorded in adjacent pixels at about the 1% level (McCullough 2008). We correct this effect with a linear deconvolution algorithm (McCullough 2008; Hilbert & McCullough 2011), although we find it makes little difference to the final results.

3.5. *Extracting the Zeroth Order Image*

The $0^{\rm th}$ order image can act as a diagnostic for tracking changes in the telescope pointing and in the shape of the instrumental PSF. We select a 10×10 pixel box around the $0^{\rm th}$ order image and fit a 2D Gaussian to it with the x position, y position, size in the x direction, size in the y direction, and total flux allowed to vary (5 parameters).

Time series of the $0^{\rm th}$ order x and y positions, sizes in both directions, and total flux are shown for all three visits in Fig. 3 (panels c-g). Thanks to the dispersion by the grism's prism, the Gaussian is typically 20% wider in the x direction than in the y direction. Even though the throughput of the $0^{\rm th}$ order image is a factor of 60 lower than the $1^{\rm st}$ order spectrum, the transit of GJ1214b is readily apparent in the $0^{\rm th}$-order flux time series.

3.6. *Extracting the First Order Spectrum*

To extract the first order spectra, we first determine the position of GJ1214 in the direct image, which serves as a reference position for defining the trace and wavelength calibration of the $1^{\rm st}$ order spectrum. We adopt the mean position GJ1214 in all of the direct images as the reference position, which we measure using the same method as in extracting the $0^{\rm th}$ order image in §3.5. The measured (x, y) reference positions for the first, second, and third visits are (498.0, 527.5), (498.6, 531.1), and (498.9, 527.1) in physical pixels.

Once the reference pixel for a visit is known, we use the coefficients stored in the WFC3/G141 aXe configuration file[10] (Kuntschner et al. 2009), to determine the geometry of the $1^{\rm st}$ order trace and cut out a 30 pixel tall extraction box centered on the trace. Within this extraction box, we use the wavelength calibration coefficients to determine the average wavelength of light that will be illuminating each pixel. We treat all pixels in the same column as having the same effective wavelength; given the spectrum's 0.5° tilt from to the x axis, errors introduced by this simplification are negligible.

Kuntschner et al. (2008) used flat-fields taken through all narrow-band filters available on WFC3/IR to construct a flat-field "cube" where each pixel contains 4 polynomial coefficients that describe its sensitivity as a function of wavelength[11]. We use this flat-field cube to construct a color-dependent flat based on our estimate of the effective wavelength illuminating each pixel, and divide each exposure by it. WFC3 wavelength calibration and flat-fielding is described in detail in the aXe manual (Kümmel et al. 2010).

[10] The aXe configuration file `WFC3.IR.G141.V2.0.conf` is available through `http://www.stsci.edu/hst/wfc3/`

[11] `WFC3.IR.G141.flat.2.fits`, through the same URL

To calculate 1D spectra from the flat-fielded images, we sum all the unmasked pixels within the extraction box over the y-axis. To estimate the uncertainty in each spectral channel, we first construct a per-pixel uncertainty model that includes photon noise from the source and sky as well as 22 e^- of read noise, and sum these uncertainties, in quadrature, over the y-axis. We do not use the `calwf3`-estimated uncertainties; they include a term propagated from the uncertainty in the nonlinearity correction that, while appropriate for absolute photometry, would not be appropriate for relative photometry. In each exposure, there are typically 1.2×10^5 e^- per single-pixel spectral channel and a total of 1.5×10^7 e^- in the entire spectrum. Fig. 3 (panel a) shows the extracted spectra summed over all wavelengths as a function of time, the "white" light curve.

For diagnostics' sake, we also measure the geometrical properties of the 1^{st} order spectra in each exposure. We fit 1D Gaussians the cross-dispersion profile in each column of the spectrum and take the median Gaussian width among all the columns as a measurement of the PSF's width. We fit a line to the location of the Gaussian peaks in all the columns, taking the intercept and the slope of that line as an estimate of the y-offset and tilt of the spectrum on the detector. Time series of these parameters are shown in Fig. 3 (panels h-j).

3.7. *Flux Calibration*

For the sake of display purposes only (see Fig. 2), we flux calibrate each visit's median, extracted, 1D spectrum. Here we have interpolated over all bad pixels within each visit (contrary to the discussion in the §3.2), and plotted the weighted mean over all three visits. The calibration uncertainty for the G141 sensitivity curve (Kuntschner et al. 2011) is quoted to be 1%.

3.8. *Times of Observations*

For each exposure, we extract the `EXPSTART` keyword from the science header, which is the Modified Julian Date at the start of the exposure. We correct this to the mid-exposure time using the `EXPTIME` keyword, and convert it to the Barycentric Julian Date in the Barycentric Dynamical Time standard using the code provided by Eastman et al. (2010).

4. ANALYSIS

In this section we describe our method for estimating parameter uncertainties (§4.1) and our strategy for modeling GJ1214's stellar limb darkening (§4.2). Then, after identifying the dominant systematics in WFC3 light curves (§4.3) and describing a method to correct them (§4.4), we present our fits to the light curves, both summed over wavelength (§4.5) and spectroscopically resolved (§4.6). We also present a fruitless search for transiting satellite companions to GJ1214b (§4.7).

4.1. *Estimating Parameter Distributions*

Throughout our analysis, we fit different WFC3 light curves with models that have different sets of parameters, and draw conclusions from the inferred probability distributions of those parameters; this section describes our method for characterizing the posterior probability distribution for a set of parameters within a given model.

We use a Markov Chain Monte Carlo (MCMC) method with the Metropolis-Hastings algorithm to explore the posterior probability density function (PDF) of the model parameters. This Bayesian technique allows us to sample from (and thus infer the shape of) the probability distribution of a model's parameters given both our data and our prior knowledge about the parameters (for reviews, see Ford 2005; Gregory 2005; Hogg et al. 2010). Briefly, the algorithm starts a chain with an initial set of parameters ($\mathbf{M}_{j=0}$) and generates a trial set of parameters ($\mathbf{M'}_{j+1}$) by perturbing the previous set. The ratio of posterior probability between the two parameter sets, given the data $\mathbf{D}$, is then calculated as

$$\frac{P(\mathbf{M'}_{j+1}|\mathbf{D})}{P(\mathbf{M}_j|\mathbf{D})} = \frac{P(\mathbf{D}|\mathbf{M'}_{j+1})}{P(\mathbf{D}|\mathbf{M}_j)} \times \frac{P(\mathbf{M'}_{j+1})}{P(\mathbf{M}_j)} \quad (1)$$

where the first term (the "likelihood") accounts for the information that our data provide about the parameters and the second term (the "prior") specifies our externally conceived knowledge about the parameters. If a random number drawn from a uniform distribution between 0 and 1 is less than this probability ratio, then $\mathbf{M}_{j+1}$ is set to $\mathbf{M'}_{j+1}$; if not, then $\mathbf{M}_{j+1}$ reverts to $\mathbf{M}_j$. The process is iterated until j is large, and the resulting chain of parameter sets is a fair sample from the posterior PDF and can be used to estimate confidence intervals for each parameter.

To calculate the likelihood term in Eq. 1, we assume that each of the N flux values d_i is drawn from a uncorrelated Gaussian distribution centered on the model value m_i with a standard deviation of $s\sigma_i$, where σ_i is the theoretical uncertainty for the flux measurement based on the detector model and photon statistics and s is a photometric uncertainty rescaling parameter. Calculation of the ratio in Eq. 1 is best done in logarithmic space for numerical stability, so we write the likelihood as

$$\ln P(\mathbf{D}|\mathbf{M}) = -N \ln s - \frac{1}{2s^2}\chi^2 + \text{constant} \quad (2)$$

where

$$\chi^2 = \sum_{i=1}^{N} \left(\frac{d_i - m_i}{\sigma_i} \right)^2 \quad (3)$$

and we have only explicitly displayed terms that depend on the model parameters. Including s as a model parameter is akin to rescaling the uncertainties by externally modifying σ_i to achieve a reduced χ^2 of unity, but enables the MCMC to fit for and marginalize over this rescaling automatically. Unless otherwise stated for specific parameters, we use non-informative (uniform) priors for the second term in Eq. 1. We use a Jeffreys prior on s (uniform in $\ln s$) which is the least informative, although the results are practically indistinguishable from prior uniform in s.

When generating each new trial parameter set $\mathbf{M'}_{j+1}$, we follow Dunkley et al. (2005) and perturb every parameter at once, drawing the parameter jumps from a multivariate Gaussian with a covariance matrix that approximates that of the parameter distribution. Doing so allows the MCMC to move easily along the dominant linear correlations in parameter PDF, and greatly increases the efficiency of the algorithm. While this procedure may

Table 1
White Light Curves from WFC3/G141

Time[a] (BJD_{TDB})	Relative Flux[b]	Uncertainty	Sky (e^-/s)	0^{th}-X[c] (pix)	0^{th}-Y[c] (pix)	0^{th}-A[d] (pix)	0^{th}-B[d] (pix)	1^{st}-Y[c] (pix)	1^{st}-B[d] (pix)	1^{st}-Slope[e] (pix/pix)	Visit
2455478.439980	0.99381	0.00031	1.9546	-187.830	-0.474	0.790	0.613	-0.075	0.7484	0.00921	1
2455478.440270	0.99713	0.00031	1.9938	-187.839	-0.481	0.792	0.616	-0.082	0.7490	0.00925	1
2455478.440559	0.99787	0.00031	1.9718	-187.846	-0.484	0.789	0.615	-0.076	0.7486	0.00914	1
2455478.440848	0.99958	0.00032	1.8808	-187.827	-0.490	0.792	0.614	-0.087	0.7476	0.00917	1
2455478.441138	0.99989	0.00032	1.9379	-187.844	-0.490	0.785	0.612	-0.085	0.7492	0.00921	1
				...							

Note. — This table is published in its entirety in the electronic edition of the Astrophysical Journal. A portion is shown here for guidance concerning its form and content.

[a] Mid-exposure time.
[b] Normalized to the median flux level of the out-of-transit observations in each visit.
[c] Position measured relative to the Gaussian center of each visit's direct image.
[d] Gaussian width of the 0^{th} or 1^{st} order spectra in the horizontal (A) or vertical (B) direction.
[e] Slope of the 1^{st} order spectrum.

seem circular (if we knew the covariance matrix of the parameter distribution, why would we need to perform the MCMC?), the covariance matrix we use to generate trial parameters could be a very rough approximation to the true shape of the parameter PDF but still dramatically decrease the computation time necessary for the MCMC.

To obtain an initial guess for parameters ($\mathbf{M_{j=0}}$), we use the `MPFIT` implementation (Markwardt 2009) of the Levenberg-Marquardt (LM) method to maximize $\ln P(\mathbf{M}|\mathbf{D})$. This would be identical to minimizing χ^2 in the case of flat priors, but it can also include constraints from more informative priors. The LM fit also provides an estimate of the covariance matrix of the parameters, which is a linearization of the probability space near the best-fit. We use this covariance matrix estimate for generating trial parameters in the MCMC, and with it, achieve parameter acceptance rates of 10-40% throughout the following sections. As expected, when fitting models with flat priors and linear or nearly-linear parameters (where the PDF should well-described by a multivariate Gaussian), the LM covariance matrix is identical to that ultimately obtained from the MCMC (see Sivia 1996, for further discussion).

MCMC chains are run until they contain 1.25×10^5 points. The first 1/5 of the points are ignored as "burn-in", leaving 1×10^5 for parameter estimation. Correlation lengths for the parameters in the MCMC chains are indicated throughout the text; they are typically of order 10 points. A chain with such a correlation length effectively contains $1 \times 10^5/10 = 1 \times 10^4$ independent realizations of the posterior PDF. We quote confidence intervals that exclude the upper and lower 16% of the marginalized distribution for each parameter (i.e. the parameter's central 68% confidence interval), using all 1×10^5 points in each chain.

4.2. *Modeling Stellar Limb Darkening*

Accurate modeling of the WFC3 integrated and spectroscopic transit light curves requires careful consideration of the stellar limb-darkening (LD) behavior. GJ1214b's M4.5V stellar host is so cool that it exhibits weak absorption features due to molecular H_2O. Because inferences of the planet's apparent radii from transit light curves depend strongly on the star's limb-darkening, which is clearly influenced by H_2O as an opacity source, inaccurate treatment of limb-darkening could potentially introduce spurious H_2O features into the transmission spectrum.

If they were sufficiently precise, transit light curves alone could simultaneously constrain both the star's multiwavelength limb-darkening behavior and the planet's multiwavelength radii (e.g. Knutson et al. 2007b). For less precise light curves, it is common practice to fix the limb-darkening to a theoretically calculated law, even if this may underestimate the uncertainty in the planetary parameters (see Burke et al. 2007; Southworth 2008). Given the quality of our data, we adopt an intermediate solution where we allow the limb-darkening parameters to vary in our fits, but with a Gaussian prior centered on the theoretical values (e.g. Bean et al. 2010).

We model the star GJ1214's limb-darkening behavior with a spherically symmetric `PHOENIX` atmosphere (Hauschildt et al. 1999), assuming stellar parameters of $T_{\rm eff} = 3026K$, $\log g = 5$, and [M/H] = 0 (Charbonneau et al. 2009). As shown in Fig. 2, the integrated flux from the PHOENIX model is in good qualitative agreement with the low-resolution, calibrated WFC3 stellar spectrum of GJ1214. From this model, we calculate photon-weighted average intensity profiles for the integrated spectrum and for each of the individual wavelength bins, using the WFC3 grism sensitivity curve and the `PHOENIX` model to estimate the photon counts. In the spherical geometry of the PHOENIX atmospheres the characterization of the actual limb (defined as $\mu = 0$, see below) is not straightforward, as the model extends beyond the photosphere into the optically thin outer atmosphere. The result is an approximately exponentially declining intensity profile from the outermost layers, that Claret & Hauschildt (2003) found not to be easily reproduced by standard limb darkening laws for plane-parallel atmospheres. These authors suggest the use of "quasi-spherical" models by ignoring the outer region. In an extension of this concept, we set the outer surface of the star to be where the intensity drops to e^{-1} of the central intensity, and measure $\mu = \cos\theta$ (where θ is the emission angle relative to the line of sight) relative to that outer radius.

We derive coefficients for a square-root limb-darkening

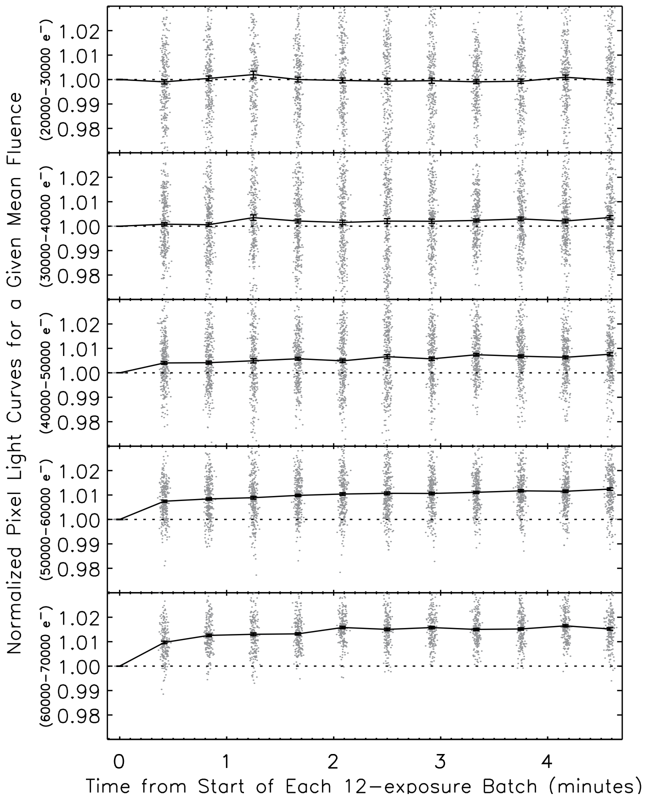

Figure 4. Single-pixel light curves within each 12-exposure batch following a buffer download (*gray points*), shown for different mean pixel illuminations. Pixel light curves have been normalized to the first exposure within each batch, and plotted with small random horizontal offsets for clarity. Only data from the first HST visit, which exhibited the smallest pointing drifts (see Fig. 3), are shown. Error bars show the mean and its standard error for each time point and each illumination. An exponential ramp begins is present in pixels with a mean recorded fluence greater than 30,000 and 40,000 e^- (50% of the detector full well). Note, the nominal fluences quoted here do not include charge accumulated during detector flushing and initial readout (see text).

law for each of these average intensity profiles using least-squares fitting. In this law, the intensity relative to the center of the star is given by

$$\frac{I(\mu)}{I(1)} = 1 - c(1-\mu) - d(1-\sqrt{\mu}), \tag{4}$$

where c and d are the two coefficients of the fit. We chose a square-root law over the popular quadratic law because it gave noticeably better approximations to the PHOENIX intensity profiles, while still having few enough free parameters that they can be partially inferred from the data. Indeed, van Hamme (1993) found the square-root law to be generally preferable to other 2-parameter limb-darkening laws for late-type stars in the near-IR. The square-root law matches the theoretical intensity profile nearly as well as the full nonlinear 4-parameter law introduced by Claret (2000) for the models we use here.

4.3. *Light Curve Systematics*

The summed light curve shown in Fig. 3 (panel a) exhibits non-astrophysical systematic trends. The most obvious of these are the sharply rising but quickly saturating "ramp"-like features within each batch of 12 exposures between buffer downloads. To the eye, the ramps are very repeatable; the flux at the end of all batches asymptotes to nearly the same level. The amplitude of the ramp is 0.4% from start to finish for most batches, except for the first batch of each orbit, where the ramp is somewhat less pronounced.

These ramps are reminiscent of those seen in high-cadence Spitzer light curves at 8 and 16 μm (e.g. Deming et al. 2006; Knutson et al. 2007a; Charbonneau et al. 2008) which Agol et al. (2010) recently proposed may be due to "charge trapping" within the detector pixels. In their toy model, charge traps within each pixel become filled throughout an exposure and later release the trapped charge on a finite timescale, thereby increasing the pixel's dark current in subsequent exposures. The model leads to exponential ramps when observing bright sources as the excess dark current increases sharply at first but slows its increase as the population of charge traps begins to approach steady state. We note this model also leads to after-images following strong exposures, i.e., persistence.

WFC3 has been known since its initial ground-testing to exhibit strong persistence behavior (McCullough & Deustua 2008; Long et al. 2010). Smith et al. (2008a) have proposed that persistence in 1.7 μm cutoff HgCdTe detectors like WFC3 is likely related to charge trapping. Measurements (McCullough & Deustua 2008) indicate that WFC3's persistence may be of the right order of magnitude (on $<$ 1 minute timescales) to supply the roughly 50 e^- s^{-1} pixel^{-1} in the brightest pixels that would be necessary to explain the observed several millimagnitude ramp, although persistence levels and decay timescales can depend in complicated ways on the strength of previous exposures (see Smith et al. 2008b).

We were aware of this persistence issue before our observations and made an effort to control its effect on our light curves. When we planned the timing of the exposures, we attempted to make the illumination history of each pixel as consistent as possible from batch to batch and orbit to orbit. In practice, this means we gathered more direct images than necessary for wavelength calibration to delay some of the grism exposures.

Whether or not the ramps are caused by the charge trapping mechanism, they are definitely dependent on the illumination that a pixel receives. To demonstrate this, we construct light curves for each individual pixel over the duration of every out-of-transit 12-exposure batch that follows a buffer download and normalize each of these pixel light curves to the first exposure in the batch. Fig. 4 shows the normalized pixel light curves, grouped by their mean recorded fluence. Because it takes a finite time to read the subarray (0.8 seconds) and reset the full array (2.9 seconds), we note that each exposure actually collects 60% more electrons than indicated by these nominal, recorded fluences (see Long et al. 2011). The appearance of the ramp clearly becomes more pronounced for pixels that are more strongly exposed.

Buried beneath the ramp features, the summed light curve exhibits subtler trends that appear mostly as orbit-long or visit-long slopes with a peak-to-peak variation of about 0.05%. These are perhaps caused by slow drifts in pointing and focus (telescope "breathing") interacting with sensitivity variations across the detector that are not perfectly corrected by the flat field.

4.4. *Correcting for Systematics*

Fortunately, these systematics are extremely repeatable between orbits within a visit; we harness this fact when correcting for them. We divide the in-transit orbit of any photometric timeseries, either the white light curve or one of the spectroscopically resolved light curves, by a systematics correction template constructed from the two good out-of-transit orbits. This template is simply the weighted average of the fluxes in the out-of-transit orbits, evaluated at each exposure within an orbit. It encodes both variations in the effective sensitivity of the detector within an orbit and the mean out-of-transit flux level.

When performing the division, we propagate the template uncertainty into the photometric uncertainty for each exposure, which typically increases it by a factor of $\sqrt{1+1/2} = 1.22$. This factor, although it may seem like an undesired degradation of the photometric precision, would inevitably propagate into measurements of the transit depth whether we performed this correction or not, since $R_p/R_\star$ is always measured relative to the out-of-transit flux, which must at some point be inferred from the data.

Throughout this work, we refer to this process of dividing by the out-of-transit orbits as the `divide-oot` method. Because each point in the single in-transit orbit is equally spaced in time between the two out-of-transit exposures being used to correct it, the `divide-oot` method also naturally removes the 0.05% visit-long slope seen in the raw photometry. As we show in §4.5, when applied to the white light curves, the `divide-oot` treatment produces uncorrelated Gaussian residuals that have a scatter consistent with the predicted photon uncertainties.

Unlike decorrelation techniques that have often been used to correct systematics in HST light curves, the `divide-oot` method does not require knowing the relationship between measured photometry and the physical state of the camera. It does, however, strictly require the systematics to repeat over multiple orbits. The `divide-oot` method would not work if the changes in the position, shape, and rotational angle of the 1$^{\rm st}$-order spectrum were not repeated in the other orbits in a visit or if the cadence of the illumination were not nearly identical across orbits. In such cases, the Gaussian process method proposed by Gibson et al. (2011b) may be a useful alternative, and one that would appropriately account for the uncertainty involved in the systematics correction.

4.5. *White Light Curve Fits*

Although the main scientific result of this paper is derived from the spectroscopic light curves presented in §4.6, we also analyze the light curve summed over all wavelengths between 1.1 and 1.7 μm. We use these white light curves to confirm the general system properties found in previous studies and quantitatively investigate the instrumental systematics.

We fit an analytic, limb-darkened transit light curve model (Mandel & Agol 2002) to the `divide-oot`-corrected white light curves. Only the in-transit orbits were fit; after the `divide-oot` correction, the two out-of-transit orbits contain no further information. Also, because the in-transit orbit's flux has already been normalized, we fix the out-of-transit flux level to unity in all the fits. Throughout, we fix the planet's period to $P = 1.58040481$ days and mid-transit time to $T_c = 2454966.525123$ $\mathrm{BJD_{TDB}}$ (Bean et al. 2011), the orbital eccentricity to $e = 0$, and the stellar mass to $0.157 M_\odot$ (Charbonneau et al. 2009).

4.5.1. *Combined White Light Curve*

First, we combine the three visits into a single light curve, as shown in Fig. 5, and fit for the following parameters: the planet-to-star radius ratio ($R_p/R_\star$), the total transit duration between first and fourth contact (t_{14}), the stellar radius ($R_\star$), and the two coefficients c and d of the square-root limb-darkening law[12]. Previous studies have found no significant transit timing variations for the GJ1214b system (Charbonneau et al. 2009; Sada et al. 2010; Bean et al. 2010; Carter et al. 2011; Désert et al. 2011a; Kundurthy et al. 2011; Berta et al. 2011; Croll et al. 2011), so we fix the time of mid-transit for each visit to be that predicted by the linear ephemeris.

As in Burke et al. (2007), we use the parameters t_{14} and $R_\star$ to ensure quick convergence of the MCMC because correlations among these parameters are more linear than for the commonly fit impact parameter (b) and scaled semi-major axis ($a/R_\star$). Because nonlinear transformations between parameter pairs will deform the hypervolume of parameter space, we include a Jacobian term in the priors in Eq. 1 to ensure uniform priors for the physical parameters R_p, $R_\star$, and i (see Burke et al. 2007; Carter et al. 2008, for detailed discussions). For the combined light curves, the influence of this term is practically negligible, but we include it for completeness. In the MCMC chains described in this section, all parameters have correlation lengths of 6-13 points.

Initially, we perform the fit with limb-darkening coefficients c and d without any priors from the `PHOENIX` atmosphere model, enforcing only that $0 < c + d < 1$, which ensures that the star is brighter at its center ($\mu = 1$) than at its limb ($\mu = 0$). Interestingly, the quantity $(c/3 + d/5)$, which sets the integral of $I(\mu)$ over the stellar surface, defines the line along which c and d are most strongly correlated in the MCMC samples (see also Irwin et al. 2011). For quadratic limb-darkening, the commonly quoted $2u_1 + u_2$ combination (Holman et al. 2006) has the same physical meaning. The integral of $I(\mu)$ can be thought of as the increase in the central transit depth over that for a constant-intensity stellar disk, so it makes sense that it is well-constrained for nearly equatorial transiting systems like GJ1214b. Planets with higher impact parameters do not sample the full range of $0 < \mu < 1$ during transit, leading to correspondingly weaker limb-darkening constraints that can be derived from their light curves (see Knutson et al. 2011). We quote confidence intervals for the linear combination $(c/3 + d/5)$ and one orthogonal to it in Table 2, along with rest of the parameters.

Heartened by finding that when they are allowed to vary freely, our inferred white-light limb-darkening coef-

[12] The square-root law is a special case of the 4-parameter law and straightforward to include in the Mandel & Agol (2002) model.

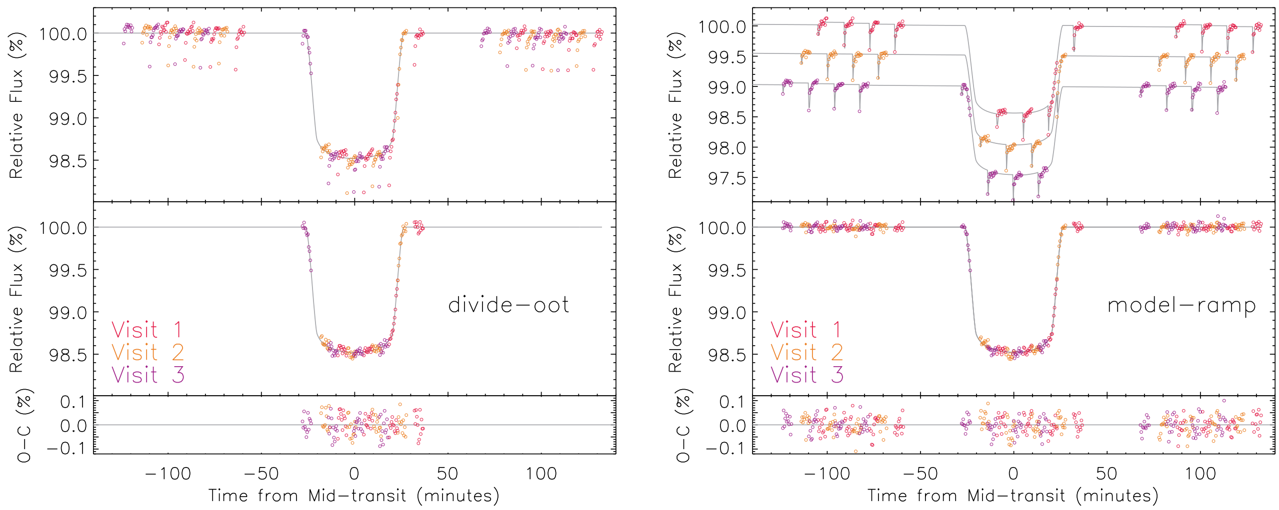

Figure 5. The white light curve of GJ1214b's transits before (*top panels*) and after (*middle panels*) removing the instrumental systematics using the `divide-oot` (*left*) and `model-ramp` (*right*, with offsets for clarity) methods described in §4.4 and §4.5.4. A transit model that was fit to the `divide-oot`-corrected light curve, constrained to the values of $a/R_\star$ and b used by Bean et al. (2010), is shown (*gray lines*), along with residuals from this model (*bottom panels*). In the left panels, the out-of-transit orbits are not shown after the correction has been applied, because they contain no further information.

ficients agree to 1σ to those derived using the `PHOENIX` stellar model, we perform a second fit that includes the `PHOENIX` models as informative priors. For this prior, we say $P(\mathbf{M})$ in Eq. 1 is proportional to a Gaussian with $(c/3 + d/5) = 0.0892 \pm 0.018$ and $(c/5 - d/3) = -0.431 \pm 0.032$, which is centered on the `PHOENIX` model. To set the 1σ widths of these priors, we start by varying the effective temperature of the star in the `PHOENIX` model by its 130K uncertainty in either direction, and then double the width of the prior beyond this, to account for potential systematic uncertainties in the atmosphere model. The results from the fit with these LD priors are shown in Table 2.

The photometric noise rescaling parameter s is within 10% of unity, implying that the 376 ppm achieved scatter in the combined white light curve can be quite well-explained from the known sources of uncertainty in the measurements, predominantly photon noise from the star. As shown in Fig. 6, for the `divide-oot`-corrected light curves, the autocorrelation function (ACF) of the residuals shows no evidence for time-correlated noise. Likewise, the scatter in binned `divide-oot` residuals decreases as the square-root of the number of points in a bin, as expected for uncorrelated Gaussian noise. If there are uncorrected systematic effects remaining in the data, they are below the level of the photon noise over the time-scales of interest here.

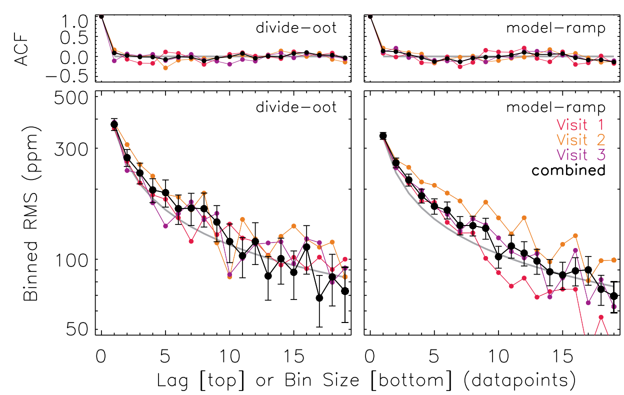

Figure 6. For the transit model in Fig. 5 and both types of systematics treatments, the autocorrelation function of the residuals (ACF; *top*) and the scatter in binned residuals as a function of bin size (*bottom*). The residuals from the combined light curve are shown (*black points*), as well as the individual visits (*colorful points*). The expectations from uncorrelated Gaussian noise (0 in the top, $\propto 1/\sqrt{N}$ in the bottom) are overplotted (*dashed lines*).

4.5.2. *Individual White Light Curves*

To test for possible differences among our WFC3 visits, we fit each of the three `divide-oot`-corrected white light curves individually. In addition to $R_p/R_\star$, t_{14}, and $R_\star$, we also allow ΔT_c (the deviation of each visit's mid-transit time from the linear ephemeris) to vary freely. We allow c and d to vary, but enforce the same `PHOENIX`-derived priors described in §4.5.1.

Table 3 shows the results, which are consistent with each other and with other observations (Charbonneau et al. 2009; Bean et al. 2010; Carter et al. 2011; Kundurthy et al. 2011; Berta et al. 2011). The three measured ΔT_c's show no evidence for transit timing variations. The uncertainties for the parameters t_{14}, $R_\star$, and ΔT_c are noticeably largest in the first visit; this is most likely because the first visit does not directly measure the timing of either 1$^{\rm st}$ or 4$^{\rm th}$ contact, on which these parameters strongly depend. Additionally, whereas the correlation lengths in the MCMC chains for these parameters in the two visits that do measure 1$^{\rm st}$/4$^{\rm th}$ contact and for $R_p/R_\star$, c, and d in all three visits are small (10-30 points), the correlation lengths for t_{14}, $R_\star$, and ΔT_c in the first visit are very large (300-400 points), indicating these weakly constrained parameters are poorly approximated by the MPFIT-derived covariance matrix. On account of the large correlation lengths for these parameters, we ran the MCMC for the first visit with a factor of 10 more points. In each of the three visits, the

Table 2
Transit Parameters Inferred from the Combined White Light Curve

Parameter	No LD Prior	LD Prior[a]
t_{14} (days)	0.03624 ± 0.00013	0.03620 ± 0.00012
$R_\star$ ($R_\odot$)	$0.2014^{+0.0038}_{-0.0025}$	$0.201^{+0.004}_{-0.003}$
$a/R_\star$	$15.30^{+0.19}_{-0.29}$	$15.31^{+0.21}_{-0.29}$
i (°)	89.3 ± 0.4	$89.3^{+0.4}_{-0.3}$
b	$0.18^{+0.09}_{-0.11}$	$0.19^{+0.08}_{-0.11}$
$R_p/R_\star$[b]	$0.1158^{+0.0007}_{-0.0006}$	0.1160 ± 0.0005
$c/3 + d/5$	0.096 ± 0.008	0.095 ± 0.007
$c/5 - d/3$	$-0.52^{+0.22}_{-0.14}$	-0.433 ± 0.032
predicted RMS	337 ppm	337 ppm
achieved RMS	373 ppm	376 ppm
s	1.12 ± 0.07	1.12 ± 0.07

[a] The Gaussian limb-darkening priors of $(c/3 + d/5) = 0.0892 \pm 0.018$ and $c/5 - d/3 = -0.4306 \pm 0.032$ were derived from `PHOENIX` stellar atmospheres, as described in the text.
[b] Confidence intervals on $R_p/R_\star$ do not include the $\Delta R_p/R_\star = 0.0006$ systematic uncertainty due to stellar variability (see text).

Table 3
Transit Parameters Inferred from Individual White Light Curves

Parameter	Visit 1	Visit 2	Visit 3
ΔT_c (days)[a]	$-0.0001^{+0.0012}_{-0.0020}$	0.00028 ± 0.00031	0.0002 ± 0.0004
t_{14} (days)	$0.037^{+0.004}_{-0.003}$	0.0357 ± 0.0007	0.0369 ± 0.0010
$R_\star$ ($R_\odot$)	$0.211^{+0.021}_{-0.014}$	$0.200^{+0.008}_{-0.006}$	$0.214^{+0.015}_{-0.011}$
$a/R_\star$	14^{+1}_{-1}	$15.4^{+0.5}_{-0.6}$	14.4 ± 0.9
i (°)	88.9 ± 0.7	89.2 ± 0.6	$88.5^{+0.8}_{-0.7}$
b	0.27 ± 0.17	0.21 ± 0.15	$0.38^{+0.13}_{-0.20}$
$R_p/R_\star$	$0.1164^{+0.0009}_{-0.0008}$	$0.1159^{+0.0011}_{-0.0009}$	$0.1175^{+0.0011}_{-0.0012}$
predicted RMS	337 ppm	337 ppm	337 ppm
acheived RMS	343 ppm	360 ppm	366 ppm
s	$1.07^{+0.13}_{-0.11}$	$1.12^{+0.13}_{-0.11}$	$1.14^{+0.14}_{-0.11}$

[a] Offset between the observed mid-transit time and that calculated from the linear ephemeris with $P = 1.58040481$ and T_c=2454966.525123 $\mathrm{BJD_{TDB}}$.

uncertainty rescaling parameter s is slightly above but consistent with unity, indicating the photometric scatter is quite well explained by known sources of noise.

4.5.3. *Stellar Variability*

GJ1214 is known to be variable on 50-100 day timescales with an amplitude of 1% in the MEarth bandpass (715-1000 nm; see Charbonneau et al. 2009; Berta et al. 2011). To gauge the impact of stellar variability in the wavelengths studied here, we plot in Fig. 7 the relative out-of-transit flux as measured by our WFC3 data. For each HST visit, we have three independent measurements of this quantity: the F130N narrow-band direct image, the 0^{th}-order spectrum, and the 1^{st}-order spectrum. Consistent variability over these measurements that sample different regions of the detector within each visit would be difficult to reproduce by instrumental effects, such as flat-fielding errors. In Fig. 7, GJ1214 appears brighter in the first visit than in the last two visits, with an overall amplitude of variation of about 1%.

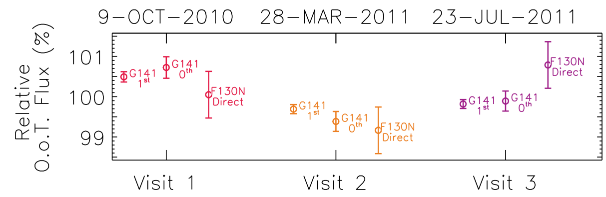

Figure 7. The relative out-of-transit (O.o.T.) flux for each HST visit, measured independently from three different groups of images: the summed 1^{st}-order spectrum, the 0^{th}-order image, and the narrow-band direct image, each normalized to its mean. Error bars denote the standard deviation of the out-of-transit measurements within each visit; they do not include the 0.5% uncertainty in the detector flat-field. The narrow-band measurements sample fewer photons, thus their larger uncertainties.

This 1% variability, if caused by unocculted spots on the stellar surface, should lead to variations in the inferred planet-to-star radius ratio on the order of $\Delta R_p/R_\star = 0.0006$ (Berta et al. 2011). This is larger than the formal error on $R_p/R_\star$ from the combined white light curve (Tab. 2), and must be considered as an important systematic noise floor in the measurement of the absolute, white-light transit depth. We do not detect this variability in the individually measured transit depths (Tab. 3) because it is smaller than the uncertainty on each. Most importantly, while the spot-induced variability influences the absolute depth at each epoch, its effect on the relative transit depth among wavelengths will be much smaller and not substantially bias our transmission spectrum estimate.

4.5.4. *Modeling Instrumental Systematics*

Before calculating GJ1214b's transmission spectrum, we detour slightly to use the white light curves's high photometric precision to investigate the characteristics of WFC3's instrumental systematics. Rather than correcting for the instrumental systematics with the simple non-parametric `divide-oot` method, in this section we describe them with an analytic model whose parameters illuminate the physical processes at play. We refer to this treatment as the `model-ramp` method.

In this model, we treat the systematics as consisting of an exponential ramp, an orbit-long slope, and a visit-long slope. We relate the observed flux (F_{obs}) to the systematics-free flux (F_{cor}) by

$$\frac{F_{\mathrm{obs}}}{F_{\mathrm{cor}}} = (C + V t_{\mathrm{vis}} + B t_{\mathrm{orb}}) \left(1 - R e^{-(t_{\mathrm{bat}} - D_{\mathrm{b}})/\tau}\right) \quad (5)$$

where t_{vis} is time within a visit (= 0 at the middle of each visit), t_{orb} is time within an orbit (= 0 at the middle of each orbit), t_{bat} is time within a batch (= 0 at the start of each batch), τ is a ramp timescale, and the term

$$D_{\mathrm{b}} = \begin{cases} D & \text{for the } 1^{\mathrm{st}} \text{ batch of an orbit} \\ 0 & \text{for the other batches} \end{cases} \quad (6)$$

allows the exponential ramp to be delayed slightly for the first batch of an orbit.

The exponential form arises out of the toy model proposed by Agol et al. (2010), where a certain volume of the detector pixels has the ability to temporarily trap charge carriers and later release them as excess dark current. In quick series of sufficiently strong exposures, the population of charge traps approaches steady state, corresponding to the flattening of the exponential. Judging by the appearance of the ramp in the 2nd-4th batches of each orbit, the release timescale seems to be short enough that the trap population completely resets to the same baseline level after each 6 minute buffer download (during which the detector was being continually flushed each 2.9 seconds). Compared to these batches, the 1st batch of each orbit appears to exhibit a ramp that is either weaker, or as we have parameterized it with the D_b term, delayed. We do not explain this, but we hypothesize that it relates to rapid changes in the physical state of the detector coming out of Earth occultation affecting the pixels' equilibrium charge trap populations.

The visit-long and orbit-long slopes are purely descriptive terms (as in Brown et al. 2001; Carter et al. 2009; Nutzman et al. 2011), but relate to physical processes in the telescope and camera. The orbit-long slope probably arises from the combination of pointing/focus drifts (see Fig. 3) with our imperfect flat-fielding of the detector. This effect of this orbital phase term could be equally well-achieved, for instance, by including a linear function of the 0th order x and y positions (see Burke et al. 2010; Pont et al. 2007; Swain et al. 2008). The visit-long slope is not mirrored in any of the measured geometrical properties of the star on the detector, and is more difficult to associate with a known physical cause.

In order to determine the parameters C, V, B, R, D, and τ, we fit Eq. 6 multiplied by a transit model to the last three orbits of each visit's uncorrected white light curve. The transit parameters are allowed to vary exactly as in §4.5.2, including the use of the informative prior on the limb-darkening coefficients. The white light curves with the best `model-ramp` fit are shown in Fig. 5, and the properties of the residuals from this model are shown in Fig. 6. The transit parameters from this independent systematics correction method are consistent with those in Tab 3. We do not quite achieve the 280 ppm predicted scatter in the `model-ramp` light curves, and the residuals show slight evidence for correlated noise (Fig. 6). More complicated instrumental correction models could almost certainly improve this, but we only present this simple model for heuristic purposes. In all sections except this one, we use the `divide-oot`-corrected data exclusively for drawing scientific conclusions about GJ1214b.

Fig. 8 shows the inferred PDF's of the instrumental systematics parameters for all three HST visits, graphically demonstrating the striking repeatability of the systematics. As expected from the nearly identical cadence of illumination within each of the three visits, the ramp has the same $R = 0.4\%$ amplitude, $\tau = 30$ second timescale, and $D = 20$ second delay time across all observations. The values of τ and D are similar to the time for a single exposure, 25 seconds (including overhead). While the visit-long slope V is of an amplitude (fading by 0.06% over an entire visit) that could conceivably be consistent with stellar variability, the fact that it is identical across all three visits argues strongly in favor of it being an instrumental systematic. B is the only parameter that shows any evidence for variability between orbits; we would expect this to be the case if this term arises out of flat-fielding errors, since the 1st order spectrum falls on different pixels in the three visits.

Figure 8. The *a posteriori* distribution of the instrumental systematics parameters from the analytic model, in each of the three visits. The MCMC results for single parameters (*diagonal;* histograms) and pairs of parameters (*off-diagonal;* contours encompassing 68% and 95% of the distribution) are shown, marginalized over all other parameters (including c and d with priors, t_{14}, and $R_\star$). V is measured in units of relative flux/(3 × 96 minutes), B in relative flux/(96 minutes), R in relative flux, and both τ and D in seconds. All visits are plotted on the same scale; for quantitative comparison, the median values and 1σ uncertainties of each parameter are quoted along the diagonal. The systematics parameters are remarkably repeatable from visit to visit; also, they are largely uncorrelated with $R_p/R_\star$ (*left column*).

4.6. *Spectroscopic Light Curve Fits*

We construct multiwavelength spectroscopic light curves by binning the extracted first order spectra into channels that are 5 pixels ($\Delta\lambda = 23$ nm) wide. We estimate the flux, flux uncertainty, and effective wavelength of each bin from the inverse-variance (estimated from the noise model) weighted average of each quantity over the binned pixels. For each of these binned spectroscopic light curves, we employ the `divide-oot` method to correct for the instrument systematics.

To measure the transmission spectrum of GJ1214b, we fit each of these 24 spectroscopic light curves from each of the three visits with a model in which $R_p/R_\star$, c, d, and s are allowed to vary. We hold the remaining parameters fixed so that $a/R_\star = 14.9749$ and $b = 0.27729$, which are the values used by Bean et al. (2010), Désert et al. (2011a), and Croll et al. (2011). For limb-darkening priors, we use the same sized Gaussians on the same linear combinations of c and d as in §4.5.1, but center them on the `PHOENIX`-determined best values for each spectroscopic bin (see §4.2). The correlation length of all parameters is < 10 in the MCMC chains.

For most spectroscopic bins, the inferred value of s is within 1σ of unity, indicating that the flux residuals show scatter commensurate with that predicted from photon noise (1400 to 1900 ppm across wavelengths). No evidence for correlated noise is seen in any of the bins, as judged by the same criterion as for the white light curves (see Fig. 6).

Fig. 9 shows the transmission spectra inferred from each of the three visits, as well as the `divide-oot`-corrected, spectrophotometric light curves from which they were derived. For the final transmission spectrum (shown as black points in Fig. 9), we combine the three values of $R_p/R_\star$, and $\sigma_{R_p/R_\star}$ in each wavelength bin by averaging them over the visits with a weighting proportional to $1/\sigma^2_{R_p/R_\star}$. Table 4 gives this average transmission spectrum, as well as the central values of the limb-darkening prior used in each bin. The wavelength grids in the three visits are offset slightly (by less than a pixel) from one another; in Table 4 we quote the average wavelength for each bin.

In §4.5.2 we found that GJ1214's 1% variability at WFC3 wavelengths causes $\Delta D = 0.014\%$ or $\Delta R_p/R_\star = 0.0006$ variations in the absolute transit depth. The starspots causing this variability would have a similar effect on measurements of the transmission spectrum, but unless GJ1214's starspot spectrum is maliciously behaved, the offsets should be broad-band and the influence on the wavelength-to-wavelength variations within the WFC3 transmission spectrum should be much smaller. Each visit's transmission spectrum is a differential measurement made with respect to the integrated stellar spectrum at each epoch; by averaging together three estimates to produce our final transmission spectrum, we average over the time-variable influence of the starspots. Importantly, if GJ1214 is host to a large population of starspots that are symmetrically distributed around the star and do not appear contribute to the observed flux variability over the stellar rotation period, their effect on the transmission spectrum will not average out (see Désert et al. 2011b; Carter et al. 2011; Berta et al. 2011).

If we fix the limb-darkening coefficients to the `PHOENIX` values instead of using the prior, the uncertainties on the $R_p/R_\star$ measurements decrease by 20%. If we use only a single pair of LD coefficients (those for the white light curve) instead of those matched to the individual wavelength bins, the transmission spectrum changes by about 1σ on the individual bins, in the direction of showing stronger water features and being less consistent with an achromatic transit depth. These tests confirm that the presence of the broad H_2O feature in the stellar spectrum (see Fig. 2) makes it especially *crucial* that we employ the detailed, multiwavelength LD treatment.

As a test to probe the influence of the `divide-oot` systematics correction, we repeat this section's analysis using the analytic `model-ramp` method to remove the instrumental systematics; every point in the transmission spectrum changes by much less than 1σ. We also experimented with combining the three visits' spectroscopic light curves and fitting for them jointly, instead of averaging together the transmission spectra inferred separately from each visit. We found the results to be practically identical to those quoted here.

Because the transmission spectrum is conditional on the orbital parameters we held fixed ($a/R_\star$, b), we underestimate the uncertainty in the *absolute* values of $R_p/R_\star$; the quoted $\sigma_{R_p/R_\star}$ are intended for *relative* comparisons only. Judging by the $R_p/R_\star$ uncertainty in the unconstrained white light curve fit (Table 2), varying $a/R_\star$ could cause the ensemble of $R_p/R_\star$ measurements in Table 4 to move up or down in tandem with a systematic uncertainty that is comparable to the statistical uncertainty on each. This is in addition to the $\Delta R_p/R_\star = 0.0006$ offsets expected from stellar variability (Berta et al. 2011).

4.7. *Searching for Transiting Moons*

Finally, we search for evidence of transiting satellite compansions to GJ1214b in our summed WFC3 light curves. The light curve morphology of transiting exomoons can be complicated, but they could generally appear in our data as shallow transit-shaped dimmings *or* brightenings offset from the planet's transit light curve (see Kipping 2011, for a detailed discussion). While the presence of a moon could also be detected in temporal variations of the planetary transit duration (Kipping 2009), we only poorly constrain GJ1214b's transit duration in individual visits due to incomplete coverage.

Based on the Hill stability criterion, we would not expect moons to survive farther than 8 planetary radii away from GJ1214b so their transits should not be offset from GJ1214b's by more than 25 minutes, less than the duration of an HST Earth occultation. We search only the data in the in-transit visit, using the `divide-oot` method to correct for the systematics. Owing to the long buffer download gaps in our light curves (see §2), the most likely indication of a transiting moon in the WFC3 light curve would be an offset in flux from one 12-exposure batch to another. Given the 376 ppm per-exposure scatter in the `divide-oot` corrected light curve, we would have expected to be able to identify transits of 0.4 $R_\oplus$ (Ganymede-sized) moons at 3σ confidence. We see no strong evidence for such an offset. Also, we note that starspot occultations could easily mimic the light curve of a transiting exomoon in the time coverage we achieve with WFC3, and such occultations are known to occur in the GJ1214b system (see Berta et al. 2011; Carter et al. 2011; Kundurthy et al. 2011).

Due to the many possible configurations of transiting exomoons and the large gaps in our WFC3 light curve, our non-detection of moons does not by itself place strict limits on the presence of exo-moons around GJ1214b.

5. DISCUSSION

The average transmission spectrum of GJ1214b from our three HST visits is shown in Fig. 9. To the precision afforded by the data, this transmission spectrum is flat; a simple weighted mean of the spectrum is a good fit, with $\chi^2 = 20.4$ for 23 degrees of freedom.

5.1. *Implications for Atmospheric Compositions*

We compare the WFC3 transmission spectrum to a suite of cloud-free theoretical atmosphere models for GJ1214b. The models were calculated in Miller-Ricci & Fortney (2010), and we refer the reader to that paper for their details. To compare them to our transmission spectrum, we bin these high-resolution ($R = 1000$) models to the effective wavelengths of the 5-pixel WFC3 spectroscopic channels ($R = 50-70$) by integrating over each bin. Generally, to account for the possible suppression of transmission spectrum features caused by the overlap of shared planetary and stellar absorption lines, this binning should be weighted by the photons detected from

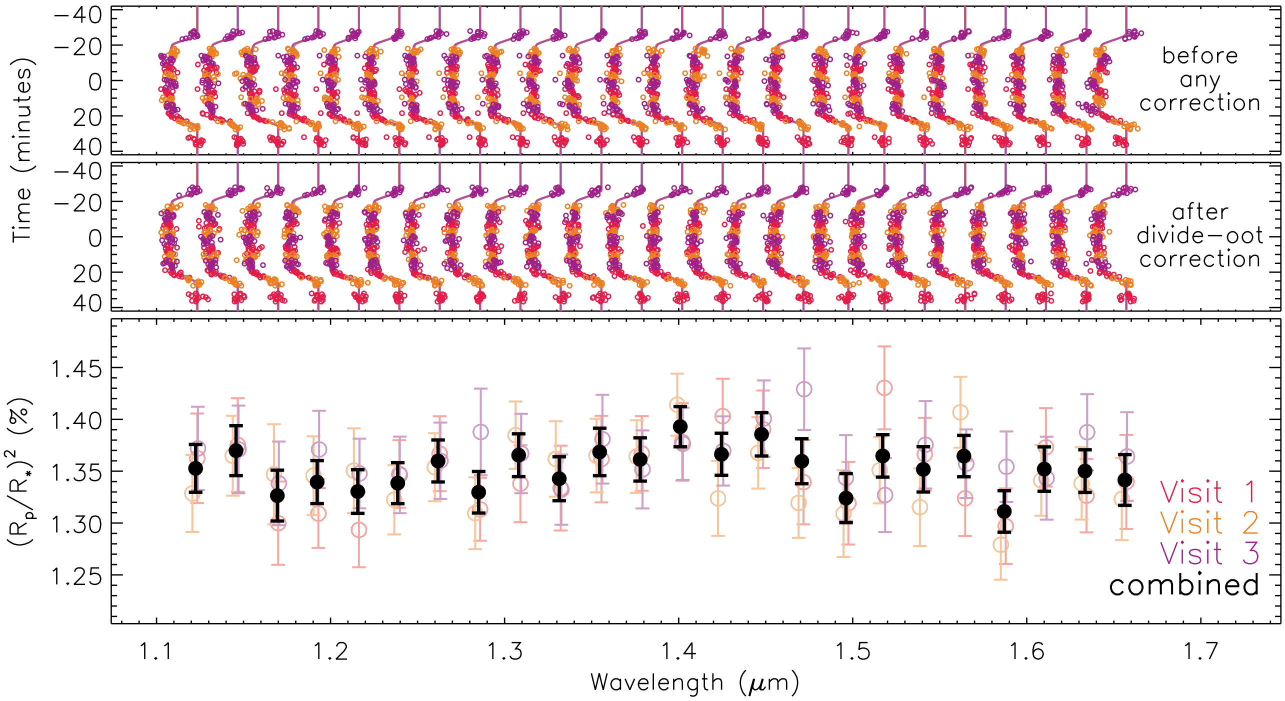

Figure 9. *Top panels:* Spectroscopic transit light curves for GJ1214b, before and after the `divide-oot` correction, rotated and offset for clarity. *Bottom panel:* The combined transmission spectrum of GJ1214b (*black circles with error bars*), along with the spectra measured for each visit (*colorful circles*). Each light curve in the top panel is aligned to its respective wavelength bin in the panel below. Colors denote HST visit throughout.

the system at very high resolution, but this added complexity is not justified for our dataset. The normalization of the model spectra is uncertain (i.e. the planet's true R_p), so we allow a multiplicative factor in $R_p/R_\star$ to be applied to each (giving 24-1=23 degrees of freedom for all models). Varying the bin size between 2 and 50 pixels wide does not significantly change any of the results we quote in this section.

A solar composition atmosphere in thermochemical equilibrium is a terrible fit to the WFC3 spectrum; it has a χ^2=126.2 (see Fig. 9) and is formally ruled out at 8.2σ confidence. Likewise, the same atmosphere but enhanced 50× in elements heavier than helium, a qualitative approximation to the metal enhancement in the Solar System ice giants (enhanced $30-50\times$ in C/H; Gautier et al. 1995; Encrenaz 2005; Guillot & Gautier 2009), is ruled out at 7.5σ ($\chi^2 = 113.2$). Both models assume equilibrium molecular abundances and the absence of high-altitude clouds; if GJ1214b has an H_2-rich atmosphere, at least one of these assumptions would have to be broken.

Suggesting, along these lines, that photochemistry might deplete GJ1214b's atmosphere of methane, Désert et al. (2011a), Croll et al. (2011), and Crossfield et al. (2011) have noted their observations to be consistent with a solar composition model in which CH_4 has been artificially removed. With the WFC3 spectrum alone, we can rule out such an H_2-rich, CH_4-free atmosphere at 6.1σ (Fig. 10). This is consistent with Miller-Ricci Kempton et al. (2011)'s theoretical finding that such thorough methane depletion cannot be achieved through photochemical processes, even when making extreme assumptions for the photoionizing UV flux from the star.

Previous spectroscopic measurements in the red optical (Bean et al. 2010, 2011) could only be reconciled with a H_2-rich atmosphere if such an atmosphere were to host a substantial cloud layer at an altitude above 200 mbar (see Miller-Ricci Kempton et al. 2011). How far the flattening influence of such a cloud layer would extend beyond 1 μm to WFC3 wavelengths would depend on both the concentration and size distribution of the scattering particles. As such, we explore possible cloud scenarios consistent with the WFC3 spectrum in an *ad hoc* fashion, using a solar composition atmosphere and arbitrarily cutting off transmission below various pressures to emulate optically thick cloud decks at different altitudes in the atmosphere. Fig. 11 summarizes the results. A cloud deck at 100 mbar, which would be sufficient to flatten the red optical spectrum, is ruled out at 5.7σ ($\chi^2 = 82.8$). Due to higher opacities between 1.1 and 1.7 μm, WFC3 probes higher altitudes in the atmosphere than the red optical, requiring clouds closer to 10 mbar to match the data ($\chi^2 = 23.4$). Note, with the term "clouds" we refer to all types of particles that cause broad-band extinction, whether they scatter or absorb, and whether they were formed through near-equilibrium condensation (such as Earth's water clouds) or through upper atmosphere photochemistry (such as Titan's haze).

Fortney (2005) and Miller-Ricci Kempton et al. (2011) identified KCl and ZnS as condensates that would be likely to form in GJ1214b's atmosphere, but found they would condense deeper in the atmosphere (200-500 mbar) than required by the WFC3 spectrum and would probably not be optically thick. While winds may be able to

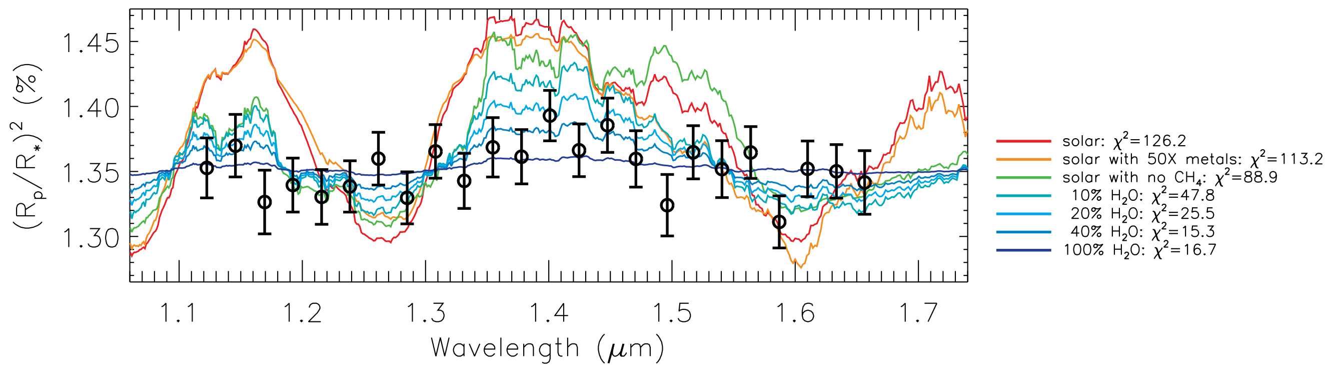

Figure 10. The WFC3 transmission spectrum of GJ1214b (*black circles with error bars*) compared to theoretical models (*colorful lines*) with a variety of compositions. The high resolution models are shown here smoothed for clarity, but were binned over each measured spectroscopic bin for the χ^2 comparisons. The amplitude of features in the model transmission spectra increases as the mean molecular weight decreases between a 100% water atmosphere ($\mu = 18$) and a solar composition atmosphere ($\mu = 2.36$).

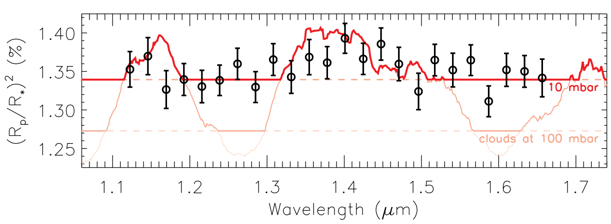

Figure 11. The WFC3 transmission spectrum of GJ1214b (*black circles with error bars*) compared to a model solar composition atmosphere that has thick clouds located at altitudes of 100 mbar (*pink lines*) and 10 mbar *red lines*). We treat the hypothetical clouds in an *ad hoc* fashion, simply cutting off transmission through that atmosphere below the denoted pressures.

loft such clouds to higher altitudes, it is not clear that the abundance of these species alone would be sufficient to blanket the entire limb of the planet with optically thick clouds. The condensation and complicated evolution of clouds has been studied within the context of cool stars and hot Jupiters (e.g Lodders & Fegley 2006; Helling et al. 2008), but further study into the theoretical landscape for equilibrium clouds on planets in GJ1214b's gravity and temperature regime is certainly warranted. The scattering may also be due to a high altitude haze formed as by-products of high-altitude photochemistry; Miller-Ricci Kempton et al. (2011) found the conditions on GJ1214b to allow for the formation of complex hydrocarbon clouds through methane photolysis.

However such clouds might form, they would either need to be optically thick up to a well-defined altitude or consist of a substantial distribution of particles acting in the Mie regime, i.e. with sizes approaching 1 μm. Neither the VLT spectra nor our observations give any definitive indications of the smooth falloff in transit depth toward longer wavelengths that would be expected from Rayleigh scattering by molecules or small particles. This is unlike the case of the hot Jupiter HD189733b, where the uniform decrease in transit depth from 0.3 to 1 μm (Pont et al. 2008; Lecavelier Des Etangs et al. 2008; Sing et al. 2011) and perhaps to as far as 3.6 μm (see Sing et al. 2009; Désert et al. 2009) has been convincingly attributed to a small particle haze.

As an alternative, the transmission spectrum of GJ1214b could be flat simply because the atmosphere has a large mean molecular weight. We test this possibility with H_2 atmospheres that contain increasing fractions of H_2O. This is a toy model, but including molecules other than H_2 or H_2O in the atmosphere would serve principally to increase μ without substantially altering the opacity between 1.1 and 1.7 μm, so the limits we place on μ are robust. We find that an atmosphere with a 10% water by number (50% by mass) is disfavored by the WFC3 spectrum at 3.1σ ($\chi^2 = 47.8$), as shown in Fig. 9. All fractions of water above 20% (70% by mass) are good fits to the data ($\chi^2 < 25.5$). The 10% water atmosphere would have a minimum mean molecular weight of $\mu = 3.6$, which we take as a lower limit on the atmosphere's mean molecular weight.

For the sake of placing the WFC3 transmission spectrum in the context of other observations of GJ1214b, we also display it alongside the published transmission spectra from the VLT (Bean et al. 2010, 2011), CFHT (Croll et al. 2011), Magellan (Bean et al. 2011), and Spitzer (Désert et al. 2011a) in Fig. 12. Stellar variability could cause individual sets of observations to move up and down on this plot by as much as $\Delta D = 0.014\%$ for measurements in the near-IR (Berta et al. 2011); we indicate this range of potential offsets by an arrow at the right of the plot. We display the measurements in Fig. 12 with no relative offsets applied and note that their general agreement is consistent with the predicted small influence of stellar variability. Depending on the temperature contrast of the spots, however, the variability could be larger by a factor of $2-3\times$ in the optical, and we caution the reader to consider this systematic uncertainty when comparing depths between individual studies. For instance, the slight apparent rise in $R_p/R_\star$ toward 0.6 μm, that would potentially be consistent with Rayleigh scattering in a low-μ atmosphere, could also be easily explained through the poorly constrained behavior of the star in the optical. Indeed, Bean et al. (2011) found a significant offset between datasets that overlap in wavelength (near 0.8 μm) but were taken in different years, suggesting variability plays a non-negligible role at these wavelengths.

Finally, we note that any model with $\mu > 4$, such as one with a $> 50\%$ mass fraction of water, would be consistent with the measurements from Bean et al. (2010), Désert et al. (2011a), Crossfield et al. (2011), Bean et al.

Table 4
GJ1214b's Transmission Spectrum from WFC3/G141

Wavelength (μm)	$R_p/R_\star$	c	d
1.123	0.11641 ± 0.00102	−0.372	1.068
1.146	0.11707 ± 0.00099	−0.397	1.088
1.169	0.11526 ± 0.00098	−0.404	1.089
1.192	0.11589 ± 0.00093	−0.410	1.090
1.215	0.11537 ± 0.00091	−0.406	1.075
1.239	0.11574 ± 0.00090	−0.403	1.063
1.262	0.11662 ± 0.00088	−0.407	1.064
1.285	0.11565 ± 0.00088	−0.403	1.045
1.308	0.11674 ± 0.00085	−0.411	1.042
1.331	0.11595 ± 0.00087	−0.390	1.068
1.355	0.11705 ± 0.00089	−0.368	1.101
1.378	0.11664 ± 0.00088	−0.371	1.110
1.401	0.11778 ± 0.00088	−0.338	1.075
1.425	0.11693 ± 0.00091	−0.295	1.029
1.448	0.11772 ± 0.00090	−0.319	1.056
1.471	0.11663 ± 0.00092	−0.322	1.061
1.496	0.11509 ± 0.00100	−0.345	1.084
1.517	0.11635 ± 0.00104	−0.305	1.013
1.541	0.11626 ± 0.00091	−0.330	1.042
1.564	0.11681 ± 0.00091	−0.351	1.059
1.587	0.11443 ± 0.00091	−0.372	1.072
1.610	0.11631 ± 0.00091	−0.399	1.082
1.633	0.11620 ± 0.00092	−0.399	1.075
1.656	0.11581 ± 0.00096	−0.415	1.081

(2011), and WFC3. The only observation it could not explain would be the deep K_s-measurement from Croll et al. (2011). Of the theoretical models we tested, we could find none that matched all the available measurements. We are uncertain of how to interpret this apparent incompatibility but hopeful that future observational and theoretical studies of the GJ1214b system may clarify the issue. In the meantime, we adopt an atmosphere with at least 50% water by mass as the most plausible model to explain the WFC3 observations.

5.2. *Implications for GJ1214b's Internal Structure*

If GJ1214b is not shrouded in achromatically optically thick high-altitude clouds, the WFC3 transmission spectrum disfavors any proposed bulk composition for the planet that relies on a substantial, unenriched, hydrogen envelope to explain the planet's large radius. Both the ice-rock core with nebular H/He envelope and pure rock core with outgassed H_2 envelope scenarios explored by Rogers & Seager (2010) would fall into this category, requiring additional ingredients to match the observations. In contrast, their model that achieves GJ1214b's large radius mostly from a large water-rich core, would agree with our observations.

Perhaps most compellingly, a high μ scenario would be consistent with composition proposed by Nettelmann et al. (2011), who found that GJ1214b's radius could be explained by a bulk composition consisting of an ice-rock core surrounded by a H/He/H_2O envelope that has a water mass fraction of 50-85%. Such a composition would be intermediate between the H/He- and H_2O-envelope limiting cases proposed by Rogers & Seager (2010). The H/He/H_2O envelope might arise if GJ1214b had originally accreted a substantial mass of hydrogen and helium from the primordial nebula but then was depleted of its lightest molecules through atmospheric escape.

5.3. *Prospects for GJ1214b*

Future observations with the James Webb Space Telescope (Deming et al. 2009; Kaltenegger & Traub 2009), one of the immense next generation ground-based telescopes (GMT, TMT, ELT; see Ehrenreich et al. 2006), or possibly even a dedicated campaign with current facilities, could detect the 0.01% transmission spectrum features of a 100% water atmosphere on GJ1214b, and potentially distinguish between clear H_2-poor and cloudy H_2-rich atmospheres. Along another front, simulations by Menou (2011) show that observations of GJ1214b's thermal phase curve, such as those for HD189733b by Knutson et al. (2007a), would probe the ratio of radiative to advective timescales in GJ1214b's outer envelope and provide an independent constraint on the atmospheric composition. Detecting the thermal emission from this 500K exoplanet is currently very difficult, and will likely have to wait until the launch of JWST.

In the meantime, we advocate further study of the GJ1214 system in general. Confirming and refining the parallax for the system (van Altena et al. 1995) will improve our knowledge of the stellar mass, and in turn, the planet's mass and radius. Likewise, further radial velocity observations will empirically constrain the hypothesis by Carter et al. (2011) that a significantly non-zero orbital eccentricity could be biasing GJ1214b's inferred density.

6. CONCLUSIONS

In this work, we made new measurements of the GJ1214b's transmission spectrum using HST/WFC3. Reaching a precision of $\sigma_{R_p/R_\star} = 0.0009$ in 24 simultaneously measured wavelength bins, we found the transmission spectrum to be completely flat between 1.1 and 1.7 μm. We saw no evidence for the strong H_2O absorption features expected from a range of H_2-rich model atmospheres.

Given the lack of a known source for clouds or hazes that could create a truly achromatic transit depth across all wavelengths, we interpret this flat WFC3 transmission spectrum to be best explained by an atmosphere with a high mean molecular weight. Based on our observations, this atmosphere would likely consist of more than 50% water by mass or a mean molecular weight of $\mu > 4$. Such an atmosphere would be consistent with observations of GJ1214b's transmission spectrum by Bean et al. (2010), Désert et al. (2011a), Crossfield et al. (2011), and Bean et al. (2011) although it would be difficult to reconcile with those by Croll et al. (2011).

Such a constraint on GJ1214b's upper atmosphere serves as a boundary condition for models of bulk composition and structure of the rest of the planet. It suggests GJ1214b contains a substantial fraction of water throughout the interior of the planet in order to obviate the need for a completely H/He- or H_2-dominated envelope to explain the planet's large radius. A high bulk volatile content would point to GJ1214b forming beyond the snow line and migrating inward, although any such statements about GJ1214b's past are subject to large uncertainties in the atmospheric mass loss history (see Rogers et al. 2011).

Finally, this paper is the first published study using WFC3 for observing a transiting exoplanet. Aside from several instrumental systematics that were straightforward to correct and did not require a detailed instrumen-

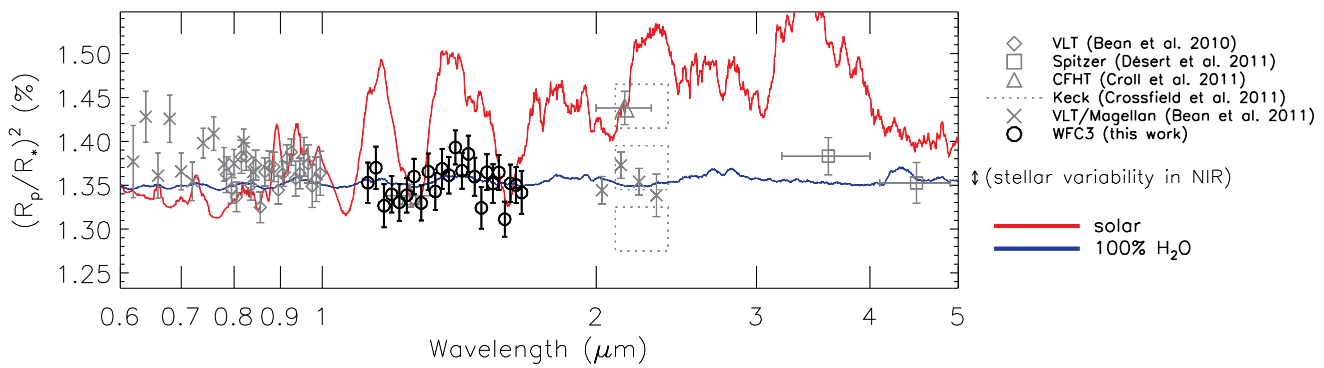

Figure 12. GJ1214b's transmission spectrum from WFC3 in the context of observations from the VLT (0.6-1 μm; Bean et al. 2010, 2011), CFHT (1.25 + 2.15 μm; Croll et al. 2011), Magellan (2.0-2.3 μm Bean et al. 2011), and Spitzer (3.6 + 4.5 μm; Désert et al. 2011a). While they do not measure an absolute transit depth, observations from NIRSPEC on Keck (2.1-2.3 μm Crossfield et al. 2011) disfavor models they tested that had amplitudes larger than 0.05% in their wavelength range; we represent these constraints with the dashed rectangular boxes. Two extremes of the models explored in this paper are shown, normalized to the MEarth-measured transit depth (see Miller-Ricci & Fortney 2010). It is important to note that stellar variability could cause individual data sets to shift up or down on this plot as much as $\Delta D = 0.014\%$ in the near-IR or $2-3\times$ more in the optical, depending on the stellar spot spectrum.

tal model, the camera delivered nearly photon-limited performance both in individual spectrophotometric light curves and in summed white light curves. We are confident that WFC3 will serve as a valuable tool for exoplanet atmospheric characterization in the years to come.

We thank Jacob Bean, Bryce Croll, Ian Crossfield, Ernst de Mooij, Leslie Rogers, Drake Deming, Avi Mandell, Martin Kümmel, Ron Gilliand, Tom Wheeler, and the CfA Summer Statistics Club for valuable discussions regarding this work. We are extremely grateful to our Program Coordinator Patricia Royle, Contact Scientist Howard Bushhouse, the WFC3 instrument team, the entire staff at STScI and NASA, and the crew of STS-125 for their crucial roles in enabling these observations. We thank the anonymous referee for a thorough reading and thoughtful comments that improved the manuscript. E.K. acknowledges funding from NASA through the Sagan Fellowship Program. P.R.M. thanks CfA director Charles Alcock for enabling a sabbatical at SAO. We gratefully acknowledge funding from the David and Lucile Packard Fellowship for Science and Engineering (awarded to D.C.), the National Science Foundation (grant AST-0807690, awarded to D.C.), and NASA (grant HST-GO-12251). This work is based on observations made with the NASA/ESA Hubble Space Telescope, obtained at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555. These observations are associated with program #GO-12251. This research has made use of NASA's Astrophysics Data System.

Facilities: HST (WFC3)

REFERENCES

Agol, E., Cowan, N. B., Knutson, H. A., et al. 2010, ApJ, 721, 1861
Batalha, N. M., Borucki, W. J., Bryson, S. T., et al. 2011, ApJ, 729, 27
Bean, J. L., Désert, J.-M., Kabath, P., et al. 2011, accepted to ApJ, (arXiv:1109.0582)
Bean, J. L., Miller-Ricci Kempton, E., & Homeier, D. 2010, Nature, 468, 669
Berta, Z. K., Charbonneau, D., Bean, J., et al. 2011, ApJ, 736, 12
Brown, T. M. 2001, ApJ, 553, 1006
Brown, T. M., Charbonneau, D., Gilliland, R. L., Noyes, R. W., & Burrows, A. 2001, ApJ, 552, 699
Burke, C. J., McCullough, P. R., Bergeron, L. E., et al. 2010, ApJ, 719, 1796
Burke, C. J., McCullough, P. R., Valenti, J. A., et al. 2007, ApJ, 671, 2115
Carter, J. A., Winn, J. N., Gilliland, R., & Holman, M. J. 2009, ApJ, 696, 241
Carter, J. A., Winn, J. N., Holman, M. J., et al. 2011, ApJ, 730, 82
Carter, J. A., Yee, J. C., Eastman, J., Gaudi, B. S., & Winn, J. N. 2008, ApJ, 689, 499
Charbonneau, D., Berta, Z. K., Irwin, J., et al. 2009, Nature, 462, 891
Charbonneau, D., Brown, T. M., Noyes, R. W., & Gilliland, R. L. 2002, ApJ, 568, 377
Charbonneau, D., Knutson, H. A., Barman, T., et al. 2008, ApJ, 686, 1341
Claret, A. 2000, A&A, 363, 1081
Claret, A., & Hauschildt, P. H. 2003, A&A, 412, 241
Croll, B., Albert, L., Jayawardhana, R., et al. 2011, ApJ, 736, 78
Crossfield, I. J. M., Barman, T., & Hansen, B. M. S. 2011, ApJ, 736, 132
Deming, D., Harrington, J., Seager, S., & Richardson, L. J. 2006, ApJ, 644, 560
Deming, D., Seager, S., Winn, J., et al. 2009, PASP, 121, 952
Désert, J.-M., Bean, J., Miller-Ricci Kempton, E., et al. 2011a, ApJ, 731, L40+
Désert, J.-M., Lecavelier des Etangs, A., Hébrard, G., et al. 2009, ApJ, 699, 478
Désert, J.-M., Sing, D., Vidal-Madjar, A., et al. 2011b, A&A, 526, A12+
Dressel, L., Wong, M., Pavlovsky, C., Long, K., et al. 2010, "Wide Field Camera 3 Instrument Handbook", Version 2.1, (Baltimore: STScI)
Dunkley, J., Bucher, M., Ferreira, P. G., Moodley, K., & Skordis, C. 2005, MNRAS, 356, 925
Eastman, J., Siverd, R., & Gaudi, B. S. 2010, PASP, 122, 935
Ehrenreich, D., Tinetti, G., Lecavelier Des Etangs, A., Vidal-Madjar, A., & Selsis, F. 2006, A&A, 448, 379
Encrenaz, T. 2005, Space Sci. Rev., 116, 99
Ford, E. B. 2005, AJ, 129, 1706
Fortney, J. J. 2005, MNRAS, 364, 649
Gautier, D., Conrath, B. J., Owen, T., de Pater, I., & Atreya, S. K. 1995, in Neptune and Triton, ed. D. P. Cruikshank, M. S. Matthews, & A. M. Schumann, 547–611
Gibson, N. P., Pont, F., & Aigrain, S. 2011b, MNRAS, 411, 2199
Gibson, N. P., Aigrain, S., Roberts, S., et al. 2011a, accepted to MNRAS, (arXiv:1109.3251)

Gregory, P. C. 2005, Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with 'Mathematica' Support, ed. Gregory, P. C. (Cambridge University Press)
Guillot, T., & Gautier, D. 2009, arXiv:0912.2019
Hauschildt, P. H., Allard, F., & Baron, E. 1999, ApJ, 512, 377
Helling, C., Woitke, P., & Thi, W.-F. 2008, A&A, 485, 547
Hilbert, B., & McCullough, P. 2011, WFC3 Instrument Science Report 2011-10 (Baltimore, MD: STScI)
Hogg, D. W., Bovy, J., & Lang, D. 2010, arXiv:1008.4686
Holman, M. J., Winn, J. N., Latham, D. W., et al. 2006, ApJ, 652, 1715
Hubbard, W. B., Fortney, J. J., Lunine, J. I., et al. 2001, ApJ, 560, 413
Irwin, J. M., Quinn, S. N., Berta, Z. K., et al. 2011, accepted to ApJ, (arXiv:1109.2055)
Kaltenegger, L., & Traub, W. A. 2009, ApJ, 698, 519
Kim Quijano, J., et al. 2009, "WFC3 Mini-Data Handbook", Version 1.0, (Baltimore: STScI)
Kipping, D. M. 2009, MNRAS, 392, 181
Kipping, D. M. 2011, MNRAS, 416, 689
Knutson, H. A., Charbonneau, D., Allen, L. E., et al. 2007a, Nature, 447, 183
Knutson, H. A., Charbonneau, D., Noyes, R. W., et al. 2007b, ApJ, 655, 564
Knutson, H. A., Madhusudhan, N., Cowan, N. B., et al. 2011, accepted to ApJ, (arXiv:1104.2901)
Kümmel, M., Kuntschner, H., Walsh, J. R., & Bushouse, H. 2011, Instrument Science Report WFC3-2011-01 (Baltimore, MD: STScI)
Kümmel, M., Walsh, J. R., & Kuntschner, H. 2010, "aXe User Manual", Version 2.1, (Baltimore: STScI)
Kümmel, M., Walsh, J. R., Pirzkal, N., Kuntschner, H., & Pasquali, A. 2009, PASP, 121, 59
Kundurthy, P., Agol, E., Becker, A. C., et al. 2011, ApJ, 731, 123
Kuntschner, H., Bushouse, H., Kümmel, M., & Walsh, J. R. 2009, ST-ECF Instrument Science Report WFC3-2009-17 (Baltimore, MD: STScI)
Kuntschner, H., Bushouse, H., Walsh, J. R., & Kümmel, M. 2008, ST-ECF Instrument Science Report WFC3-2008-16 (Baltimore, MD: STScI)
Kuntschner, H., Kümmel, M., Walsh, J. R., & Bushouse, H. 2011, ST-ECF Instrument Science Report WFC3-2011-05 (Baltimore, MD: STScI)
Lecavelier Des Etangs, A., Pont, F., Vidal-Madjar, A., & Sing, D. 2008, A&A, 481, L83
Léger, A., Rouan, D., Schneider, J., et al. 2009, A&A, 506, 287
Lissauer, J. J., Fabrycky, D. C., Ford, E. B., et al. 2011, Nature, 470, 53
Lodders, K., & Fegley, Jr., B. 2006, Chemistry of Low Mass Substellar Objects, ed. Mason, J. W. (Springer Verlag), 1–+
Long, K. S., Baggett, S., Deustua, S., & Riess, A. 2010, WFC3 Instrument Science Report 2010-17 (Baltimore, MD: STScI)
Long, K. S., Wheeler, T., & Bushouse, H. 2011, WFC3 Instrument Science Report 2011-09 (Baltimore, MD: STScI)
Mandel, K., & Agol, E. 2002, ApJ, 580, L171
Markwardt, C. B. 2009, in Astronomical Society of the Pacific Conference Series, Vol. 411, Astronomical Data Analysis Software and Systems XVIII, ed. D. A. Bohlender, D. Durand, & P. Dowler, 251–+
McCullough, P. 2008, Instrument Science Report WFC3 2008-26 (Baltimore, MD: STScI)
McCullough, P., & Deustua, S. 2008, Instrument Science Report WFC3 2008-33 (Baltimore, MD: STScI)
McCullough, P. R., & MacKenty, J. W. 2011, WFC Space Telescope Analysis Newletter 6 (Baltimore, MD: STScI)
Menou, K. 2011, submitted to ApJL, (arXiv:1109.1574)
Miller-Ricci, E., & Fortney, J. J. 2010, ApJ, 716, L74
Miller-Ricci, E., Seager, S., & Sasselov, D. 2009, ApJ, 690, 1056
Miller-Ricci Kempton, E., Zahnle, K., & Fortney, J. J. 2011, submitted to ApJ, (arXiv:1104.5477)
Nettelmann, N., Fortney, J. J., Kramm, U., & Redmer, R. 2011, ApJ, 733, 2
Nutzman, P., Gilliland, R. L., McCullough, P. R., et al. 2011, ApJ, 726, 3
Pallé, E., Zapatero Osorio, M. R., & García Muñoz, A. 2011, ApJ, 728, 19
Pirzkal, N., Mack, J., Dahlen, T., & Sabbi, E. 2011, Instrument Science Report WFC3-2011-11 (Baltimore, MD: STScI)
Pont, F., Gilliland, R. L., Moutou, C., et al. 2007, A&A, 476, 1347
Pont, F., Knutson, H., Gilliland, R. L., Moutou, C., & Charbonneau, D. 2008, MNRAS, 385, 109
Rogers, L. A., Bodenheimer, P., Lissauer, J. J., & Seager, S. 2011, ApJ, 738, 59
Rogers, L. A., & Seager, S. 2010, ApJ, 716, 1208
Sada, P. V., Deming, D., Jackson, B., et al. 2010, ApJ, 720, L215
Seager, S., & Sasselov, D. D. 2000, ApJ, 537, 916
Seager, S., Kuchner, M., Hier-Majumder, C. A., & Militzer, B. 2007, ApJ, 669, 1279
Sing, D. K., Désert, J.-M., Lecavelier Des Etangs, A., et al. 2009, A&A, 505, 891
Sing, D. K., Pont, F., Aigrain, S., et al. 2011, MNRAS, 1159
Sivia, D. S. 1996, Data Analysis: A Bayesian Tutorial (Oxford University Press), xi, 189 p.
Smith, R. M., Zavodny, M., Rahmer, G., & Bonati, M. 2008a, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7021, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
Smith, R. M., Zavodny, M., Rahmer, G., & Bonati, M. 2008b, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7021, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series
Southworth, J. 2008, MNRAS, 386, 1644
Swain, M. R., Vasisht, G., & Tinetti, G. 2008, Nature, 452, 329
van Altena, W. F., Lee, J. T., & Hoffleit, E. D. 1995, The general catalogue of trigonometric [stellar] parallaxes, ed. van Altena, W. F., Lee, J. T., & Hoffleit, E. D.
van Hamme, W. 1993, AJ, 106, 2096
Viana, A. C., & Baggett, S. 2010, WFC3 TV3 Testing: IR Crosstalk, Tech. rep.
Winn, J. N., Matthews, J. M., Dawson, R. I., et al. 2011, ApJ, 737, L18+