



The Hubble PanCET Program: A Metal-rich Atmosphere for the Inflated Hot Jupiter HAT-P-41b

Kyle B. Sheppard^{1,2} , Luis Welbanks³ , Avi M. Mandell² , Nikku Madhusudhan³ , Nikolay Nikolov⁴ , Drake Deming¹ , Gregory W. Henry⁵ , Michael H. Williamson⁵, David K. Sing^{6,7} , Mercedes López-Morales⁸ , Jegug Ih¹ , Jorge Sanz-Forcada⁹ , Panayotis Lavvas¹⁰ , Gilda E. Ballester¹¹, Thomas M. Evans¹² , Antonio García Muñoz¹³ , and Leonardo A. dos Santos¹⁴ 

¹ Department of Astronomy, University of Maryland at College Park, College Park, MD 20742, USA; kshep23@umd.edu

² Solar System Exploration Division, NASA's Goddard Space Flight Center, Greenbelt, MD 20771, USA

³ Institute of Astronomy, University of Cambridge, Madingley Road Cambridge CB3 0HA, UK

⁴ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

⁵ Center of Excellence in Information Systems, Tennessee State University, Nashville, TN 37209, USA

⁶ Department of Physics & Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA

⁷ Department of Earth & Planetary Sciences, Johns Hopkins University, Baltimore, MD 21218, USA

⁸ Center for Astrophysics, Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 01238, USA

⁹ Centro de Astrobiología (CSIC-INTA), Spain

¹⁰ Groupe de Spectrométrie Moléculaire et Atmosphérique, Université de Reims Champagne Ardenne, Reims, France

¹¹ Lunar & Planetary Laboratory, Department of Planetary Sciences, University of Arizona, Tucson, AZ, USA

¹² Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Avenue, 37-241, Cambridge, MA 02139, USA

¹³ Zentrum für Astronomie und Astrophysik, Technische Universität Berlin, Berlin, Germany

¹⁴ Observatoire Astronomique de l'Université de Geneve, 51 chemin des Maillettes, 1290 Versoix, Switzerland

Received 2020 March 17; revised 2020 October 19; accepted 2020 October 30; published 2021 January 6

Abstract

We present a comprehensive analysis of the 0.3–5 μm transit spectrum for the inflated hot Jupiter HAT-P-41b. The planet was observed in transit with Hubble STIS and WFC3 as part of the Hubble Panchromatic Comparative Exoplanet Treasury (PanCET) program, and we combine those data with warm Spitzer transit observations. We extract transit depths from each of the data sets, presenting the STIS transit spectrum (0.29–0.93 μm) for the first time. We retrieve the transit spectrum both with a free-chemistry retrieval suite (AURA) and a complementary chemical equilibrium retrieval suite (PLATON) to constrain the atmospheric properties at the day–night terminator. Both methods provide an excellent fit to the observed spectrum. Both AURA and PLATON retrieve a metal-rich atmosphere for almost all model assumptions (most likely O/H ratio of $\log_{10} Z/Z_{\odot} = 1.46^{+0.53}_{-0.68}$ and $\log_{10} Z/Z_{\odot} = 2.33^{+0.23}_{-0.25}$, respectively); this is driven by a 4.9 σ detection of H₂O as well as evidence of gas absorption in the optical (>2.7 σ detection) due to Na, AlO, and/or VO/TiO, though no individual species is strongly detected. Both retrievals determine the transit spectrum to be consistent with a clear atmosphere, with no evidence of haze or high-altitude clouds. Interior modeling constraints on the maximum atmospheric metallicity ($\log_{10} Z/Z_{\odot} < 1.7$) favor the AURA results. The inferred elemental oxygen abundance suggests that HAT-P-41b has one of the most metal-rich atmospheres of any hot Jupiters known to date. Overall, the inferred high metallicity and high inflation make HAT-P-41b an interesting test case for planet formation theories.

Unified Astronomy Thesaurus concepts: [Exoplanet atmospheric composition \(2021\)](#); [Exoplanet atmospheres \(487\)](#)

Supporting material: data behind figures, machine-readable table

1. Introduction

Transit spectroscopy has been fundamental in understanding the physics and chemistry of hot exoplanet atmospheres. Transit observations with the Hubble Space Telescope (HST) and the Spitzer Space Telescope have been especially fruitful in illuminating the composition and atmospheric structure of close-in planets, starting with the first measurements of sodium absorption (Charbonneau et al. 2002) and the first detection of thermal emission (Deming et al. 2005) for the atmosphere of HD209458b.

The installation of the Wide Field Camera 3 (WFC3) instrument and the refurbishment of the Space Telescope Imaging Spectrograph (STIS) on HST opened up a new era of transit spectroscopy measurements for hot Jupiters. WFC3 has provided the first repeatable and well-validated detections of the presence of water vapor (Deming et al. 2013; Huitson et al. 2013; Mandell et al. 2013; Wakeford et al. 2013), and has

opened the field to population studies looking at H₂O abundance and metallicity as a function of stellar and planetary properties (Sing et al. 2016; Tsiaras et al. 2018; Pinhas et al. 2019; Welbanks et al. 2019). The upgraded STIS instrument has been a key contributor in illuminating the critical role that aerosols play in driving the continuum opacity for transit measurements of hot planets (Pont et al. 2013; Nikolov et al. 2014; Sing et al. 2016; Chachan et al. 2019).

One of the most intriguing topics from these studies is the question of atmospheric metallicity. Studies of individual planets suggested a wide diversity of atmospheric metallicity as a function of planetary mass (e.g., Kreidberg et al. 2014; Madhusudhan et al. 2014b; Wakeford et al. 2017, 2018). However, recent homogeneous statistical analyses of many planets reveal that a paucity of water vapor in hot planet atmospheres is the norm (Barstow et al. 2017; Pinhas et al. 2018; Welbanks et al. 2019). One strategy to investigate the relationships between mass and atmospheric metallicity is to

study the best targets within the Saturn and Jupiter mass range, in order to achieve high signal-to-noise ratio and leverage the expectation of a high primordial gas fraction and large transit signals.

First discovered in 2012 (Hartman et al. 2012), the inflated hot Jupiter HAT-P-41b ($T_{\text{eq}} = 1940$ K, $P = 2.7$ days) is a strong candidate to inform these trends. It is among the most inflated hot Jupiters ($R = 1.69R_{\text{Jup}}$, $M = 0.8M_{\text{Jup}}$), and it orbits a relatively inactive mid-F dwarf ($R = 1.68R_{\odot}$, $T_{\text{eff}} = 6390$ K). HAT-P-41b’s extended atmosphere and its host star’s lack of significant variability make it highly amenable to characterization through transit spectroscopy. Johnson et al. (2017) determined the spin-orbit misalignment of the system to be moderate (-22°), while the host star appears to be part of a multi-stellar system, with a wide-orbit late-type companion discovered at ~ 1000 au (Hartman et al. 2012; Wöllert & Brandner 2015; Evans et al. 2016).

Tsiaras et al. (2018) retrieved the WFC3 G141 grism spectrum (1.1–1.7 μm) with τ -Rex (Waldmann et al. 2015a, 2015b), confirming an 4.2σ atmospheric detection and finding no evidence of contributions from either high-altitude clouds or photochemical hazes (e.g., Zahnle et al. 2009). Tsiaras et al. (2018) also found evidence of and abundance constraints for water vapor ($\log_{10}(X_{\text{H}_2\text{O}}) = -2.77 \pm 1.09$), though due to narrow wavelength coverage, the abundance uncertainties are large. Still, they are able to rule out upper atmospheric water depletion. Fisher & Heng (2018) built upon this result by retrieving from the same data set with a focus on cloud opacity and other near-infrared opacity sources (NH_3 , HCN). They find a water abundance of $-0.9^{+0.28}_{-1.20}$, which they note is generally consistent with that of Tsiaras et al. (2018). However, this reported abundance is for a cloud-free model with only H_2O and NH_3 as opacity sources, and this simplified treatment is not necessarily directly comparable with more comprehensive atmospheric models. They also find weak evidence of NH_3 , though they are unable to favor the NH_3 and H_2O model over a model with gray clouds and H_2O .

Wide spectral baselines provide the potential for a more complete and constrained understanding of atmospheric properties (Benneke & Seager 2012; Griffith 2014; Welbanks & Madhusudhan 2019). For example, Line & Parmentier (2016) demonstrated how individual WFC3 spectra are unable to constrain mean molecular weight, due to a degeneracy with partial clouds. Furthermore, for a fully homogeneous cloud cover, the cloud-top pressure is degenerate with the chemical abundance (e.g., Deming et al. 2013). Welbanks & Madhusudhan (2019) showed that optical data help alleviate such degeneracies and improve the precision with which planetary radius, cloud properties, and molecular/atomic abundances are inferred. As a practical example, Sing et al. (2016) utilized optical-to-infrared spectra to jointly constrain cloud, haze, and chemistry parameters for a sample of ten hot Jupiters. STIS data have been specifically useful in constraining the atmospheric metallicity of giant exoplanets (Wakeford et al. 2017, 2018; Chachan et al. 2019).

Bayesian spectral retrievals are the most reliable way to interpret exoplanet spectra, due to their flexibility in describing diverse exoplanet atmospheres and their ability to evaluate the full posterior distribution of a forward model’s parameters (Madhusudhan 2018). This allows for understanding not only the properties of an exoplanet’s atmosphere, but also the uncertainties on those properties. Consequently, such retrieval

codes are common in atmospheric characterization (e.g., Madhusudhan & Seager 2009, 2010; Benneke & Seager 2012; Lee et al. 2012; Line et al. 2013; Amundsen et al. 2014; Waldmann et al. 2015b; Barstow et al. 2017, and many others). Nested sampling (Skilling 2004) is a particularly powerful Bayesian sampler, as it naturally determines the Bayesian evidence of the fitted model (with the posterior distribution being a byproduct), which is necessary for correct model comparison (e.g., justifying more complicated models, reporting correct detection significances).

Despite their ubiquity, each retrieval code is necessarily unique, given the assumptions and modeling choices that must be made. Though these retrievals generally agree, subtle discrepancies can lead to different conclusions for the same data (Fisher & Heng 2018; Kilpatrick et al. 2018; Barstow et al. 2020). Examples include different chemical parameterizations (i.e., enforcing chemical equilibrium), cloud parameterizations, opacity sources, and prior assumptions. Therefore, it is important to understand the effect of modeling assumptions on the retrieved atmospheric parameters (e.g., Welbanks & Madhusudhan 2019). Testing a suite of models for a given retrieval code—and, even better, different modeling paradigms altogether—more accurately captures the uncertainty in the atmospheric parameters. It is important to be transparent about the assumptions made in a retrieval analysis in order to best contextualize and understand the results.

In this paper, we derive the 0.3–5 μm transit spectrum of HAT-P-41b using transit observations from HST/STIS, HST/WFC3, and Spitzer (Section 2). Section 3 characterizes both the variability (incorporating both X-ray and visible photometric monitoring; Section 3.1) and the parameters (Section 3.2) of the host star. Section 4 describes the data analysis to derive the transit spectrum. Section 5 describes the two different retrieval methods that we use. First, we use a chemical-equilibrium framework (PLANETARY ATMOSPHERIC TRANSMISSION FOR OBSERVER NOOBS; PLATON; Zhang et al. 2019, Section 6), and those results are described in Section 7. We also explore a more flexible free-chemistry retrieval using the AURA framework (Pinhas et al. 2018, Section 8). The two retrieval analyses were independently done by different members of the team, to allow for an unbiased comparison. We conclude that a high, supersolar atmospheric metallicity (on the order of 30–200 \times solar O/H) best describes the observed spectrum, and—though median retrieved values differ—this result is not sensitive to model assumptions. Section 9 discusses the comparison between retrievals (Section 9.1), the comparison to interior modeling constraints (Section 9.2), and the implications for planet formation (Section 9.3). Section 10 provides a summary of our conclusions.

2. Observations

2.1. HST

We observed one transit of HAT-P-41b with the WFC3 instrument on HST and three transits with the STIS instrument as part of the PanCET Program (ID 14767, P.I. Sing). The WFC3 observations were taken on 2016 October 16 using the G141 prism, which covers a wavelength range of approximately 1.1–1.7 μm with a spectral resolving power of $R \sim 150$. The STIS observations were taken with the G430L and G750L grisms, which cover a wavelength range of approximately 0.3–1.0 μm with a spectral resolving power of $R \sim 500$. The

STIS data were acquired on 2017 September 4 (G430L visit 83), 2018 May 7 (G430L, visit 84), and 2018 June 11 (G750L, visit 85). For each visit, the target was observed for 7 hr over five consecutive HST orbits. An HST gyro issue prevented the acquisition of the third orbit for visit 85. A collection of the HST observations is available via MAST: doi:[10.17909/t9-h6m5-h670](https://doi.org/10.17909/t9-h6m5-h670).

For the WFC3 observations, data were taken in spatial scan spectroscopic mode with a forward scanning rate of $0''.065 \text{ s}^{-1}$ along the cross-dispersion axis, resulting in scans across approximately 46 pixel rows. We utilized the 256×256 pixel subarray and the SPARS-10 sampling sequence, with 12 nondestructive reads (NSAMP = 12) resulting in a total integration time of 81 s for each exposure. We obtained a total of 17 exposures in the first HST orbit following acquisition and 19 exposures in each subsequent HST orbit. Typical peak frame counts were $\sim 33,000$ electrons per pixel, which is within the linear regime of the WFC3 detector.

For the STIS observations, each visit consisted of five orbits, with gaps of ~ 45 minutes due to Earth occultations. We utilized the wide $52'' \times 2''$ slit to minimize slit light losses and an integration time of ~ 253 s for each exposure, for a total of 48 spectra for each visit. Data acquisition overheads were minimized by reading out a subarray of the CCD with a size of 1024×128 pixels.

2.2. Spitzer

The Spitzer Infrared Array Camera (IRAC) observations were taken in 2017 January and February as part of Program 13044 (P.I. D. Deming). A single transit of HAT-P-41b was observed in each of the IRAC1 ($3.6 \mu\text{m}$) and IRAC2 ($4.5 \mu\text{m}$) channels. Each transit was preceded by a 30 minute peakup sequence that also mitigates the steepest portion of a temporal ramp due to the detector. The transit was observed over ~ 12 hr, with equivalent in-transit and out-of-transit coverage. Three hundred and thirty-eight exposures were obtained for each transit, and each exposure consisted of 64 subarray frames of 32×32 pixels, using an exposure time of 2.0 s per frame.

2.3. Photometric Monitoring Observations

To better diagnose the likelihood of stellar variability impacting the transit spectrum, we complemented our transit observations with monitoring observations at both visible (AIT) and X-ray wavelengths (XMM-Newton). XMM-Newton observed HAT-P-41 on 2017 April 7 with an overall 17 ks exposure time (Proposal ID 80479, P.I. J. Sanz-Forcada). The target was not detected in any of the EPIC detectors; we discuss the implications of this in Section 3.1.

We obtained nightly ground-based photometric observations of HAT-P-41 during its 2018 and 2019 observing seasons with the Tennessee State University Celestron 14 inch (C14) automated imaging telescope (AIT) located at Fairborn Observatory in the Patagonia Mountains of southern Arizona (see, e.g., Henry 1999; Eaton et al. 2003). The AIT is equipped with an SBIG STL-1001E CCD camera; observations were made through a Cousins R filter. Details of our observing, data reduction, and analysis procedures are described in Sing et al. (2015).

We collected a total of 207 successful nightly observations (excluding a few isolated transit observations) over the two observing seasons. Our observing activities at Fairborn must

Table 1
AIT Photometric Observations of HAT-P-41

Hel. Julian Date (HJD–2,400,000) (1)	Delta R (mag) (2)	Sigma (mag) (3)
58175.0255	–0.56702	0.00122
58176.0214	–0.56986	0.00109
58180.0102	–0.56636	0.00052
58181.0158	–0.56444	0.00293
58182.0022	–0.56479	0.00185
58184.9980	–0.56390	0.00182

Note. Table 1 is presented in its entirety in the electronic edition of *ApJ*. A portion is shown here for guidance regarding its form and content.

(This table is available in its entirety in machine-readable form.)

Table 2

Results of the Analysis of Photometric Monitoring Observations for HAT-P-41

Observing Season (1)	N_{obs} (2)	Date Range (HJD–2,400,000) (3)	Sigma (mag) (4)	Seasonal Mean (mag) (5)
2018 A	110	58175–58295	0.00224	-0.56496 ± 0.00021
2018 B	34	58386–58451	0.00291	-0.56979 ± 0.00051
2019 A	42	58577–58657	0.00264	-0.56966 ± 0.00041
2019 B	21	58756–58802	0.00394	-0.56602 ± 0.00088

come to a halt each year during the southern Arizona rainy season, which typically lasts from approximately July 1 to September 10. Since HAT-P-41 comes to opposition around July 18, each observing season is broken into two intervals, which we designate as intervals A and B. Information for a portion of the AIT observations are shown in Table 1; the full table is available in the electronic edition of *ApJ*.

3. Stellar Properties

3.1. Analysis of Stellar Variability

The results of analysis of the AIT photometric observations (Section 2.3) are given in Table 2. The low numbers of observations in 2018 B, 2019 A, and 2019 B are the result of the unusually cloudy weather at Fairborn for the past two years. This cloudy weather pattern continues to the present. Column 4 of the table gives the standard deviation of the individual observations with respect to their corresponding seasonal mean. The standard deviations range between 0.00224 and 0.00394 mag for the four observing intervals. This is near the limit of our nightly measurement precision with the C14, as determined from the constant comparison stars in the field. Periodogram analyses of the four intervals reveal no significant periodicities. The scatter in the seasonal means given in column 5 is consistent with the expected photometric precision considering the small number of observations in the last three intervals and the marginal photometric conditions prevalent at Fairborn Observatory over the past two years. Therefore, HAT-P-41 appears to be constant on night-to-night and year-to-year timescales to the limit of our precision.

Additionally, HAT-P-41 was not detected in any of the XMM Newton’s EPIC detectors. Given the distance of the object, we can set an upper limit of $L_X = 1 \times 10^{29} \text{ erg s}^{-1}$ on the stellar X-ray luminosity. This implies a value of $\log L_X/L_{\text{bol}} < -5.2$,

indicating that the star has a moderate activity level at most (Wright et al. 2011).

The photometric observations of HAT-P-41 describe a relatively quiet star. Furthermore, the Ca II chromospheric activity index ($S = 0.18$) and the corresponding estimated parameter flux $\log R'_{HK}$ (-5.04) for HAT-P-41's spectral type ($B - V = 0.29$) are not indicative of high activity (Noyes et al. 1984; Hartman et al. 2012) and may indicate instead a basal-level activity (Isaacson & Fischer 2010). Rackham et al. (2019) show that, for all but the most active F-dwarfs, variability does not result in any detectable change to the transit spectrum. Specifically, the impact of potential complications such as false TiO/VO detections, false water detections, and optical offsets are all determined to be less than ~ 10 ppm. Therefore, we conclude that stellar variability is unlikely to contaminate HAT-P-41b's transit spectrum.

3.2. Stellar Parameters

Inferred atmospheric planetary parameters are directly dependent on host star parameters. For our analyses, we incorporate the stellar parameters from TESS Input Catalog—version 8 (TIC-8; Stassun 2019; Stassun et al. 2019). TIC-8 provides reliable stellar parameters for planetary host stars based primarily on Gaia Data Release 2 (GDR2) point sources (Gaia Collaboration et al. 2016, 2018). The algorithm for HAT-P-41's parameters is as follows: distance is first derived from Gaia DR2 parallax, using a correct inference procedure (Bailer-Jones et al. 2018). HAT-P-41's galactic longitude (-10.6) puts it in a region where uncertainty on reddening makes determining effective temperature from Gaia photometry difficult. As a result, a spectroscopically derived effective temperature (from the PASTEL catalog; see Soubiran et al. (2016)) is preferred. The stellar radius and mass are then self-consistently derived from the distance and effective temperature (Andrae et al. 2018). Finally, $\log g_s$ is calculated from the stellar radius and mass.

It is important to recalculate R_p , M_p , and semimajor axis a based on the R_s value from TIC-8, since those are derived in the discovery paper assuming a certain value for R_s . As a simple example, R_p is derived by constraining R_p/R_s in transit and multiplying by R_s . To rederive the planetary parameters, we follow the methodology of Stassun et al. (2017). The resulting values and 1σ ranges are shown in Table 3, along with the values from the discovery paper (Hartman et al. 2012).

We favor the TIC-8 stellar parameters over the discovery paper values (derived using isochrones and high-resolution spectroscopy; see Hartman et al. (2012)) because they are based on more recent and comprehensive data. We emphasize that the two sets of parameters are consistent to better than 1σ , and using the discovery paper values has no impact on the conclusions of this paper.

A planet's composition is directly linked to its host star's composition. Brewer & Fischer (2018a) determined the stellar abundance of 15 different elements for HAT-P-41 as part of the Spectral Properties of Cool Stars (SPOCS) catalog. Table 4 gives the abundances, relative to solar, for the relevant elements. Brewer & Fischer (2018a) find an effective temperature, metallicity, and $\log g_s$ consistent with both TIC-8 and the discovery paper. HAT-P-41 is a metal-enriched star, and notably has an elemental oxygen abundance of $2.3\times$ solar. Carbon is the only depleted element at $\sim 0.8\times$ solar, resulting in a subsolar C/O ratio of 0.19 ($0.36\times$ solar).

Table 3
System Parameters for HAT-P-41

Parameter	TIC-8 ^a	Discovery Paper ^b
R_s [R_{Sun}]	$1.65_{-0.06}^{+0.08}$	$1.683_{-0.036}^{+0.058}$
M_s [M_{Sun}]	$1.32_{-0.16}^{+0.25}$	1.42 ± 0.047
$\log g_s$ [cgs units]	$4.12_{-0.06}^{+0.11}$	4.14 ± 0.02
$T_{s,\text{eff}}$ [K]	6480_{-100}^{+100}	6390 ± 100
R_p [R_{Jup}]	$1.65_{-0.07}^{+0.08}$	$1.685_{-0.051}^{+0.076}$
M_p [M_{Jup}]	$0.76_{-0.12}^{+0.14}$	0.80 ± 0.10
ρ_p [g cm^{-3}]	0.21 ± 0.05	0.20 ± 0.03
$T_{p,\text{eq}}$ [K]	1960_{-35}^{+40}	1940 ± 38
a [au]	$0.0418_{-0.0019}^{+0.0021}$	0.0426 ± 0.0005
Distance [pc]	348 ± 4.5	344_{-8}^{+12}

Notes.

^a Provided by or derived from TESS Input Catalog; see Stassun et al. (2019).

^b Hartman et al. (2012).

Table 4
HAT-P-41 Elemental Abundances

Elemental Ratio	Abundance (log Solar Unit)
[O/H]	0.37 ± 0.04
[C/H]	-0.08 ± 0.03
[Na/H]	0.17 ± 0.01
[Ti/H]	0.22 ± 0.01
[V/H]	0.09 ± 0.03
[Al/H]	0.07 ± 0.03
[M/H]	0.18 ± 0.01
[C/O]	-0.45 ± 0.05

Note. All values are from Brewer & Fischer (2018b).

4. Data Analysis

4.1. STIS

4.1.1. Data Reduction

Our data analysis procedures follow the general methodology detailed in Nikolov et al. (2014, 2015). We commenced analysis from the fit.fits files, which were reduced (bias-, dark-, and flat-field-corrected) using the latest version of the CALSTIS pipeline and the latest calibration frames. We used median combined difference images to identify and correct for cosmic ray events in the images as described by Nikolov et al. (2014). We found that $\sim 4\%$ of the detector pixels were affected by cosmic ray events. We also corrected pixels identified by CALSTIS as bad with the same procedure. Together with the cosmic ray identified pixels, this resulted in a total of $\sim 14\%$ interpolated pixels.

We performed spectral extraction with the IRAF procedure APALL using aperture sizes in the range from 6 to 18 pixels with a step of 0.5. The best aperture for each grating was selected based on the resulting lowest light-curve residual scatter after fitting the white light curves. We found that aperture sizes of 13.5, 13.5, and 10.5 pixels satisfy this criterion for visits 83, 84, and 85, respectively.

We cross-correlated and interpolated all spectra with respect to the first spectrum, to prevent subpixel wavelength shifts in the dispersion direction. The STIS spectra were then used to extract both white light spectrophotometric time series and

custom wavelength bands after summing the appropriate flux from each bandpass.

The raw STIS light curves exhibit instrumental systematics on the spacecraft orbital timescale, which are attributed to thermal contraction/expansion (referred to as the “breathing effect”) as the spacecraft warms up during its orbital day and cools down during orbital night. We take into account the systematics associated with the telescope temperature variations in the transit light-curve fits by fitting a baseline function depending on various parameters.

4.1.2. Light-curve Analysis

White and spectroscopic light curves were created from the time series of each visit by summing the flux of each stellar spectrum along the dispersion axis. We fit each transit light curve using a two-component function that simultaneously models the transit and systematic effects. The transit model was computed using the analytical formulae given in Mandel & Agol (2002), which are parameterized with the mid-transit times (T_{mid}), orbital period (P) and inclination (i), normalized planet semimajor axis (a/R_s), and planet-to-star radius ratio (R_p/R_*).

Stellar limb-darkening was accounted for by adopting the four-parameter nonlinear limb-darkening law with coefficients c_1 , c_2 , c_3 , and c_4 , computed using a three-dimensional stellar atmosphere model grid (Magic et al. 2015). We adopted the closest match to the effective temperature, surface gravity, and metallicity values for HAT-P-41 determined by Hartman et al. (2012).

As in our past STIS studies, we applied orbit-to-orbit flux corrections by fitting for a fourth-order polynomial to the spectrophotometric time series phased on the HST orbital period and a linear time term. We also used a low-order polynomial (up to a third degree with no cross terms) of the spectral displacement in the dispersion and cross dispersion directions. The first exposures of each HST orbit exhibit lower fluxes and have been discarded in the analysis. Similar to our past HST STIS analyses (Nikolov et al. 2014; Sing et al. 2016), we intended to discard the entire first orbit to minimize the space craft thermal breathing trend, but found that, for two of the three HST visits, a few of the exposures taken toward the end of the first orbit can be used in the analysis.

We then generated systematics models that spanned all possible combinations of detrending variables and performed separate fits using each systematics model included in the two-component function. The Akaike information criterion (AIC; Akaike 1974) was calculated for each attempted function and used to marginalize over the entire set of functions following Gibson (2014). Our choice to rely on the AIC instead of the Bayesian information criterion (BIC; Schwarz 1978) was determined by the fact that the BIC is more biased toward simple models than the AIC. The AIC therefore provides a more conservative model for the systematics and typically results in larger or more conservative error estimates, as demonstrated by Gibson (2014). Marginalization over multiple systematics models assumes equal prior weights for each model tested.

For the white light curves, we fixed the orbital period, inclination, and a/R_s to the values reported in Table 5 and fit for the transit mid-time and planet-to-star radius ratio. We find central transit times $T_c[\text{MJD}] = 58000.6958 \pm 0.0029$ (visit 83), $T_c[\text{MJD}] = 58245.85414 \pm 0.00039$ (visit 84), and

Table 5
Transit Parameters for HAT-P-41b

Parameter	Value
R_p/R_s	0.1028 ± 0.0016
a/R_s	$5.44^{+0.09}_{-0.15}$
i [Degrees]	87.7 ± 1.0
T_c [BJD]	$2454983.86167 \pm 0.00107$
$\log g_p$ [cgs units]	2.84 ± 0.06
P [days]	$2.694047 \pm 4 \times 10^{-6}$

Note. All values are from the discovery paper by Hartman et al. (2012).

$T_c[\text{MJD}] = 58280.87484 \pm 0.00036$ (visit 85). We derive the white light transit depths to be 10200 ± 104 ppm and 10320 ± 85 ppm for G430L and G750L, respectively.

For the spectroscopic light curves, a common-mode systematics model was established by simply dividing the white light curve by a transit model (Berta et al. 2012; Deming et al. 2013). We computed the transit model using the orbital period, inclination, and a/R_s from Table 5, and the central times for each orbit from the white light analysis. The common-mode factors from each night were then removed from the corresponding spectroscopic light curves before model fitting.

We then performed fits to the spectroscopic light curves using the same set of systematics models as in the white light curve analysis and marginalized over them as described above. For these fits, R_p/R_s was allowed to vary for each spectroscopic channel, while the central transit time and system parameters were fixed. We assumed the nonlinear limb-darkening law with coefficients fixed to their theoretical values, determined in the same way as for the white light curve. The detrended spectrophotometric light curves are shown in Figures 19–21. The derived STIS transit spectrum is shown along with the entire transit spectrum in Table 6.

4.2. WFC3

4.2.1. Data Reduction

The WFC3 data reduction generally follows the methodology of Sheppard et al. (2017), using programs adapted from IDL for Python. We download the “ima” data files from the HST archive, and remove background contamination following the “difference reads” methods of Deming et al. (2013), which allows us to easily resolve and remove potential contamination from the nearby companion (Evans et al. 2016). We determine a wavelength solution by taking the zero point from the F140W photometric observation and fitting for the wavelength coefficients that allow an out-of-transit spectrum to match the appropriate ATLAS stellar spectrum (Castelli & Kurucz 2003).

We then divide the background-subtracted “ima” science frame by the WFC3 flat-field calibration file, and return the dark-, bias-, and flat-field-corrected flux array, in units of electrons. The uncertainty of the flux at each pixel is taken from the “ima” file’s error frame, which accounts for gain-adjusted Poisson noise, read noise, and noise from dark current subtraction. This is further adjusted via error propagation for the added uncertainty from background removal and flat-field correction. We use the “ima” file’s data quality frame to mask (i.e., give zero weight to) pixels that are flagged as bad in every exposure in the time series. We then correct for cosmic rays

Table 6
Transit Spectrum

Instrument	Wavelength (μm)	Depth (ppm) ^a	Instrument	Wavelength (μm)	Depth (ppm)	
STIS G430L ^b	0.290–0.350	10091 \pm 230	STIS G750L	0.711–0.731	10499 \pm 364	
	0.350–0.370	10006 \pm 249		0.731–0.750	10038 \pm 302	
	0.370–0.387	10397 \pm 198		0.750–0.770	10142 \pm 224	
	0.387–0.404	10273 \pm 163		0.770–0.799	10124 \pm 281	
	0.404–0.415	9999 \pm 137		0.799–0.819	10618 \pm 317	
	0.415–0.426	9980 \pm 185		0.819–0.838	9910 \pm 239	
	0.426–0.437	10324 \pm 145		0.838–0.884	10252 \pm 238	
	0.437–0.443	10428 \pm 287		0.884–0.930	10217 \pm 298	
	0.443–0.448	10245 \pm 168		WFC3 ^c	1.122–1.141	10297 \pm 107
	0.448–0.454	10411 \pm 165			1.141–1.159	10620 \pm 115
	0.454–0.459	10367 \pm 192			1.159–1.178	10347 \pm 130
	0.459–0.465	10227 \pm 145			1.178–1.196	10479 \pm 113
	0.465–0.470	10640 \pm 164			1.196–1.215	10497 \pm 111
	0.470–0.476	10429 \pm 165			1.215–1.233	10644 \pm 111
	0.476–0.481	10453 \pm 144			1.233–1.252	10289 \pm 100
	0.481–0.492	10584 \pm 125			1.252–1.271	10360 \pm 107
	0.492–0.498	10224 \pm 192	1.271–1.289		10488 \pm 122	
	0.498–0.503	10289 \pm 166	1.289–1.308		10405 \pm 97	
	0.503–0.509	10459 \pm 149	1.308–1.326		10396 \pm 94	
	0.509–0.514	10519 \pm 183	1.326–1.345		10581 \pm 92	
	0.514–0.520	10506 \pm 168	1.345–1.364		10684 \pm 105	
	0.520–0.525	10429 \pm 186	1.364–1.382		10622 \pm 113	
	0.525–0.531	10558 \pm 128	1.382–1.401		10477 \pm 112	
	0.531–0.536	10247 \pm 174	1.401–1.419		10689 \pm 98	
	0.536–0.542	10451 \pm 180	1.419–1.438	10535 \pm 108		
	0.542–0.547	10476 \pm 177	1.438–1.456	10626 \pm 113		
	0.547–0.552	10422 \pm 206	1.456–1.475	10564 \pm 121		
	0.552–0.558	10794 \pm 153	1.475–1.494	10686 \pm 113		
	0.558–0.563 ^d	9221 \pm 279	1.494–1.512	10799 \pm 129		
	0.563–0.569	10444 \pm 188	1.512–1.531	10551 \pm 124		
	STIS G750L ^e	0.526–0.555	10356 \pm 220	1.531–1.549	10566 \pm 111	
		0.555–0.575	10497 \pm 278	1.549–1.568	10515 \pm 134	
0.575–0.594		10317 \pm 202	1.568–1.587	10436 \pm 110		
0.594–0.614		10061 \pm 236	1.587–1.605	10492 \pm 145		
0.614–0.633		10386 \pm 142	1.605–1.624	10340 \pm 122		
0.633–0.653		10542 \pm 219	1.624–1.642	10338 \pm 142		
0.653–0.672		10418 \pm 276	1.642–1.661	10331 \pm 138		
0.672–0.692		10182 \pm 225	Spitzer IRAC1	3.2–4.0	10191 \pm 102	
0.692–0.711		10168 \pm 192		Spitzer IRAC2	4.0–5.0	10679 \pm 145

Notes.

^a Transit depth = $R_{\text{planet}}^2/R_{\text{star}}^2$.

^b Typical STIS G430L bin size = 0.0055 μm (median resolution \sim 350).

^c WFC3 bin size = 0.0186 μm (median resolution \sim 75).

^d Outlier bin strongly affected by systematics and ignored in retrieval analyses.

^e Typical STIS G750L bin size = 0.0196 μm (median resolution \sim 130).

using a conservative time series sigma cut of 8σ , while accounting for changes in flux that occur due to uneven scan rates and the transit itself, and set the affected pixels to the median value of that pixel in the time series. The average amount of pixels either impacted by cosmic ray events or flagged as bad pixels is 1.8% of all pixels in an exposure. The reduced exposures are summed over the spatial scan direction to give a 1D spectrum at each observation time.

4.2.2. Light-curve Analysis

We use a marginalization light-curve analysis similar to that of Sheppard et al. (2017), applied to transit curves. This is a Bayesian model averaging method, first described by Gibson (2014) and applied by Wakeford et al. (2016), with further detrending by use of band-integrated (white light) residuals in

spectral light-curve fitting (Mandell et al. 2013; Haynes et al. 2015). We first analyzed the band-integrated light curve to simplify the spectral light-curve analysis, then we fit the spectral light curves to derive the WFC3 transit spectrum.

We model the observed light curve as a *BATMAN*¹⁵ transit model (Kreidberg 2015) combined with an instrumental systematic model. For the transit model, we assume nonlinear limb darkening and derive the coefficients by interpolating the 3D values from Magic et al. (2015) to the central wavelength of WFC3 (1.4 μm). The limb-darkening derivation is consistent with that used in the STIS analysis. We only fit for transit depth and central transit time, since the incomplete coverage of HST makes it difficult to improve constraints on other transit parameters, such as a/R_s or inclination. We fix these values in

¹⁵ <https://github.com/lkreidberg/batman>

the light curve analyses of each instrument (STIS, WFC3, and Spitzer), which ensures consistent orbital parameters are used when analyzing different data sets. The transit and system parameters are shown in Table 5.

The systematic model grid comprises a set of 50 polynomial models, which account for a visit-long linear slope, HST orbital phase-dependent systematics (i.e., “breathing,” ramp), and subpixel wavelength shifts (for full model grid, see Wakeford et al. (2016)). It is computationally difficult to fully sample the parameter space of all 50 models using Markov Chain Monte Carlo (MCMC) samplers, so we instead fit each model using *KMPFIT*¹⁶ (Terlouw & Vogelaar 2015), a Python implementation of the Levenberg–Markwardt least squares minimization algorithm, to more quickly determine parameter values and uncertainties. Wakeford et al. (2016) found that uncertainties derived from these two methods typically agree within 10%. We then weight each model by its Bayesian evidence—approximated from AIC—and marginalize over the model grid (assuming a prior that each model is equally likely) to derive the light-curve parameters and uncertainties while inherently accounting for uncertainty in model choice.

As is common practice, we ignore the systematic-dominated first orbit in the white light analysis; however, the use of common-mode detrending allows us to include that orbit in the spectral light-curve analysis. The raw light curve, the light curve with instrumental systematics removed, and the residuals from the highest-weight systematic model are shown in Figure 1. We derive the white light depth to be 10490 ± 51 ppm.

To derive the transit spectrum, we bin the 1D spectra from each exposure between the steep edges of the grism response curve (1.12–1.66 μm), deriving a flux time series for each spectral bin. We use bins of width 0.0186 μm (4 pixels) to maximize resolution without drowning the signal in noise. We note that the atmospheric retrieval is not sensitive to the choice of bin size. We fit each spectral light curve using a model similar to that used for the band-integrated light curve, to derive the transit depth at each wavelength.

The spectral light-curve analysis mimics white light analysis with a few exceptions: all system parameters except transit depth are fixed to the white light values, limb-darkening coefficients are interpolated to the bin’s central wavelength, and residuals from the band-integrated light curve are incorporated as a scalable parameter. The band-integrated residuals encode possible instrumental systematics that were not captured by the model grid, and so they can be used to further detrend the binned spectral light curves. The shapes of the residuals are assumed to be constant with wavelength, though the amplitude is allowed to vary. This allows for removal of any wavelength-independent red noise from spectral bin curves, at the penalty of slightly increasing the white noise. Note that the band-integrated uncertainty is sufficiently small, relative to spectral light-curve uncertainty, that the added noise has only a minor effect. The wavelength range, transit depth, and depth uncertainty for each WFC3 bin are shown in Table 6.

The shape of our derived spectrum is in excellent agreement with the literature spectrum (Tsiaras et al. 2018), though it is shifted to higher depths by ~ 90 ppm, indicating that we derive a larger white light depth. This difference persists even if we derive the spectrum without using white light residuals. This

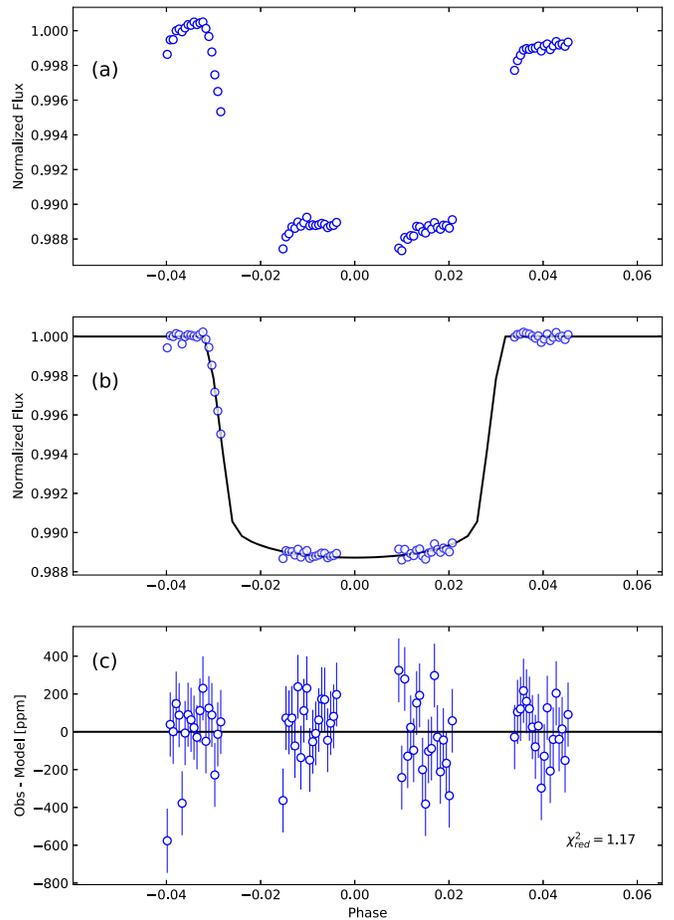


Figure 1. Top panel: pre-processed band-integrated light curve for WFC3 observations. This is the band-integrated flux vs. planet phase derived from the reduced data. The first orbit is excluded, as it is dominated by instrumental systematics. Middle panel: band-integrated light curve divided by the highest-weighted systematic model (i.e., the detrended light curve). Bottom panel: residuals between light-curve data and highest-weighted joint transit and systematic model. Reduced χ^2 for the highest-weighted model is 1.17, which is consistent with the model being a good fit for 67 degrees of freedom.

could plausibly be due to different limb-darkening treatments or different systematic modeling choices. This difference emphasizes the importance of considering offsets between instruments in retrieval analyses (Section 5.1.2). Further, the white light depth is subject to the choice of orbital parameters, which are typically fixed in spectroscopic light-curve fits. We run sensitivity tests to determine that accounting for orbital parameter uncertainties increases the scatter between HST STIS and HST WFC3 by roughly 60 ppm for HAT-P-41b. We capture this scatter by including it in WFC3’s depth uncertainty, increasing it from 50 to 80 ppm.

As a check, we performed retrievals using the published spectrum from Tsiaras et al. (2018) in combination with the derived STIS and Spitzer depths, and found differences well within the 1σ uncertainties for the retrieved parameters. The major results and conclusions in this paper are not sensitive to the WFC3 spectrum choice.

4.2.3. WFC3 Transit Spectrum Verification

Marginalization is only reliable if at least one model is a good representation of the data (Gibson 2014; Wakeford et al. 2016). We therefore checked the goodness-of-fit of the

¹⁶ <https://github.com/kapteyn-astro/kapteyn/blob/master/doc/source/kmpfittutorial.rst>

highest-weighted systematic model for each light curve using both reduced χ^2 and residual normality tests. Further, we explored whether red noise is present in the light-curve residuals, as that can bias inferred depth accuracy and precision (Cubillos et al. 2017).

Though χ^2 cannot prove that a model is correct, it can demonstrate that the fit of a particular model is consistent with that of the “true” model with “true” parameter values (Andrae et al. 2010). Therefore, it is an informative goodness-of-fit diagnostic, and it is particularly useful due to its familiarity and simplicity. The “true” model with “true” parameters will have a reduced χ^2 of one, with uncertainty defined by the χ^2 distribution. For both the band-integrated and spectral light curves (66 and ~ 88 degrees of freedom, respectively), this results in an acceptable reduced χ^2 range of roughly 0.7–1.3.

The band-integrated analysis ($\chi^2_{\nu} = 1.17$) and all but one of the spectral bins (median $\chi^2_{\nu} = 0.9$) fall within this range. The exception is the $1.299 \mu\text{m}$ light curve, for which the highest-weighted model fit has a reduced χ^2 of 0.56. This low value indicates that the uncertainties in this light curve are over-estimated. This is likely due to incorporating white light residuals, which both inflate uncertainties and can potentially interpret random white noise as structure. However, it is not flagged by the normality or correlated noise analyses (described below), and fitting the light curve without incorporating white light residuals finds a consistent depth with a more reasonable $\chi^2_{\nu} = 0.71$. Further, the derived depth and uncertainty exhibit good agreement with the published (Tsiaras et al. 2018) transit depth at this wavelength (accounting for the white light offset). Therefore, we include it in the transit spectrum. For the other 28 spectral bins, the reduced χ^2 values provide no evidence against the validity of the derived transit depths and uncertainties.

A residual normality test checks whether the residuals for a model are Gaussian-distributed in order to determine goodness of fit, as this is expected for the “correct” model. Like reduced χ^2 , a normality test cannot prove that a model is correct; rather, it can only diagnose incorrect models. We use the *scipy* implementation of the common Shapiro–Wilk test for normality (Shapiro & Wilk 1965), and determine for which light curves the highest-evidence model has normality ruled out at the 5% significance level. At a sample size of around 90, this is by no means rigorous, but it is still a useful heuristic for flagging potentially problematic light-curve models.

Normality is rejected at the 5% significance level only for the band-integrated residuals and the $1.243 \mu\text{m}$ spectral bin residuals. In both cases, normality is ruled out due to a single outlier in the time-series. Normality is rejected at 3% significance for the band-integrated residuals, due entirely to the first exposure in the first orbit. When this exposure is ignored, we recover a consistent depth and uncertainty, and the residuals are consistent with normality. We therefore keep this exposure in the analysis.

A possible cause of the spectral bin’s outlier is a minor cosmic ray event or bad pixel that was small enough to both avoid the detection by the sigma cut and not affect the band-integrated curve, but large enough to impact the much smaller bin flux. Removing this spectral bin from the retrieval had no noticeable effect on the results. Further, the derived depth and uncertainty are consistent with literature values (Tsiaras et al. 2018). We therefore decide to include this bin in the retrieval.

Finally, we test for correlated noise in the residuals following the methodology of Cubillos et al. (2017) (also see Pont et al. 2006) and using MC³.¹⁷ Though this method is not necessarily rigorous for HST, due to the incomplete phase coverage and relatively small number of exposures, it is still a practical diagnostic. We find no evidence of correlated noise up to the timescale of half an HST orbit (Figure 2). Beyond this timescale (nine exposures per bin) the standard deviation is based on a light curve with fewer than eight points, and so the relationship is less informative and more subject to small number statistics.

With the caveats noted above, marginalization does an excellent job in fitting the spectral light curves. We emphasize that removing any of the flagged bin spectra has no effect on the retrieval. Together, these tests support the validity of the derived transit depths and uncertainties in the WFC3 bandpass.

4.3. Spitzer

4.3.1. Data Reduction

The Spitzer data consist of cubes of 64 subarray frames in each band, each of size 32×32 pixels. We extracted aperture photometry for each frame, totaling 21,632 frames at both 3.6 and $4.5 \mu\text{m}$. To extract photometry, we used 11 numerical apertures with radii ranging from 1.6 to 3.5 pixels, and we centered those apertures on the position of the host star determined using both a 2D Gaussian fit to the stellar point-spread function and also a center-of-light calculation. Since HAT-P-41 has a companion star $3''.6$ distant (Hartman et al. 2012; Evans et al. 2016), we measured the flux of the companion scattered into each of our numerical apertures, using the method described by Garhart et al. (2020). We adopted the magnitude difference in the Spitzer bands as deduced by Garhart et al. (2020). Accounting for our having used different aperture radii than Garhart et al. (2020), we derive dilution correction factors of 1.0171 and 1.0106 at 3.6 and $4.5 \mu\text{m}$, respectively. The transit depths are then multiplied by those factors in order to correct for the presence of the companion star.

4.3.2. Light-curve Analysis

We fit transit curves to the 22 sets of photometry at each wavelength (eleven apertures, each with two centering methods). Our default fitting procedure fixes the orbital parameters at the values in the discovery paper by Hartman et al. (2012), fitting only for the central time and depth of the transit. The shapes of the Spitzer transits are well-matched when fixing the orbital parameters to those values. However, we also explored including the orbital inclination and a/R_s in the fit (see below), those being the orbital parameters that most strongly affect the shape of the transit. We adopt quadratic limb-darkening coefficients calculated by least squares for the Spitzer bands by Claret et al. (2013), using 2 km s^{-1} microturbulence. We choose the values for $T_{\text{eff}} = 6400 \text{ K}$ and $\log g = 4.0$, without interpolation. We fix those coefficients in the fitting process, and we deem these choices to be appropriate given that the limb darkening is minimal at these infrared wavelengths. Our fits to the transit account for the intrapixel sensitivity variations of the Spitzer photometry using pixel-level decorrelation (PLD; Deming et al. 2015), including

¹⁷ <https://github.com/pcubillos/mc3>

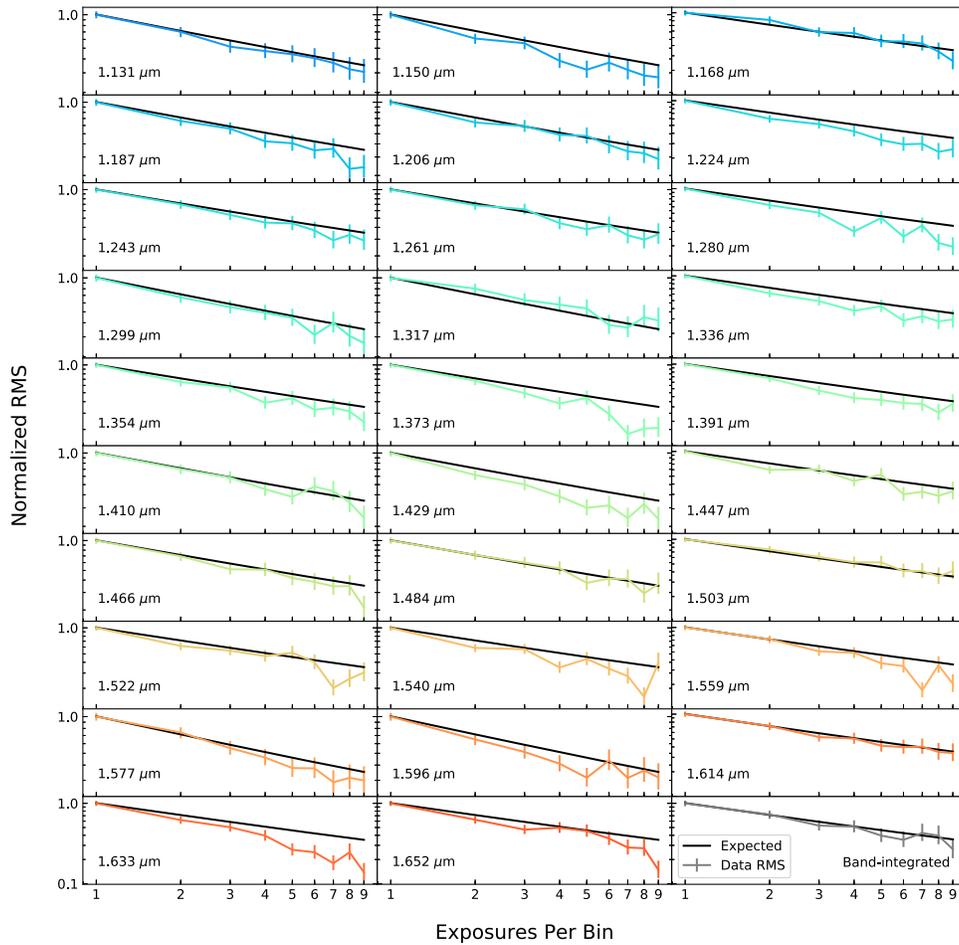


Figure 2. Correlated noise analysis for each of the 29 WFC3 spectral bins and the band-integrated light curve. For each bin’s light curve, we find the rms of the residuals for an increasing number of exposures per bin. Pure white noise would scale with the black line, while correlated noise would increase with binning. Though not exact, given the gaps between WFC3 data, this is a useful heuristic to search for correlated noise. See Cubillos et al. (2017) and Pont et al. (2006) for more details.

a linear baseline (ramp) in time. We use a Bayesian information criterion to decide between a linear versus quadratic ramp. The details of the PLD fit are the same as described for secondary eclipses by Garhart et al. (2020), except that we are fitting transits, so we include limb darkening. Briefly, the fitting code bins the photometry and pixel basis vectors to various degrees, and selects the optimal bin size, aperture radius, and centering method, based on the smallest difference from an ideal Allan deviation relation (Allan 1966). The Allan deviation relation expresses that the standard deviation of the residuals should decrease as the square root of the bin time. Operating on binned data allows the PLD algorithm to concentrate on the longer timescales that characterize the red noise (and also the transit duration), as opposed to the 0.4 s cadence time of the raw photometry.

We determine the errors on our transit depths and times using an MCMC procedure, with a burn-in phase of 10,000 steps, followed by 800,000 steps to explore parameter space. We calculate multiple chains for each transit, and verify convergence using a Gelman–Rubin (GR) statistic (Gelman & Rubin 1992). Our GR values for our PLD fits are very close to unity, being 1.0027 at 3.6 μm and 1.0004 at 4.5 μm , indicating good convergence. Our transit depths and times are listed in Table 7.

Our derived transit times are in excellent agreement with measurements of the same Spitzer transits by Wakeford et al. (2020). Specifically, using our uncertainties, our fitted times

Table 7
Transit Times and Depths for HAT-P-41b in the Spitzer Bands

Wavelength	BJD(TDB)	R_p^2/R_s^2 (ppm)
3.6 μm	2457772.20440 \pm 0.00033	10020 \pm 100
4.5 μm	2457788.36860 \pm 0.00032	10568 \pm 135

Note. These are “as observed” transit depths, not corrected for dilution by the companion star. To correct for dilution, multiply the values of R_p^2/R_s^2 by 1.0171 at 3.6 μm , and by 1.0106 at 4.5 μm . The corrected values are shown with the rest of the transit spectrum in Table 6.

differ by 1.1σ and 0.6σ at 3.6 and 4.5 μm , respectively. Wakeford et al. (2020) use eight HST transits, as well as the discovery results and the Spitzer transits, to derive a new ephemeris. Our fitted times (Table 7) differ from that ephemeris by insignificant amounts (0.2 and 7.8 s).

We explored the effect of uncertainties in the orbital parameters, since those can affect the derived transit depth (Alexoudi et al. 2018). Adopting uniform priors for a/R_s and inclination, we find that they are degenerate when fitting only the Spitzer transits. That degeneracy is illustrated in Figure 3, where it is apparent that inclination and a/R_s can trade off to maintain the observed transit duration and the sharp ingress/egress that characterizes the Spitzer transits. Changing the inclination changes the chord length across the stellar disk, and

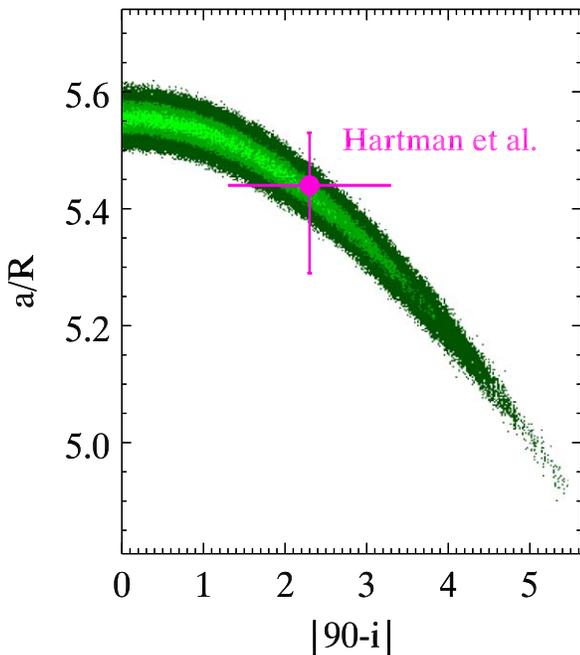


Figure 3. Likelihood distribution from the fit to the $3.6\ \mu\text{m}$ transit of HAT-P-41b, based on an MCMC using uniform priors, and shown vs. a/R_s and the orbital inclination. These two orbital parameters are degenerate when using only a Spitzer transit, and the values derived by Hartman et al. (2012) are indicated by the point with error ranges. The Spitzer transits at both 3.6 and $4.5\ \mu\text{m}$ are fully consistent with a/R_s and i from Hartman et al. (2012).

(when limb-darkening is minimal) that can be compensated by changing a/R_s to maintain the same transit duration. The orbital solution from Hartman et al. (2012) is entirely consistent with our likelihood distribution for those parameters, as shown in Figure 3 for $3.6\ \mu\text{m}$ ($4.5\ \mu\text{m}$ is similar). We therefore freeze the orbital parameters at the Hartman et al. (2012) values when fitting our Spitzer transits.

Figure 4 illustrates the transits at 3.6 and $4.5\ \mu\text{m}$. The residuals from the best-fit model are included in the figure, and the right panel shows the residuals binned over increasing timescales, a so-called Allan deviation relation (Allan 1966). The slopes of those relations are close to the -0.5 value expected for photon noise.

5. Atmospheric Retrieval

There are two common frameworks to retrieve physical parameters from transmission spectra. The first is by assuming chemical equilibrium, where the abundance of a molecule is only dependent on local temperature, local pressure, and global elemental abundances such as O/H and C/H (e.g., Kreidberg et al. 2015). The second is by instead fitting for molecular abundances based on observed spectral features, then determining global elemental abundances from the molecular abundance values (e.g., Madhusudhan 2012). Since carbon and oxygen-based molecules are typically the most spectroscopically active species over the wavelengths covered by HST and Spitzer, the elemental abundances are commonly parameterized by metallicity—defined as the enhancement of metal elements relative to hydrogen compared to solar values (see Section 5.3 for more detail)—and carbon-to-oxygen ratio (C/O). Some retrievals improve flexibility by allowing other elements—such as sodium or vanadium—to also vary from their solar ratios

(Amundsen et al. 2014; Tremblin et al. 2015, 2016; Sing et al. 2016).

The open source code PLATON¹⁸ (Zhang et al. 2019) is able to perform quick retrievals that assume chemical equilibrium, whereas AURA (Pinhas et al. 2018) is able to capture possible disequilibrium chemistry by not assuming chemical equilibrium. We retrieve the atmospheric parameters with both frameworks, in order to see how interpretations compare as well as to explore how sensitive the conclusions are to retrieval assumptions.

5.1. PLATON

PLATON is a fast, open-source retrieval code developed by Zhang et al. (2019). Like many retrieval codes, it comprises a forward model and an algorithm for Bayesian inference. Though there are some differences, it essentially uses the same forward model as Exo-Transmit (Kempton et al. 2017). Here, we summarize the forward model: To calculate a spectrum, it first determines the abundances of 34 potentially relevant chemical species for a given atmospheric metallicity and C/O. These include Na and K as well as molecules CH_4 , CO, CO_2 , HCN, H_2O , MgH, NH_3 , TiO, and VO; for a complete list, see Kempton et al. (2017). The metallicity and C/O provide elemental abundances, which are combined with a temperature–pressure grid as input into GGchem (Woitke et al. 2018) to compute equilibrium molecular and atomic abundances at each pressure layer in the atmosphere, accounting for the effects of condensation on equilibrium abundances. PLATON allows for a gray cloud deck, below which no light can penetrate, and the abundance–temperature–pressure grid facilitates the determination of total opacity at each pressure layer in the atmosphere that lies above this cloud top. PLATON includes the same opacity sources as Exo-Transmit, and accounts for opacity from gas absorption, CIA, and scattering (either parametric Rayleigh scattering or Mie scattering). The forward model converts the opacity–pressure grid to an opacity–height grid using hydrostatic equilibrium with $P_{\text{ref}} = 1$ bar, which is then used as an input to a radiative transfer code in order to determine the uncorrected transit depths. After correcting for possible stellar activity, due to either unocculted spots or faculae, PLATON’s forward model outputs the corrected transit spectrum. The largest source of computational uncertainty is opacity sampling error, which is a source of white noise from using a relatively low resolution ($R = 1000$) that cannot resolve individual lines (Zhang et al. 2019). Accounting for opacity sampling for the transit spectrum of HAT-P-41b typically increases the depth uncertainty by 1.5% (2.5 ppm), which is sufficiently small that it does not affect interpretation. For more details on PLATON, see Zhang et al. (2019). We note that the version of PLATON we describe and use in this paper is Platon 3.1. A newer version, PLATON 5.1, has since been released with additional features as described in Zhang et al. (2020).

Our PLATON analysis does not retrieve individual abundances. Instead, it fits for the isothermal temperature, atmospheric metallicity as a multiple of solar values for atomic species, and C/O ratio; the retrieved equilibrium abundances for atomic and molecular absorbers are a natural consequence of those values. This is in contrast with the AURA analysis (Section 5.2), which retrieves individual molecular and atomic abundances.

¹⁸ <https://github.com/ideasrule/platon>

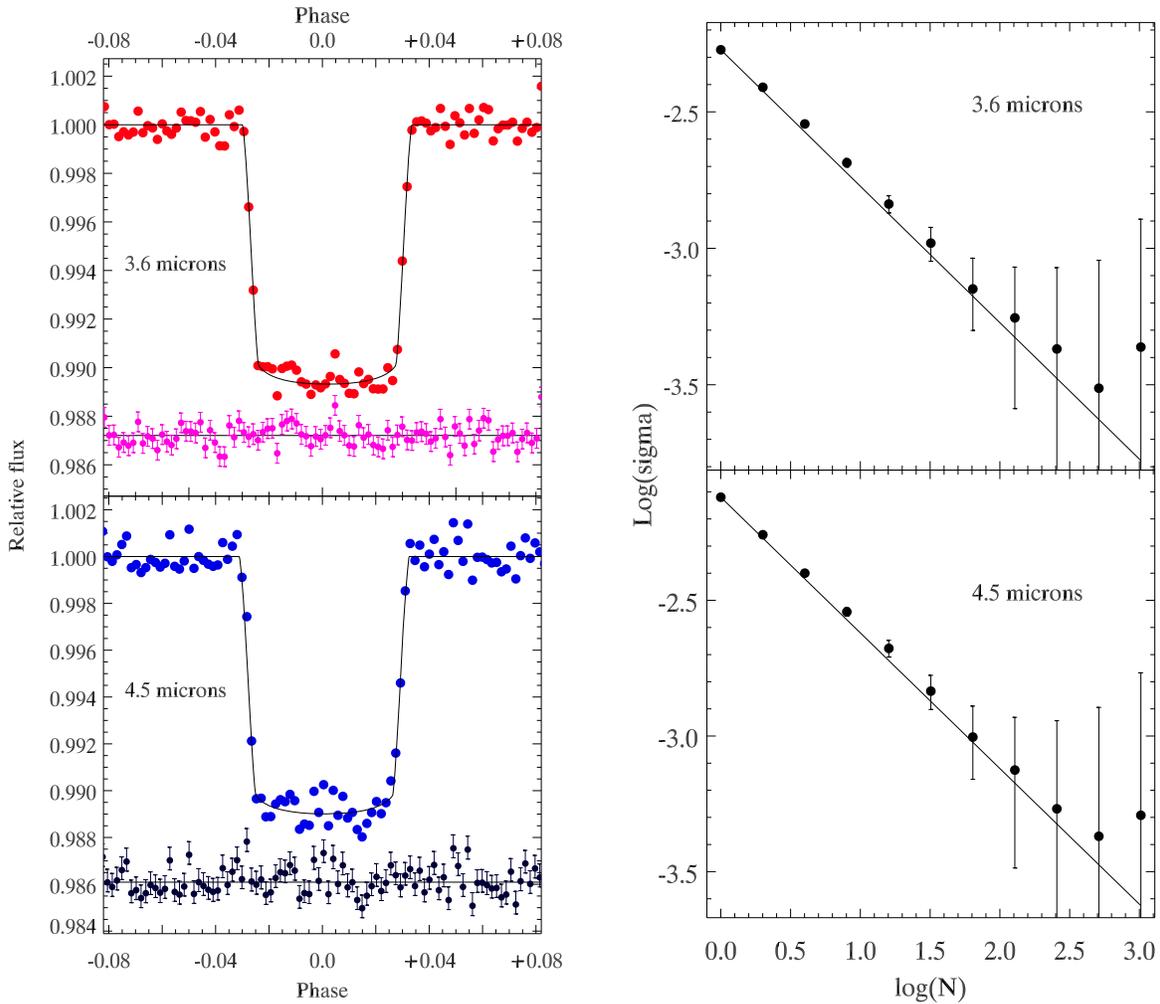


Figure 4. Left: Spitzer transit light curves of HAT-P-41b at 3.6 and 4.5 μm after correction of the intrapixel effects of the detector and temporal ramps. The data are binned to 100 points per transit, for clarity of illustration. Residuals (data minus fitted model) are shown below the transit curves, and have error bars added. Right: Allan deviation relations for the binned residuals, i.e., standard deviation of the residuals when the original data are binned over an increasing number of points, N . Solid lines project the single-point ($N = 1$) scatter to larger bin sizes with a slope of -0.5 , as expected for photon noise.

The algorithm PLATON uses for Bayesian inference is nested sampling (Skilling 2004). Specifically, PLATON uses multimodal nested sampling from the Python implementation *nestle*.¹⁹ Like MCMC samplers, nested sampling efficiently samples posterior distributions with dimensionalities typical of atmospheric retrievals ($n = 5\text{--}20$), and so it is effective at atmospheric parameter estimation. Unlike MCMC routines, it automatically calculates the Bayesian evidence for a model, which is necessary for model comparison. The Bayesian evidence intrinsically accounts for overfitting by punishing too much model structure and thus determining whether extra parameters are warranted. We use this to justify excluding parameters that add structure but do not significantly improve the fit. Nested sampling also has a well-defined stopping criteria, so there is no need to check for convergence. For an excellent write-up on this algorithm, especially about using it in practice, see the documentation of *Dynesty*²⁰ (Higson et al. 2019).

In addition to the standard set included in PLATON’s forward model, we added three new fittable parameters: a partial cloud

parameterization and three instrumental transit depth bias parameters (henceforth referred to as instrumental offsets).

5.1.1. Partial Clouds

The partial cloud parameter is motivated by work by Line & Parmentier (2016) and MacDonald & Madhusudhan (2017), which showed that if the gray cloud deck were inhomogeneous, then the spectrum we observe (D) would be a weighted average of the clear atmosphere transit spectrum (D_{clear}) and the cloudy atmosphere spectrum (D_{cloudy}), with weights given by the cloud fraction (f_c). We implement this as $D = f_c * D_{\text{cloudy}} + (1 - f_c) * D_{\text{clear}}$. Since high-altitude gray clouds are seen in spectra as flat lines, averaging a spectrum with features with this line will shrink the features and can mimic the effect of an atmosphere with high mean-molecular mass and small scale height. Thus, including this parameter allows us to account for this possible degeneracy and prevents us from overconfidently claiming a high-metallicity atmosphere.

5.1.2. Instrumental Offsets

The instrumental offsets are nuisance parameters that can capture the extent to which transit depths from STIS G430L, STIS G750L, or WFC3 are biased relative to the depths from

¹⁹ <https://github.com/kbarbary/nestle>

²⁰ <https://dynesty.readthedocs.io/en/latest/>

the other instruments. This is motivated by the use of common-mode corrections in the light-curve analysis of each instrument, which can potentially introduce a uniform bias for the depth at each spectral bin for that instrument. Offsets are also able to account for inter-instrumental transit depth scatter introduced by uncertainty in orbital parameters a/R_s and inclination (Section 4.2.2).

We explore three offset scenarios. The first is physically motivated. In this scenario, we use Gaussian priors with sigmas determined by the uncertainties in the band-integrated transit depths to try to reflect the correlated uncertainty that exists between spectral bins for each instrument, whether due to common mode corrections or orbital parameter uncertainties. This essentially propagates the white light depth uncertainty into the retrieval. Derived in Section 4, these uncertainties are 105 ppm, 85 ppm, and 80 ppm for STIS G430L, STIS G750L, and WFC3, respectively. The second scenario investigates the potential impact of unknown sources of bias by setting a large, uniform offset prior for each instrument’s offset. The third scenario extends this by setting two large, uniform priors: a WFC3 offset and a single offset for both STIS instruments. The third scenario allows the absolute depths at STIS to vary while preserving the optical spectral shape.

We caution that offsets—especially penalty-free, uniform prior offsets—can cloak missing physics in a model. We do not think they should act as a safety net to achieve a good fit to a spectrum, and the inferred atmospheric properties should be understood in context. However, offsets offer a way to both investigate potential instrumental biases and account for absolute depth uncertainty for each instrument. It is valuable to include offsets as model parameters and marginalize over these possible values in order to understand how the uncertainty in the absolute transit depth for each instrument affects the marginalized posterior distributions of the other model parameters.

5.2. AURA

We complement our analysis of HAT-P-41b by performing retrievals on the STIS, WFC3, and Spitzer observations without the assumption of chemical equilibrium. We employ an adaptation of the retrieval code AURA (Pinhas et al. 2018), as described in Welbanks & Madhusudhan (2019).

The code computes line by line radiative transfer in a transmission geometry and assumes hydrostatic equilibrium. We consider a one-dimensional model atmosphere consisting of 100 layers uniformly spaced in $\log_{10}(P)$ from 10^{-6} – 10^2 bar. The pressure–temperature (P – T) profile in the atmosphere is retrieved using the P – T parameterization of Madhusudhan & Seager (2009). The measured radius of the planet R_p is assigned to a reference pressure level in the atmosphere through a free parameter P_{ref} .

The model atmosphere assumes uniform mixing ratios for the chemical species and treats them as free parameters. We consider sources of opacity expected to be present in hot Jupiter atmospheres (e.g., Madhusudhan 2012) and include H_2O (Rothman et al. 2010), Na (Allard et al. 2019), K (Allard et al. 2016), CH_4 (Yurchenko & Tennyson 2014), NH_3 (Yurchenko et al. 2011), HCN (Barber et al. 2014), CO (Rothman et al. 2010), CO_2 (Rothman et al. 2010), TiO (Schwenke 1998), AlO (Patrascu et al. 2015), VO (McKemmish et al. 2016), and H_2 – H_2 and H_2 –He collision-induced absorption (CIA; Richard et al. 2012). The opacities for the chemical

Table 8
Prior Distributions for AURA Retrievals

Parameter	Prior distribution	Prior range
X_i	Log-uniform	10^{-12} – 10^{-1}
T_0	Uniform	800–2000 K
$\alpha_{1,2}$	Uniform	0.02 – $2.00 \text{ K}^{-1/2}$
$P_{1,2}$	Log-uniform	10^{-6} – 10^2 bar
P_3	Log-uniform	10^{-2} – 10^2 bar
a	Log-uniform	10^{-4} – 10^{10}
γ	Uniform	–20–2
P_{cloud}	Log-uniform	10^{-6} – 10^2 bar
$\bar{\phi}$	Uniform	0–1
STIS _{shift}	Uniform	–500–500 ppm
STIS G430L _{shift}	Uniform/Gaussian	–500–500 ppm/105 ppm
STIS G750L _{shift}	Uniform/Gaussian	–500–500 ppm/85 ppm
WFC3 _{shift}	Uniform/Gaussian	–500–500 ppm/80 ppm

Note. Instrument shifts were employed in a subset of the retrievals and had either uniform or Gaussian priors as explained in Section 8.2. For Gaussian priors, a mean of 0 ppm was assumed, and the sigma is indicated in parentheses.

species are computed following the methods of Gandhi & Madhusudhan (2017). The CO_2 abundance is restricted to remain below the H_2O and CO abundances as expected at these temperatures for H-rich atmospheres (Madhusudhan 2012).

We allow for the presence of clouds and/or hazes following the parameterization in Line & Parmentier (2016) and MacDonald & Madhusudhan (2017). Nonhomogeneous cloud coverage is considered through the parameter $\bar{\phi}$, corresponding to the fraction of cloud cover at the terminator. Hazes are incorporated as $\sigma = a\sigma_0(\lambda/\lambda_0)^\gamma$, where γ is the scattering slope, a is the Rayleigh-enhancement factor, and σ_0 is the H_2 Rayleigh scattering cross section ($5.31 \times 10^{-31} \text{ m}^2$) at the reference wavelength $\lambda_0 = 350 \text{ nm}$. Opaque regions of the atmosphere due to clouds are included through an opaque (gray) cloud deck with cloud-top pressure P_{cloud} .

Finally, we allow for the same three instrumental offset scenarios as described in Section 5.1.2. In these model runs, a constant offset in transit depth is applied to the data set of choice. The offset priors for each scenario are given in Table 8.

In summary, our retrievals consist of up to 25 parameters: 11 chemical species, six parameters for the P – T profile, one for the reference pressure, four for clouds and hazes, and up to three extra parameters for instrumental shifts. Table 8 shows the parameters and priors used in our retrievals.

5.3. A Note on Metallicity and C/O

Atmospheric metallicity is a broadly used term that does not always have the same definition or assumptions built into its derivation (Kreidberg et al. 2014; Madhusudhan et al. 2014b). Here, we define C/O and global atmospheric metallicity explicitly. For AURA, the abundances of different elements are derived independently from the corresponding gaseous absorbers—O/H from oxygen-bearing molecules such as H_2O , Na from gaseous Na, and so on. This retrieval approach allows for different elements to be enhanced or depleted in different quantities. As such, there is no single metric for metallicity in this approach. Nonetheless, as described below, we use the O/H ratio from the retrieved H_2O abundance using AURA as a proxy for metallicity in order to facilitate comparisons with

Table 9
Prior Distributions for Fiducial PLATON Model

Parameter	Symbol	Distribution	Range/Width ^a	Default Value
Planet Radius	R_p	Uniform	0.83–2.48 R_{Jup}^b	1.65 R_{JupAIC}
Limb Temperature	T	Uniform	850–2550 K ^b	1700 K
Carbon–Oxygen Ratio	C/O	Uniform	0.2–2.0	0.53 ^c
Metallicity	Z	Log-uniform	0.1–1000 Z_{\odot}	1 Z_{\odot}
Planet Mass	M_p	Gaussian	0.14 M_{Jup}	0.76 M_{Jup}
Stellar Radius	R_s	Gaussian	0.08 R_{\odot}	1.65 R_{\odot}
Cloud-top Pressure	P_{cloud}	Log-uniform	10^{-3} – 10^8 Pa	1 Pa

Notes.

^a Range for uniform or log-uniform; width is sigma of a Gaussian.

^b Range is 50%–150% of the default value.

^c Solar C/O.

PLATON retrievals in which all the elements are enhanced by a single metallicity parameter.

In PLATON, metallicity is a factor that scales the solar elemental abundances, denoted M/H. The ratios between metals (e.g., Fe/O, Ti/V) are fixed to solar metal ratios, but the solar metal-to-hydrogen ratio is allowed to vary. Thus, all metals are scaled by the same factor. Next, the elemental carbon abundance is determined by C/O \times metallicity. Therefore, only carbon is allowed to differ from its solar ratio compared to other metals. While allowing carbon instead of oxygen to vary is arbitrary, C/O variation is motivated: it is the only metal-to-metal ratio for which predicted molecular abundances are typically sensitive to the wavelength range covered by HST/Spitzer transit spectra. In this paradigm, metallicity is effectively a heuristic for O/H, since that dominates the retrieval both because of molecular opacity (e.g., H₂O, TiO, VO) and because over 99% of the mean molecular weight is due to C, O, and H. In reality, we retrieve O/H and C/O, then set all other elements to $X = X_{\odot} * \frac{\text{O/H}}{\text{O/H}_{\odot}}$. Therefore, PLATON’s derived metallicity is reasonably comparable to AURA’s derived O/H.

6. PLATON Retrieval Analysis

The relative importance of each physical process that affects an observed transit spectrum is not clear ahead of time. PLATON, though less flexible than free-chemistry retrievals or retrievals that allow elemental abundances to vary from solar ratios, is able to quickly perform chemically constrained retrievals (~30 minute runtime for fiducial model retrieval on full data set). This makes it well-suited to testing an array of models, which is important in order to determine how different model assumptions impact the conclusions of a retrieval. To explore this, we choose the fiducial model to be the set of parameters necessary to fully describe the simplest physical processes that we know affect the spectrum: opacity from gas absorption, CIA, and Rayleigh scattering. We then detail the effect of incorporating more complicated physics. In Section 6.3, we use both physical and statistical arguments to determine the “best” model. However, it is important to show the sensitivity of the results to each model assumption, in order to not provide overconfident constraints and to be able to predict how new observations might affect the conclusions.

6.1. Fiducial Model

The priors for the fiducial model are shown in Table 9. The metallicity of the atmosphere, the temperature of the limb, and the C/O ratio are necessary to include in order to calculate molecular and atomic abundances, which determine the gas absorption, CIA, and Rayleigh scattering opacities. While we cannot improve constraints on M_p and R_s , it is best practice to include them as parameters with Gaussian priors in order to propagate the uncertainties on those measurements (Zhang et al. 2019). Otherwise, we would mistakenly assume that M_p and R_s are precisely known. We also include the cloud-top pressure of a gray cloud deck in the fiducial model. We fix the reference pressure to 1 bar and retrieve the planetary radius at that pressure. Note that Welbanks & Madhusudhan (2019) demonstrated that it is justified to assume a reference pressure and retrieve the planetary radius without affecting the ability to constrain atmospheric composition.

Although R_p and M_p are both allowed to vary independently in PLATON retrievals, their uncertainties are not independent: the uncertainties for M_p are derived from $\log(g_p)$ (from transit observations) and R_p (derived from R_p/R_s from transit and R_s from TIC-8). PLATON does not constrain M_p and R_p^2 to match $\log(g_p)$ *a priori*, but this is only an issue if regions of high likelihood extend to combinations of values that should not be allowed (i.e., more than 3σ from observed $\log(g_p)$). We rederive $\log(g_p)$ using values at R_p and M_p at the edge of significant likelihood, and find good agreement with the prior, well within 3σ .

We use uninformative priors where appropriate in order to fully explore the possible parameter space. For quantities that can range over many orders of magnitude, such as the cloud-top pressure or the metallicity, this means a log-uniform prior is necessary to avoid bias toward higher values. Otherwise—for T_{limb} , R_p , and C/O—we use uniform priors with limits either set by the functionality of the code (e.g., C/O) or conservatively derived from previous observations. Widening the prior for any parameter in the fiducial model has no significant effect on the result of the retrieval.

The retrieved median fiducial model with uncertainty contours is shown with the observed spectrum in Figure 5. The model is an excellent fit (reduced $\chi^2 = 1.09$; consistent with the χ^2 of the true model for 70 degrees of freedom to 1σ). We clearly detect water vapor ($>5\sigma$ significance) via the 1.4 μm water band in the WFC3 data. The bump in the STIS data is indicative of TiO, and the lack of any optical slope or flat line indicates that gray clouds and scattering haze do not

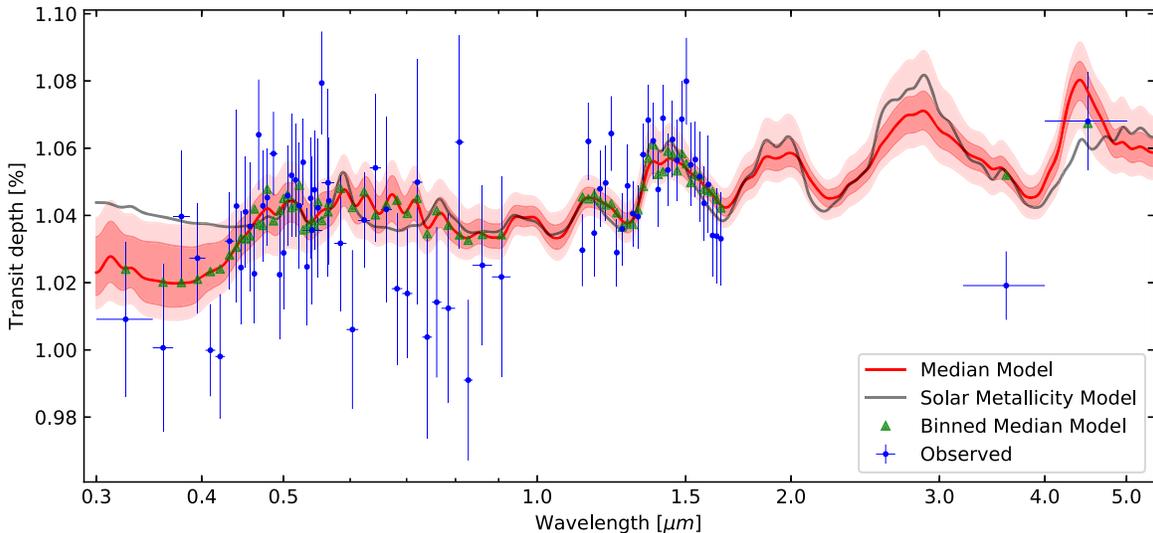


Figure 5. Median retrieved model with 1σ and 2σ uncertainty contours for the fiducial model. Also shown is the median retrieved model when metallicity is fixed to solar. Median model and uncertainties are derived by generating 100 samples from the correctly weighted posterior and calculating the depth at each bin for each sample. Contours are given by the 2nd, 16th, 50th, 84th, and 98th percentile depths at each bin. Continuous model is smoothed with a Gaussian filter with $\sigma = 15$, which approximates the resolution of HST WFC3 (Zhang et al. 2019).

(The data used to create this figure are available.)

contribute significant opacity in the planet’s spectroscopically active region. The difference between the two Spitzer points is attributed to CO_2 , though because they are photometric observations, we do not resolve any features.

The posterior probability distribution is represented by the corner plot (Foreman-Mackey 2016)²¹ in Figure 6. This figure provides the marginalized posterior distribution for each parameter (with median, 16th, and 84th percentile values indicated by vertical dashed lines), as well as every two-dimensional projection of the posterior (with 0.5, 1, 1.5, and 2σ contours) to reveal covariances. Well-constrained parameters have narrow distributions with clear peaks, and slanted or diagonal shapes are indicative of correlated sets of parameters (e.g., the R_p - R_s shape). We find a supersolar metallicity ($259^{+174}_{-114} \times$ solar metallicity, henceforth Z_\odot), a likely subsolar C/O ($\text{C/O} < 0.6$) that is consistent with stellar C/O (0.19), a clear atmosphere ($P_{\text{cloud}} > 0.5$ mBar), and $T_{\text{limb}} = 1650^{+70}_{-120}$ K. The temperature is driven primarily by the STIS data, mostly because PLATON interprets the bump in the STIS data as a metallic oxide, which is only the dominant opacity source above ~ 1500 K in chemical equilibrium. Below ~ 1500 K, the optical spectrum would be dominated by an atomic sodium line, and this is not seen in the data. The upper limit on C/O is related mainly to the H_2O : in equilibrium chemistry for $T \sim 1650$ K and $P \sim 1$ bar, H_2O opacity decreases exponentially when $\text{C/O} > 0.6$ (Madhusudhan 2012). Any model with $\text{C/O} > 0.6$ struggles to capture the water feature and relatively high infrared baseline opacity (compared to the optical) and is thus a poor fit to the data.

The high metallicity is constrained by the size of both the STIS and WFC3 features, as well as the lack of a significant Rayleigh scattering slope. While the metallicity affects chemistry, it is primarily constrained via its effect on the mean molecular mass of the atmosphere. Increasing metallicity, by definition, increases the ratio of metals to hydrogen, which

increases the atmosphere’s mean molecular mass. This lowers the atmospheric scale height and consequently decreases the predicted feature size. The equation for approximate feature size, δ_λ (Kreidberg 2018), where μ is the mean molecular mass, clarifies its dependencies:

$$\delta_\lambda \propto \frac{TR_p}{\mu g_p R_s^2} \propto \frac{TR_p^3}{\mu M_p R_s^2}. \quad (1)$$

PLATON can decrease the feature size by changing R_p or R_s , but both are well-constrained by the continuum baseline as well as priors and thus are relatively fixed. It can also be lowered by decreasing the temperature, increasing the metallicity (and thus the mean molecular weight), or by increasing M_p . The temperature is strongly constrained by chemistry, and M_p is constrained by previous observations, so only metallicity can vary enough to explain the observed feature sizes. Note that this relationship explains the correlations between M_p , T_p , and metallicity seen in Figure 6: as mass increases or temperature decreases, metallicity decreases, because a lower value is necessary to achieve the scale height that predicts the observed feature sizes. For reference, the median derived mean molecular weight is 5.8 AMU and the derived scale height is roughly 322 km.

At solar metallicity, the model predicts features that are much larger than what the data show. Consequently, solar metallicity atmospheres in the fiducial model can only explain the observed feature by invoking a cloud to mask the troughs of the features. The median retrieved model for metallicity fixed to solar is shown in Figure 5. Note that this model is dispreferred by 3.5σ , since the cloud leads to a poor fit to the bluest transit depths.

The same metallicity value is an excellent fit for both the water feature in the WFC3 data and the TiO feature in the STIS data. Retrieving only on the STIS data or only on the WFC3 data recovers supersolar atmospheric metallicities. Additionally, due to predicting a greater abundance of CO_2 , it is better

²¹ <https://github.com/dfm/cornerpy>

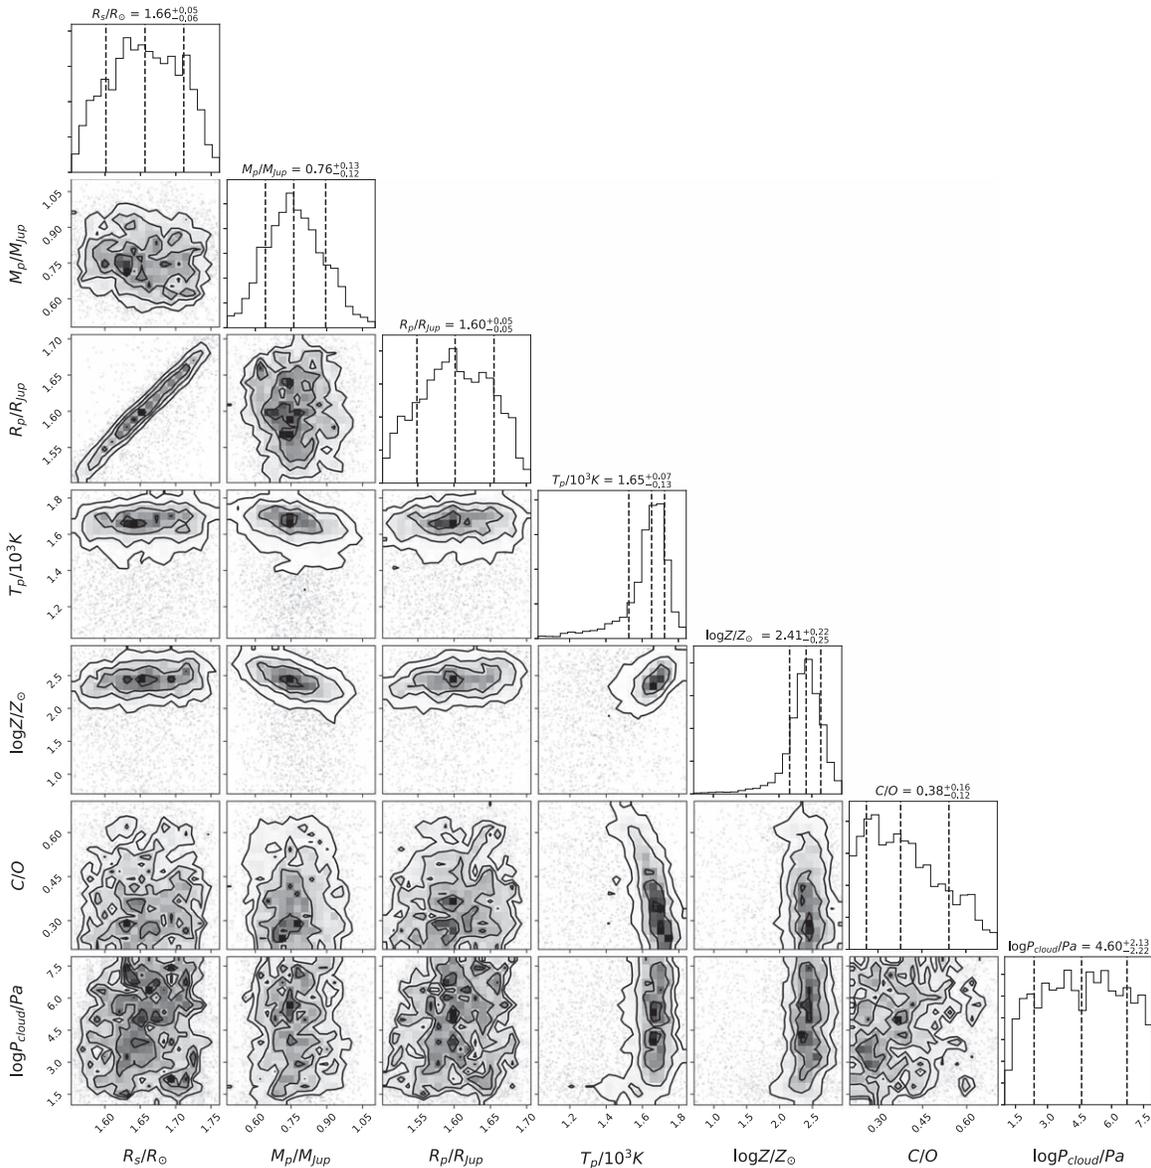


Figure 6. Corner plot illustrating posterior probability distributions from PLATON for the fiducial model. The 16th, 50th, and 84th percentile values are indicated by vertical dashed lines and stated in the title of each parameter’s 1D marginalized posterior distribution. Contours indicate the joint 0.5, 1, 1.5, and 2 σ levels for each 2D distribution. The 1 σ metallicity range is $\log_{10} Z/Z_\odot = 2.41^{+0.23}_{-0.25}$ (145–437 \times solar metallicity), the isothermal limb temperature is well-constrained around 1650 K, and the C/O ratio is likely subsolar. Spectrum is consistent with a cloud-free atmosphere, and the marginalized posterior distributions of M_p and R_s are dominated by their priors. The slight correlations between T – $\log Z$ and M_p – $\log Z$ are due to their relation in the scale height equation.

than low-metallicity solutions at explaining the large change in depth between the Spitzer points. Observations from all three instruments support the high-metallicity solution.

6.2. More Complex Models

In this section, we incorporate additional model parameters to explore whether having more complex physics impacts the inferred atmospheric parameters. We demonstrate the insensitivity of our results to model assumptions.

6.2.1. Partial Cloud Coverage

Line & Parmentier (2016) showed that partial cloud coverage (i.e., clouds at the same height but not uniformly covering the limb azimuthally) could mimic the effect of a high

mean molecular mass atmosphere for WFC3 spectra. When partial clouds are present, the observed spectrum would be the weighted average of the cloudy and clear spectra. The transit depth of an atmosphere dominated by gray clouds does not vary with wavelength, and so the cloudy spectrum is a straight line. Averaging a clear spectrum with molecular features and a cloudy, flat spectrum reduces the size of the features by an amount proportional to the cloud fraction. Given that we find a significantly supersolar metallicity, we investigate whether this possible mean molecular mass–cloud fraction degeneracy affects the results from the fiducial case.

When fit independently and allowing the cloud fraction to vary, both WFC3 and STIS spectra retrievals no longer constrain the metallicity to be supersolar. However, fitting the infrared and optical data jointly breaks this degeneracy, as

Table 10
Prior Distributions for More Complicated PLATON Models

Parameter	Symbol	Distribution	Range/Width ^a	Default Value
Cloud Fraction	f_c	Uniform	0–1	1
Scattering Slope	γ	Uniform	–2–20	4
Scattering Factor	a_0	Log-uniform	10^{-4} – 10^8	1
WFC3 Offset	...	Uniform/Gaussian	–500–500 ppm/80 ppm	0 ppm
STIS G430L Offset	...	Uniform/Gaussian	–500–500 ppm/105 ppm	0 ppm
STIS G750L Offset	...	Uniform/Gaussian	–500–500 ppm/85 ppm	0 ppm
STIS Offset	...	Uniform	–500–500 ppm	0 ppm
Stellar Effective Temperature	T_{star}	Fixed	0	6480 K
Faculae Temperature	T_{fac}	Fixed ^b	0	6580 K
Faculae Covering Fraction	f_{fac}	Uniform	0–0.10	0
Mie Particle Size	r_{part}	Log-uniform	0.01–1 μm	0.1 μm
Mie Number Density	n	Log-uniform	10 – 10^{15} m^{-3}	10^5 m^{-3}
Fractional Scale Height	$H_{\text{cloud}}/H_{\text{gas}}$	Log-uniform	0.1–10	1.0

Notes.

^a Range for uniform distribution; width is sigma of a Gaussian.

^b $T_{\text{fac}} = T_{\text{star}} + 100$ K (Rackham et al. 2019).

predicted by Line & Parmentier (2016). Effectively, a low mean molecular mass and nonuniform cloud solution should be impacted by Rayleigh scattering in the optical data, especially the bluest six wavelength bins. The dominance of gas absorption opacity over Rayleigh scattering opacity in the STIS data disallows this solution, breaking the degeneracy in favor of the high mean molecular mass solution.

However, it is possible that removing assumptions made in the fiducial model—such as fixed Rayleigh scattering or no instrumental offsets—could muddle this decisive degeneracy break and allow for a low-metallicity solution. We investigate this below.

6.2.2. Parametric Rayleigh Scattering

The fiducial model assumes Rayleigh scattering. In lieu of complicated microphysics, PLATON allows parametric scattering, in which the slope and the magnitude of Rayleigh scattering vary in order to capture the possible signature of many hazes. For a more detailed explanation, see Zhang et al. (2019). Though there is no obvious signature of haze in the optical data (i.e., no linear slope decreasing with increasing wavelength), it is worth exploring whether loosening the assumption of exact Rayleigh scattering affects the results.

Allowing the full scattering parameter space (see Table 10) has little effect: the clear lack of slope in the STIS data conclusively leads to a haze-free atmosphere. Further, the median scattering factor is 0.01, implying that the data is easiest to fit when opacity from Rayleigh scattering is muted. This complicates the mean molecular weight–cloud fraction ($\mu - f_c$) degeneracy. Lower values of μ are now possible, since the model no longer expects scattering opacity to be important at optical wavelengths. The lower the magnitude of Rayleigh scattering—and the shallower the scattering slope—the lower μ can be. This is because decreases in Rayleigh opacity allow for gas absorption to still be dominant at larger scale heights. As a result, a patchy cloud and low-metallicity solution is viable. Though possible, the low μ solution requires a specific combination of cloud-top pressure, scattering slope, scattering factor, and cloud fraction, and does not improve the fit. Therefore, it is much less likely than the high-metallicity solution. The marginalized posterior probability distribution for

metallicity has the same maximum likelihood value as the fiducial model. The difference is that the distribution has a tail extending to lower metallicities (Figure 7). The resulting median log metallicity and 1σ range (as determined by the 16th and 84th percentile values) is $\log_{10} Z/Z_{\odot} = 2.34^{+0.27}_{-0.64}$.

Though all cloud fractions and cloud pressures are allowed, the posterior is consistent with a clear atmosphere due to the likelihood desert in the upper left corner of the cloud fraction–cloud-top pressure pairs plot: clouds are only seen above the altitude corresponding to the ~ 10 Pa pressure level at fractions below 0.50.

6.2.3. Instrumental Offsets

We model an instrumental offset as a constant value added to the forward model’s binned transit depth in the wavelength range of the instrument of interest.

For the physically motivated scenario (Scenario 1 from Section 5.1.2), we set the priors for STIS G430L, STIS 750L, and WFC3 to be Gaussians centered on zero ppm, with widths set to the uncertainty on the transit depth from their white light curves (105, 85, and 80 ppm, respectively; see Table 10).

The retrieved median WFC3 offset is nontrivial, with a median of about $1.5\times$ the white light uncertainty (130 ± 50 ppm). The offsets in the STIS G430L and G750L are less significant, at 58 ± 58 ppm and -65 ± 55 ppm, both well within their white light uncertainties. However, there is a significant median offset of ~ 120 ppm between the two instruments. This is driven primarily by the retrieval attempting to align transit depths in the overlapping wavelength region between the instruments.

The Spitzer 3.6 μm point drives the WFC3 offset: shifting the WFC3 depths down necessitates a smaller radius ratio, which better captures the relatively low transit depth at 3.6 μm . The ability to better capture the Spitzer 3.6 μm results in a greater Bayesian evidence, indicating that this model is strongly preferred over the fiducial model (for a more detailed discussion, see Section 6.3).

When combined with partial clouds, the instrumental offsets cause a small decrease in the retrieved median metallicity ($\log_{10} Z/Z_{\odot} = 2.33^{+0.23}_{-0.25}$). This is for reasons similar to those explained in the parametric scattering section: whereas

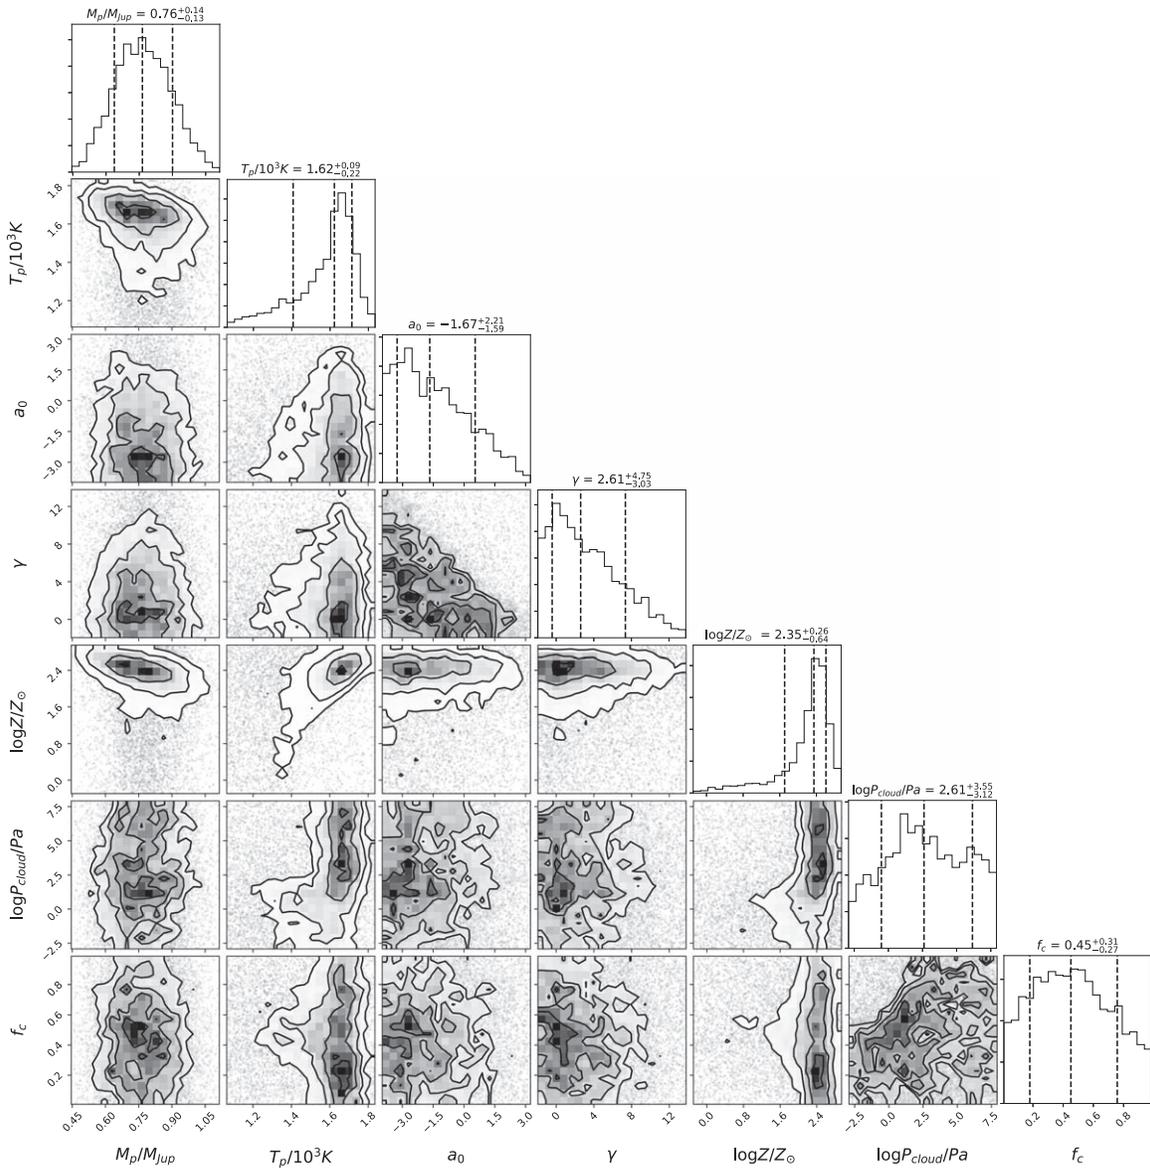


Figure 7. Corner plot illustrating the posterior probability distribution from PLATON for the partial cloud and parametric scattering case. For clarity, R_p , R_s , and C/O are not shown, since their marginalized posterior distributions are the same as in the fiducial case. The 1σ metallicity range is shifted down to $\log_{10} Z/Z_{\odot} = 2.34^{+0.27}_{-0.64}$. Note that, at $f_c = 0$ or $\log P_{\text{cloud}} > 10^{2.5}$, we recover the fiducial marginalized posteriors.

parametric scattering justified the absence of an optical scattering slope in the low-metallicity solution by effectively removing Rayleigh scattering opacity, the instrumental offset model can decrease the WFC3 depths relative to the STIS depths in order to artificially allow for it.

Note that increasing STIS depths (instead of decreasing WFC3 depths) has the same effect on Rayleigh opacity and thus metallicity. However, it is not a viable solution: this is because, unlike decreasing WFC3 depths, it does not improve the forward model’s ability to capture the low Spitzer $3.6 \mu\text{m}$ point.

Allowing Gaussian-prior instrumental offsets had no significant effect on the results. However, it is not impossible that there is some unknown wavelength-independent systematic that biases the absolute transit depths of the instruments relative to one another. Though it is unlikely, we chose to explore this by allowing offsets in the STIS G430L, STIS G750L, and WFC3 data to vary by about 5% (500 ppm) in either direction

(Scenario 2 from Section 5.1.2). Due to the model preferring a lower radius ratio to best explain the Spitzer $3.6 \mu\text{m}$ point, the median WFC3 offset is a 250 ppm decrease, about $3\times$ the white light uncertainty. Surprisingly, this large offset does not significantly change the 1σ ranges for metallicity ($\log_{10} Z/Z_{\odot} = 2.26^{+0.24}_{-0.40}$). The size of the molecular features and the large differential between Spitzer photometric points drive the supersolar metallicity in this case. Regardless of the magnitude of the offset, we retrieve a high metallicity. The two large, uniform offsets case, where both STIS instruments are offset by the same amount (Scenario 3 from Section 5.1.2), retrieves posterior distributions effectively identical to those of Scenario 2.

While it is worthwhile to understand the effect on the retrieval, there is no reason to expect such large instrumental offsets for HAT-P-41b. A transit depth offset can be caused by the necessity of analyzing each instrument differently. For example, not handling limb darkening consistently and not

using consistent orbital parameters (i.e., inclination) for each analysis might cause an offset, but this is easily fixed and is not an issue for our data set. Since the instruments’ observations are from different dates, it is also possible that stellar variability could cause an offset. However, we have long-term photometry (Section 2.3) that shows no such variability. Additionally, the STIS depths are in good agreement with HST UVIS observations (Wakeford et al. 2020). There is no indication that this particular observation is biased in any way, and unresolved companions are confidently ruled out (Evans et al. 2018). The most plausible source is unaccounted-for uncertainties or bias from the spectral analysis, as the WFC3 spectrum derived in this paper is shifted up ~ 90 ppm relative to the literature spectrum (Tsiaras et al. 2018), as noted in Section 4.2.2. However, this is still well below the 250 ppm value preferred by the large, uniform offset models. We determine that offsets beyond the physically motivated values are unlikely.

6.2.4. Stellar Activity

Section 3.1 demonstrated that HAT-P-41 is consistent with a quiet star and stellar activity is not expected to impact the transit spectrum. However, to be conservative, we investigated whether allowing for greater stellar variability impacted our conclusions.

The typical signature of unocculted, cool starspots is to mimic a haze-like slope in the transit spectrum, and such a signature is clearly absent in the derived transit spectrum of HAT-P-41b. On the other hand, the signature of hot faculae is a steep optical drop-off toward shorter wavelengths (Rackham et al. 2019). Given that we see a drop in transit depths in the optical, the retrieval could plausibly be affected if faculae dominate over star spots, and so we focus on a faculae overabundance.

We assumed that the temperature of the stellar photosphere equals the stellar effective temperature. We modeled the faculae following the prescription from Rackham et al. (2019), and accordingly fixed the faculae temperature to $T_{\text{phot}} + 100$ K. PLATON weights the contributions from the different temperature regimes via the fractional coverage parameter, which represents the overabundance of faculae in the unocculted regions. Rackham et al. (2019) states that moderately active F5V-dwarfs will have around 1% faculae coverage, and up to about 7% on the more active end. This is the faculae fraction, which is much higher than the faculae overabundance. However, we set a conservative uniform prior on the fractional coverage of 0%–10% in order to determine whether high activity would significantly alter our conclusions.

We find that including stellar activity has no effect on the posterior probability distribution. It may seem that the STIS data could be explained by a featureless flat line and stellar activity instead of a TiO feature. However, the overabundance of faculae necessary to explain the drop in the bluest six points (0.32–0.42 μm) produces a poor fit to the rest of the STIS data. Therefore, even when including stellar variability, TiO is necessary to explain the STIS depths. Allowing a wider range of faculae temperatures also had no effect.

In summary: there is no evidence of stellar variability from prior observations, and allowing for activity does not affect the retrieval.

6.2.5. Mie Scattering

Benneke et al. (2019) recently invoked Mie scattering to explain anomalously low Spitzer transit depths. Given the relatively low value of HAT-P-41b’s Spitzer 3.6 μm depth relative to the rest of the spectrum—the fiducial model’s predicted depth at 3.6 μm is about 3.3σ away from the observed depth—we included Mie scattering in our analysis. In PLATON, Mie scattering can be used in lieu of parametric Rayleigh scattering.

Each condensate is described by a wavelength-dependent complex refractive index, $n-ik$, where n is the real part and k is the imaginary part of the index. This index explains how that particular condensate interacts with light with wavelengths similar to the particle size. PLATON assumes log-normal distribution in particle size, with geometric standard deviation 0.5, to determine the abundance of different radii condensates for a given mean particle size (Zhang et al. 2019). The other relevant factors are cloud height (condensates are only relevant above that pressure; below it, gray cloud opacity dominates), particle density at the cloud-top pressure, and condensate scale height. The condensate scale height is parameterized as a fraction of the gas scale height, and it describes how the abundance of Mie scattering particles decreases with height. The refractive index is fixed for a given condensate, and the other four parameters are fit for in the retrieval (see Table 10).

PLATON tests one Mie scattering species at a time. Only a few species expected to form clouds in hot Jupiter atmospheres have condensation temperatures above HAT-P-41b’s limb temperature ($T \sim 1600$ K; Wakeford & Sing 2015). Those five (SiO_2 , Al_2O_3 , CaTiO_3 , FeO , and Fe_2O_3) fall into two phenotypes: “low- n ” with real refractive index $n \approx 1.5$, and “high- n ” with $n \approx 2.5$. Though the k values vary more significantly, we find that they do not have a significant impact on the absorption cross section of the condensates. We use n and k values from Kitzmann & Heng (2018), which we average over our wavelength range (0.3–5 μm). The n values are flat over this range, and so the average is an excellent approximation. We tested retrievals with both the low- n (corundum; Al_2O_3) and high- n (hematite; Fe_2O_3) phenotypes.

The priors for the fittable parameters are shown in Table 10. The prior for cloud-top pressure is the same as the fiducial model. Since the condensate radii must be such that they cause a relative drop in opacity around 4 μm (i.e., they increase opacity in the optical by more than in the mid-infrared), we can constrain the mean particle size reasonably well. We set the prior to be log-uniform with a range that contains all plausible values. The number density is not known ahead of time, so we set an uninformative log-uniform prior; widening the prior further did not affect the retrieval. Finally, it is unclear what physical constraints there are on condensate scale height. Fortney (2005) finds that condensate scale heights can be one-third of the gaseous scale height for hot Jupiters, and Benneke et al. (2019) found $H_{\text{part}}/H_{\text{gas}} \approx 3$ for a sub-Neptune. Using these values as guides, we set a conservative uniform prior on the fractional scale height and constrain it to be in the range 0.1–10.

Including Mie scattering opacity does not noticeably affect the results of the retrieval, and the Mie scattering parameters are not constrained by the retrieval. The inferred small gaseous scale height—which dampens features and is necessary to explain STIS and WFC3 feature sizes—makes it difficult to

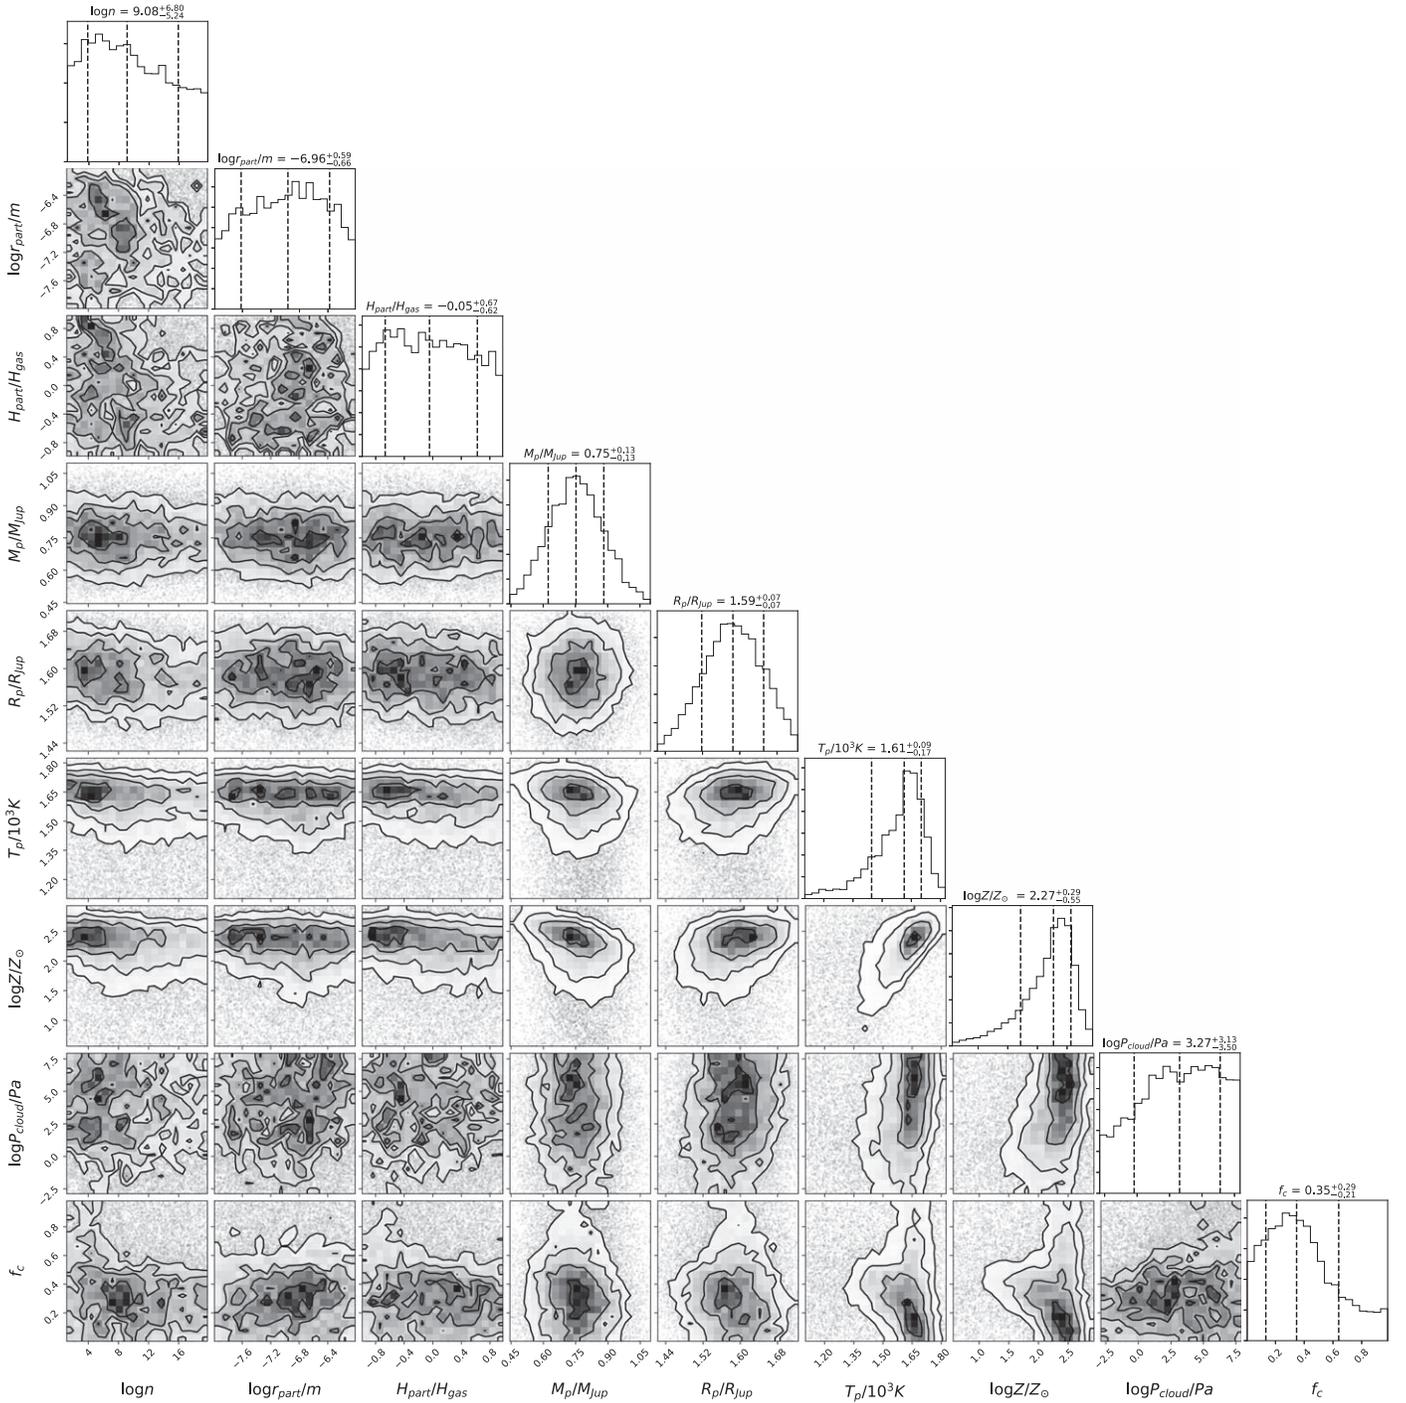


Figure 8. Corner plot illustrating posterior probability distribution from PLATON for the fiducial plus Mie scattering and partial clouds model. For clarity, R_p , and C/O are not shown, since their marginalized posterior distributions are the same as in the fiducial case. The 1σ metallicity range is shifted down a bit to $\log_{10} Z/Z_{\odot} = 2.27^{+0.30}_{-0.55}$. Note that we recover the fiducial marginalized posteriors when any of the mean particle size, the particle number density, or the condensate scale height are too small.

explain the large variations in the radius ratio. Combining Mie scattering with partial clouds—physically, an atmosphere with patchy clouds and Mie scattering particles distributed only above those clouds—alleviates the issue of explaining the large transit depth variation. Since partial clouds allow for higher scale heights, Mie scattering could then cause a larger drop in transit depth near Spitzer without needing to invoke an unreasonably high fractional scale height.

Figure 8 shows the corner plot for this model. Though the Mie scattering parameters are not constrained, at number densities above $\sim 10^8 \text{ m}^{-3}$, particle radii around $0.15 \mu\text{m}$, and condensate scale heights greater than the gaseous scale height, lower metallicity and temperature values are possible. This is because added Mie opacity tends to mute features near its peak opacity. This provides a physical reason to expect smaller spectral features, and so a less small scale height is necessary to

Table 11
Evidence and Metallicity Ranges for Each Plausible PLATON Model

Model	$\log_{10} Z^a$	Z^b	$\ln Z^c$	Weight	\mathcal{O}^d	Interpretation ^e
Fiducial (F)	$2.41^{+0.23}_{-0.25}$	145–437	551.6	0.0	Ref	Default model
F + Partial Clouds (PC)	$2.38^{+0.22}_{-0.37}$	102–398	551.7	0.0	1.1	Inconclusive
F + PC + Parametric Scattering	$2.34^{+0.27}_{-0.64}$	50–407	550.7	0.0	0.4	Inconclusive
F + PC + Stellar Activity	$2.41^{+0.27}_{-0.41}$	100–479	551.8	0.0	1.3	Inconclusive
F + Mie Scattering	$2.41^{+0.24}_{-0.29}$	132–447	551.0	0.0	0.6	Inconclusive
F + PC + Mie	$2.27^{+0.3}_{-0.55}$	52–372	551.8	0.0	1.3	Inconclusive
F + PC + 3 Gaussian Offsets	$2.33^{+0.23}_{-0.25}$	120–363	556.4	0.33	122	Strongly preferred
F + PC + 3 Uniform Offsets	$2.26^{+0.24}_{-0.4}$	72–316	556.9	0.58	213	Strongly preferred
F + PC + 2 Uniform Offsets	$2.30^{+0.26}_{-0.39}$	81–363	554.9	0.08	29	Moderately preferred

Notes. Only the models including instrumental offsets are preferred over the fiducial model. No model assumption changes the conclusion of a supersolar atmospheric metallicity.

^a Median log metallicity with 16% and 84% quantiles, in units of log solar metallicity.

^b 68% credible interval for metallicity, in units of solar metallicity.

^c Natural log of Bayesian Evidence.

^d Odds ratio between model and the fiducial model.

^e According to Jeffreys’ scale (Trotta 2008).

fit the features. The net impact is a decreased—but still supersolar—median metallicity of $\log_{10} Z/Z_{\odot} = 2.27^{+0.30}_{-0.55}$.

6.3. Model Selection

Section 6.2 stepped through the PLATON retrieval for increasingly complex models, examining both how each additional parameter affected the posterior and why it affected it in that way. While knowing the effect of each model assumption is useful, it is important to determine a preferred model in order to effectively convey the results. In this section, we use Bayesian model comparison to select the best model.

Model selection is as important as parameter estimation in atmospheric retrievals. We determine the preferred model by using a combination of physical arguments and Bayesian statistics. Specifically, we check whether it is necessary to consider more complicated physics using the odds ratio, which is the Bayes factor between models (defined as the ratio of their evidences) multiplied by their prior probability ratio. The prior probability ratio is typically assumed to be one (i.e., the models are assumed to be equally likely). The odds ratio determines whether one model should be preferred over another by intrinsically rewarding better fits while punishing overcomplicated structure (Trotta 2008). This is entirely defined by the data and model, assuming appropriately uninformative priors are used. We compare the Bayesian evidences of each model in order to determine which should be favored.

The Bayesian evidences and 1σ metallicity ranges for every model discussed in Section 6 are shown in Table 11. The 1σ range is represented by the median metallicity with quantiles (i.e., the central 68% of metallicity values). The 1σ metallicity ranges are included to illustrate the uncertainty caused by model choice. The retrieved atmospheric metallicities are remarkably consistent across the models, and a supersolar metallicity is ubiquitous. This demonstrates that, under PLATON’s assumptions, supersolar metallicity is a robust conclusion.

Figure 9 emphasizes the insensitivity of the atmospheric parameters to model assumptions. This shows the one-dimensional marginalized posterior distributions for metallicity, temperature, and C/O for five of the models we examined.

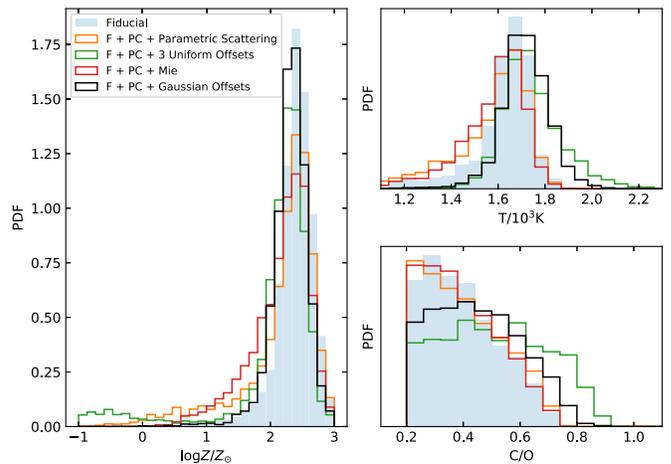


Figure 9. Marginalized posterior distributions for metallicity, temperature, and C/O from select models compared to solar values. Note that stellar C/O = 0.19 (Table 4). Though some models have low-metallicity tails, the 68% credible interval for metallicity is robust (Table 11). Offsets allow for higher temperatures and C/O ratios, while both parametric and Mie scattering allow for lower temperatures.

These specific models are shown because they are “interesting” in that they differ from the fiducial model’s posteriors the most. We emphasize that these are these are the models that most differ from the fiducial case. While including instrumental offsets tends to flatten the distributions, the peaks of all of the models are astoundingly consistent.

We define the columns relevant to model selection here:

1. **$\ln Z$** is the natural log of the Bayesian evidence. A higher value indicates the model is better able to describe the data without overfitting.
2. **Weight.** Weight assigned to each model for Bayesian model averaging based on Bayesian evidence and prior probability (Section 6.3.1). This can be roughly interpreted as the probability of each model.
3. **\mathcal{O}** is the odds ratio in favor of a model over the fiducial model. It is the product of their Bayes factor and their prior probability ratio. The prior probability ratio is often assumed to be one, as is the case here. This can be

directly interpreted: an odds ratio of 100 indicates 100:1 odds in favor of the more complex model. Values less than one indicate evidence against the corresponding model.

4. **Interpretation.** This is the empirically derived interpretation of the odds ratio based on the Jeffreys’ scale (Trotta 2008).

Table 11 contains every notable model we considered. We did not do an iterative combination of every model scenario, for two primary reasons. Most importantly, we are wary of overfitting the data. The fiducial model is already an excellent fit to the data ($\chi^2_\nu = 1.09$), so we must be careful about adding complications. Layering multiple parameter physical processes, such as Mie scattering and stellar activity, involves an extra five parameters and significantly overcomplicates the fit. Instead, we only combine complications when there is a physically motivated reason to do so, e.g., partial clouds. The second reason is computational difficulty. Some model combinations have enough free parameters to describe the data with many different combinations, and so the retrieval does not converge on the timescale of weeks. Mie scattering combined with offsets falls in this category. However, given the overfitting concerns and the lack of evidence for just Mie scattering, we do not think this is a worrying omission.

It is generally best practice to assume the simplest model unless there is evidence in favor of extra parameters. That is why we list the fiducial model as the reference, and determine the evidence of the more complicated models. If the evidence of the model with extra parameters is not significantly greater, it means that the ability to explain to data was not improved enough to justify the added complexity. This essentially quantifies Occam’s razor.

Following this logic, we determine the “fiducial + partial clouds + 3 Gaussian offsets” model to be the best model. Only the Gaussian offset and uniform offset models are preferred over the fiducial model. While the three uniform offsets model has the highest evidence/weight, the odds ratio between that and the Gaussian offsets model is 1.75. This is inconclusive on the Jeffrey’s scale, meaning we are unable to distinguish between the models by evidence alone. Instead, we favor the Gaussian offsets model as more plausible, since its Gaussian priors are physically motivated by common-mode corrections.

The evidence for partial clouds is inconclusive; however, partial clouds are more plausible than assuming 100% cloud coverage, as argued by MacDonald & Madhusudhan (2017) as well as Welbanks & Madhusudhan (2019). Therefore, to be conservative, we choose the model that accounts for partial clouds as the “best” model.

The odds ratio only works in a direct comparison of two models and is not a statement on the absolute goodness-of-fit. The reduced chi-squared test statistic is a useful sanity check to ensure that the model is able to explain the variance in the data. The value for the best model is an ideal $\chi^2_\nu = 1.0$. The results section (Section 7)—and the abstract values—are based on parameter estimation from the “fiducial + partial cloud + Gaussian offsets” model.

6.3.1. Bayesian Model Averaging

Instead of model selection, it is possible to take a weighted average of the results from each model and therefore

automatically take their respective evidences into account (Gibson 2014; Wakeford et al. 2016, 2018). The benefit of Bayesian model averaging is the ability to quantify uncertainty in model selection, as well as avoiding having to arbitrarily choose between models with slightly different evidences. However, it requires a few assumptions: it is only valid if the set of models comprises the full model space, i.e., at least one model is a good description of the data. The weight-averaged uncertainties assume Gaussian-distributed posteriors, which is not strictly correct. However, it is useful in combining information from every model.

Here, we show the assumptions we make to use Bayesian model averaging. The χ^2_ν values for the models we tested are clustered around one, so it is fair to assume that a “correct” model is contained in the set. Figure 9 shows that, although the posteriors are not perfectly Gaussian, they have sharp, unimodal peaks, and so the uncertainty derived from marginalization is informative.

The model weights are defined by Equation 2 (adapted from Gibson (2014)). Here, W_q is the weight assigned to model q , $P(M_q|D)$ is the the likelihood of model q given the data, and $P(D|M_q)$ is the likelihood of the data given model q , which is equivalent to the Bayesian evidence of model q , E_q . The denominator is a normalization term, summed over N models. We assume a conservative prior that each model is equally likely ($P(M_i) = 1$ for all i).

$$\begin{aligned} W_q &= \frac{P(M_q|D)}{\sum_{i=1}^N P(M_i|D)} = \frac{P(D|M_q)P(M_q)}{\sum_{i=1}^N P(D|M_i)P(M_i)} \\ &= \frac{E_q P(M_q)}{\sum_{i=1}^N E_i P(M_i)}. \end{aligned} \quad (2)$$

The marginalized log metallicity with 1σ uncertainties is calculated from Equations (15) and (16) from Wakeford et al. (2016). The result is $\log_{10} Z/Z_\odot = 2.29^{+0.24}_{-0.36}$ ($194^{+144}_{-109} \times Z_\odot$). As expected, the highest-weighted models are the offset models. Bayesian model averaging demonstrates that, in PLATON’s chemical equilibrium framework, a supersolar metallicity is the most likely result even after accounting for uncertainty in model selection.

The marginalized metallicity is useful as a reference, but it is valuable to give the metallicity distribution for each specific model assumption. Marginalization is most appropriate when the specific model parameters are unimportant; however, we are interested in the impact that modeling assumptions have on the atmospheric parameters. We emphasize that, even for apparently “data-defined” methods, many assumptions have to be made, and those should be explicitly stated for an appropriate interpretation.

7. Results for the Favored PLATON Model

In Section 6.3, we argued that the best PLATON model scenario is the fiducial model with partial clouds and physically motivated, Gaussian prior instrumental offsets added. The retrieved median spectrum with uncertainty contours is shown in Figure 10. It is an excellent fit to the data, with $\chi^2_\nu = 1.0$. In this section, we discuss the details of the retrieved atmospheric parameter values.

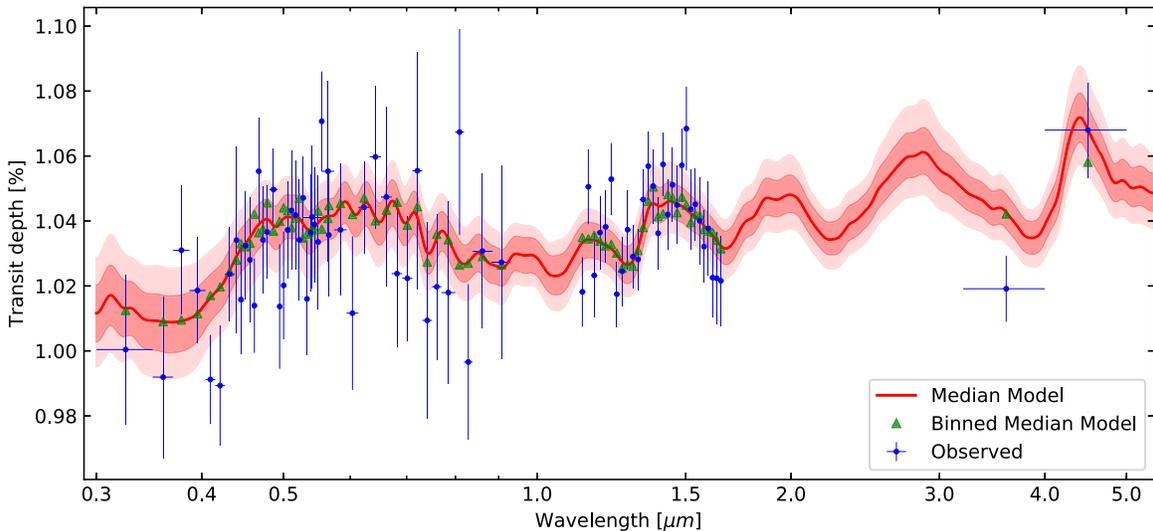


Figure 10. Median retrieved model with 1σ and 2σ uncertainty contours for the favored PLATON model (fiducial model with partial clouds and instrumental shifts with physically motivated Gaussian priors).

(The data used to create this figure are available.)

7.1. Summary of Retrieved Parameters

The posterior distribution corner plot is shown in Figure 11. Both atmospheric metallicity and temperature are well-constrained, and the C/O ratio, though relatively flat, has a strict upper limit. The median retrieved metallicity is supersolar ($\log_{10} Z/Z_{\odot} = 2.33^{+0.23}_{-0.25}$), and solar metallicity is inconsistent to 3σ (lower limit $4.8 \times Z_{\odot}$). As noted in Section 5.3, PLATON’s derived metallicity is a proxy for [O/H], enabling comparison with its host star’s oxygen abundance ([O/H] = 0.37; Table 4). PLATON determines HAT-P-41b to be metal-enriched relative to its host star ($\log_{10} Z/Z_{\text{star}} = 1.97^{+0.23}_{-0.25}$), and it is inconsistent with the stellar metallicity to 3σ (lower limit $2.1 \times Z_{\text{star}}$). The planetary C/O ($0.44^{+0.18}_{-0.15}$) has a 3σ upper limit of 0.83. Though the planetary C/O is technically inconsistent with the stellar C/O to 1σ (0.19; Table 4), the comparison is not valid, as the planetary C/O prior had a computational lower limit of 0.20 and the posterior has significant likelihood at that limit. This “piling” at the prior boundary implies that the planetary C/O is consistent with the stellar C/O. The median isothermal limb temperature ($T_{\text{limb}} = 1710^{+100}_{-80}$ K) is close to the equilibrium temperature of the planet ($T_{\text{eq}} = 1960$ K), which implies an efficient heat recirculation. These parameters lead to a high mean molecular weight ($\mu \sim 5.5$ AMU) atmosphere with a scale height of about 320 km.

The retrieved results are consistent with a clear atmosphere. Though cloud-top pressure and cloud fraction are unconstrained, their joint marginalized posterior is constrained. A uniform gray cloud is only allowed deeper than ~ 10 Pa (0.1 mBar), and clouds above that pressure are only possible if they cover less than about 40% of the limb. Hazes are dispreferred by model selection, and the median scattering opacity was $50\times$ weaker than Rayleigh scattering in the model that allowed parametric scattering.

The retrieved relative shift between the STIS G430L and G750L instruments is 120 ppm, due in part to the model attempting to align their overlapping regions. A downshift for the WFC3 data is preferred (WFC3 offset = -132 ± 50 ppm). The stellar radius, the planetary mass, and the planetary radius

are consistent with the prior values. The planetary mass and stellar radius are, as expected, dominated by their priors. The planetary radius ($R_p = 1.59 \pm 0.06$) is at the reference pressure of 1 Bar, and when calculated at the planet’s photosphere, it is consistent with the planetary radius derived based on stellar parameters from TIC-8.

7.2. Evidence of Water and Optical-wavelength Absorbers

While the spectral features in STIS, WFC3, and Spitzer are attributed by PLATON to TiO, H₂O, and CO₂, respectively, the retrieval only robustly detects H₂O—the H₂O abundance is constrained by observations, while the abundances of other species are primarily constrained by the assumption of chemical equilibrium. We note that, while CO is more abundant than CO₂, CO₂ has a much larger cross section at $4.5 \mu\text{m}$, such that even with a smaller abundance, its opacity dominates over that of CO at the temperatures and C/O ratios inferred by the retrieval.

We determine whether a species is detected by finding the odds ratio between the best model with and without opacity from a particular species. This breaks the assumption of chemical equilibrium, so it is not strictly correct, but it is a useful heuristic nonetheless. A species is considered detected only when the odds ratio significantly favors the model with the species’ opacity. Table 12 shows the odds ratios—and their more familiar frequentist analog, the detection significances (Benneke & Seager 2013)—for several relevant spectroscopically active species.

The odds ratio in favor of H₂O is ~ 46630 , indicating that the model with water is $46630\times$ more likely than the model without water opacity. This is equivalent to a 5.0σ detection in frequentist terms. The odds ratio in favor of CO₂ is 3.3, which is barely enough evidence to claim a weak detection. PLATON finds no evidence of Na, and CO is dispreferred. The odds ratios for TiO and VO are 2.1 and 2.3, respectively, and these are not favored enough to claim detections (less than 2σ). However, TiO and VO are only seen as nondetections because they have similar cross sections. When TiO opacity is ignored,

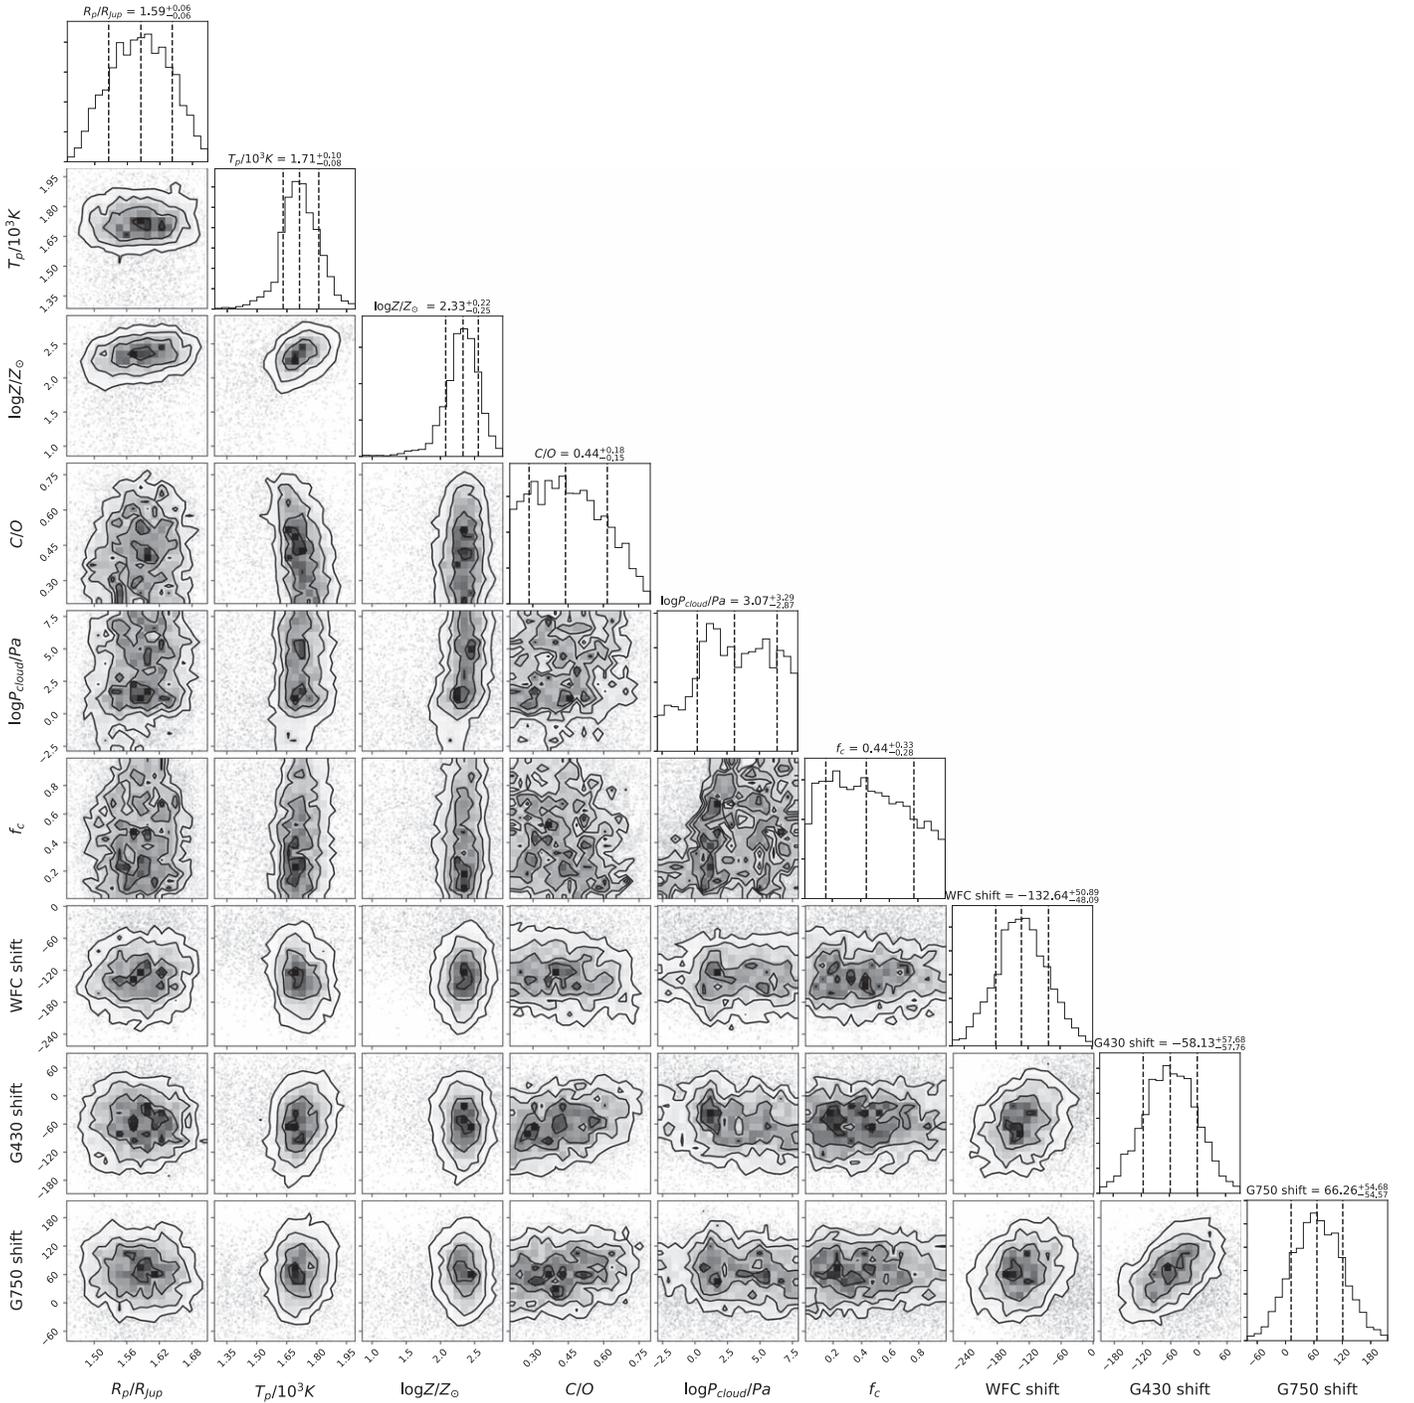


Figure 11. Corner plot for the best PLATON model (fiducial with partial clouds and Gaussian offsets). Here, R_p and M_p are prior-dominated and are excluded for clarity. Offsets are given in parts per million; for example, the median WFC3 offset indicates the retrieval favors shifting the WFC3 depths down by ~ 132 ppm.

the retrieval can compensate because VO opacity is able to describe the STIS feature just as well as TiO. If we ignore both VO and TiO, then the model cannot describe the STIS data as well, and so the odds ratio in favor of TiO/VO is 9.4 (2.7σ). Therefore, we find suggestive evidence of metallic oxide opacity, but we are unable to discern whether it is due to TiO or VO. Based on the assumption of chemical equilibrium at the retrieved temperatures, PLATON attributes the STIS feature to TiO because it is more abundant and opaque in the spectroscopically active region for a solar Ti/V ratio.

8. AURA Retrieval Analysis and Results

We perform a second, complementary atmospheric retrieval analysis: a series of free-chemistry retrievals on HAT-P-41b using AURA (5.2) to constrain the atmospheric properties at the day-night terminator of the planet while allowing for deviations from chemical equilibrium. First, we consider the presence of different chemical species in the atmosphere of HAT-P-41b using its full broadband spectrum. Next, we consider the presence of possible transit depth offsets between data sets and their possible impact in the derived chemical abundances and associated metallicities.

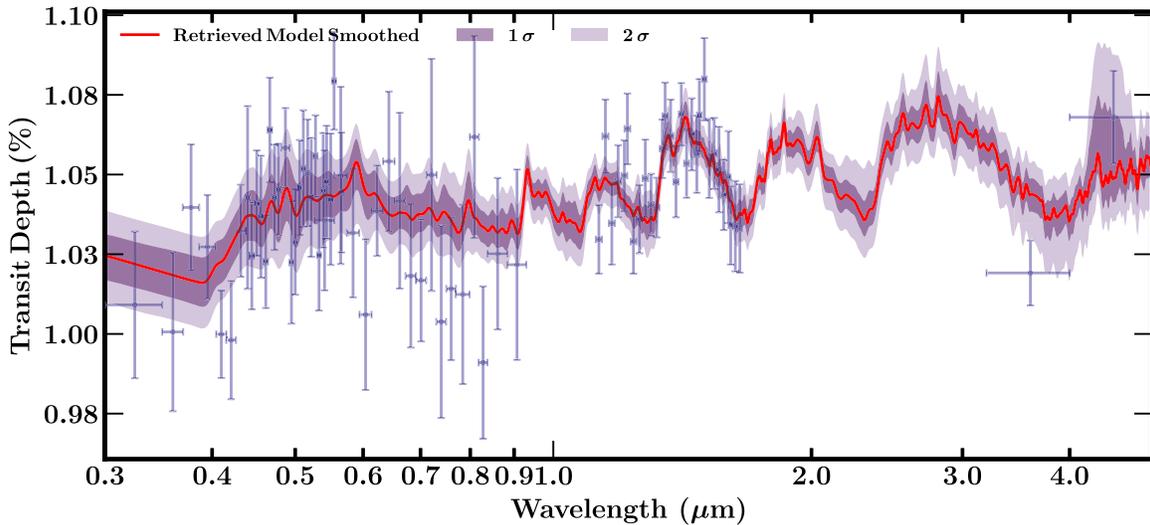


Figure 12. Retrieved spectrum of HAT-P-41b using STIS, WFC3, and Spitzer data. Observations are shown using blue markers. Retrieved median spectrum is shown in red, while the 1σ and 2σ regions are shown using shaded purple areas.

(The data used to create this figure are available.)

Table 12
PLATON Species Detection Evidences

Species	\mathcal{O}^a	Detection Significance ^b
H ₂ O	46630	5.0σ
TiO	2.1	1.9σ
VO	2.3	1.9σ
TiO/VO	9.4	2.7σ
Na	1.1	1.2σ
CO ₂	3.3	2.1σ
CO	0.4	N/A

Notes.

^a Odds ratio between model and the preferred PLATON model (fiducial model with partial clouds and Gaussian-prior instrumental shifts included).

^b Benneke & Seager (2013).

8.1. Evidence of Water and Optical-wavelength Absorbers

We perform a full retrieval on the broadband spectrum of HAT-P-41b and present the observations and retrieved median spectrum in Figure 12. The full retrieval provides constraints on the presence of H₂O, and provides indications for the presence of Na and/or AlO in the optical. The full retrieval finds $\log_{10}(X_{\text{H}_2\text{O}}) = -1.65^{+0.39}_{-0.55}$, $\log_{10}(X_{\text{Na}}) = -3.09^{+1.03}_{-1.83}$, and $\log_{10}(X_{\text{AlO}}) = -6.44^{+0.66}_{-0.91}$. While the retrieval with PLATON prefers TiO/VO to explain the STIS observations, the retrieval with AURA does not. Instead, the latter prefers a combination of Na and AlO. The retrieved TiO abundance is low and unconstrained ($\log_{10}(X_{\text{TiO}}) = -9.58^{+1.37}_{-1.50}$). Neither the CO nor CO₂ abundances are constrained by the retrieval. While the cloud/haze parameters are not tightly constrained, our retrieval indicates a coverage fraction of $\bar{\phi} = 0.25^{+0.26}_{-0.16}$ consistent with a mostly clear atmosphere. The temperature profile of the atmosphere is mostly unconstrained. We infer the temperature near the photosphere, at 100 mbar, to be $T = 1345^{+349}_{-206}$ K. The posterior distributions for the relevant parameters are shown in Figure 13.

We utilize this full retrieval as a reference model to perform a Bayesian analysis and assess the impact of not considering some of these parameters in the models. This change in model evidence is then converted to its more familiar frequentist counterpart, a detection significance (DS), following Benneke & Seager (2013). Table 13 shows the different models considered, their model evidence, DS, and $\bar{\chi}^2$. We find a robust detection of H₂O at a 4.89σ confidence. There is suggestive evidence of Na and/or AlO with confidence levels of 2.09σ and 2.58σ , respectively. The removal of TiO from the models results in an increase in the model evidence, indicating that this molecule is disfavored to be present in our models. VO is similarly undetected. However, removing opacity from the three primary metal oxides (TiO, VO, and AlO), finds a moderate-to-strong “detection,” with 3.59σ confidence. This is similar to PLATON, which did not find evidence of TiO or VO individually, but found weak-to-moderate evidence of their combined presence (Section 7.2). This can be interpreted as follows: AURA is confident (to 3.6σ) that the sharp dip in the blue STIS data ($0.4\text{--}0.5\ \mu\text{m}$) is a real molecular feature due to a metallic oxide. The retrieval finds that the most likely candidate for the metallic oxide is AlO, as shown by its 2.6σ preference, whereas TiO and VO are individually dispreferred.

We assess the retrieved H₂O abundance relative to expectations from thermochemical equilibrium for solar elemental compositions (Asplund et al. 2009). Assuming a solar composition and 50% of the available oxygen in H₂O, the retrieved H₂O abundance corresponds to a log metallicity ([O/H]) of $\log_{10} Z/Z_{\odot} = 1.72^{+0.39}_{-0.55}$ (metallicity of $53^{+82}_{-38} \times Z_{\odot}$). We also compare the retrieved H₂O abundance to the stellar metallicity of the host star ([O/H] = 0.37, Table 4) and obtain a value of $\log_{10} Z/Z_{\text{star}} = 1.35^{+0.39}_{-0.55}$ (metallicity $23^{+33}_{-17} \times Z_{\text{star}}$).

We consider the possibility of fitting the data using a simpler model consisting mainly of the parameters that are reasonably constrained by the full model. The simpler model considers the chemical abundances of H₂O, Na, CO, AlO, an isothermal P – T profile, and a clear atmosphere. The retrieved median fit and confidence contours are shown in Figure 14. The simplified model retrieves values consistent with the full model. The

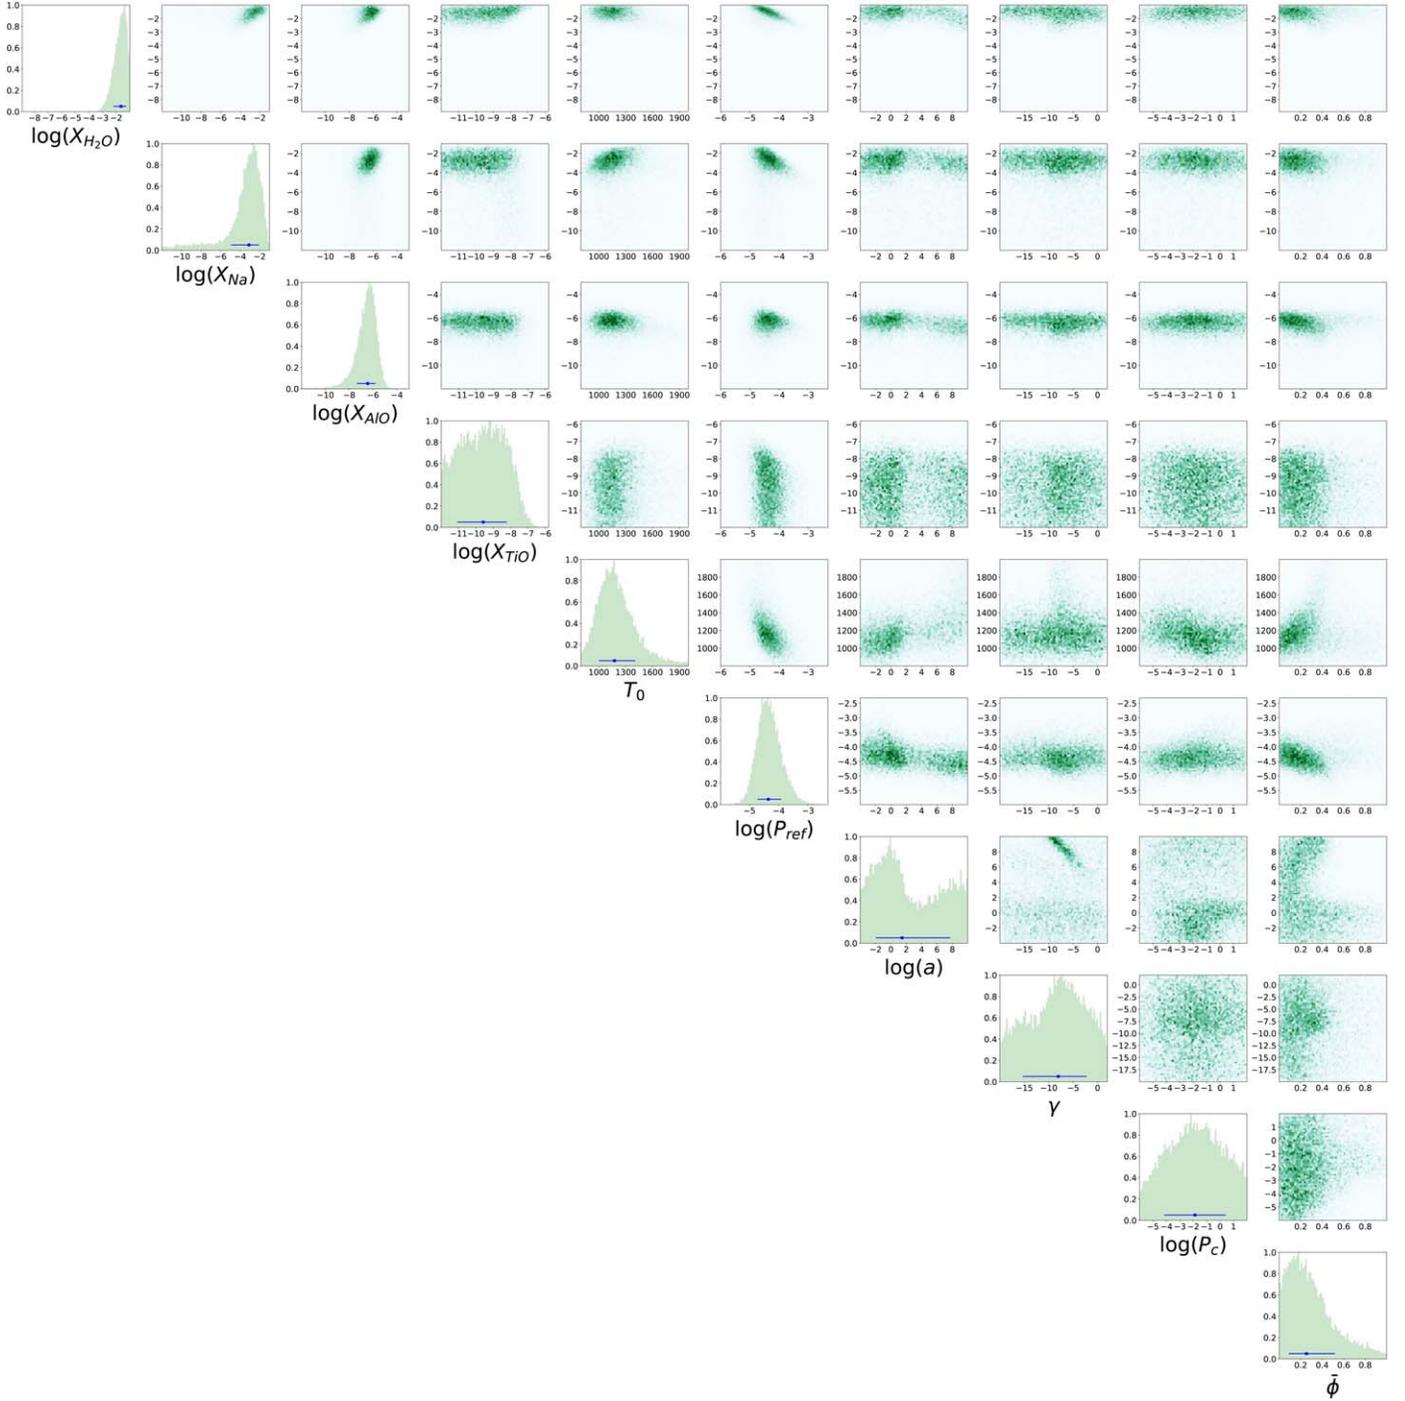


Figure 13. Posterior distributions of the relevant parameters for the full retrieval (Model 1 in Table 13) using STIS, WFC3, and Spitzer data. The abundances of H₂O, Na, and AIO are constrained, while the cloud and haze parameters are not constrained. Parameter T_0 , the temperature at the top of the atmosphere (10^{-6} bar), is shown as a subset of the P - T parameters used in the model.

retrieved values are $\log_{10}(X_{\text{H}_2\text{O}}) = -1.65^{+0.40}_{-0.63}$, $\log_{10}(X_{\text{Na}}) = -2.60^{+0.94}_{-1.10}$, $\log_{10}(X_{\text{AIO}}) = -5.81^{+0.51}_{-0.66}$, and $\log_{10} Z/Z_{\odot} = 1.72^{+0.40}_{-0.63}$. The retrieved isothermal temperature is $T = 1120^{+170}_{-140}$ and consistent with the inferred temperature at 100 mbar from the full retrieval. The posterior distribution for the retrieved parameters is shown in Figure 15.

We use these retrieved parameters to generate a set of forward models to assess the spectroscopic contribution from each chemical species. Figure 14 shows that the WFC3

observations are better explained by the H₂O absorption feature at $\sim 1.4 \mu\text{m}$ driving its strong detection in the spectrum of HAT-P-41b. On the other hand, a series of chemical species in the optical can provide some degree of fit to the STIS observations. In the optical, between ~ 0.5 and $0.7 \mu\text{m}$, the broadened wings of Na along with its absorption peak provide a fit to observations. AIO provides some fit to the substructure present in the STIS observations, particularly the increased transit depth between 0.4 and $0.5 \mu\text{m}$. Finally, the abundance of CO is not constrained, and its contribution to the spectrum is

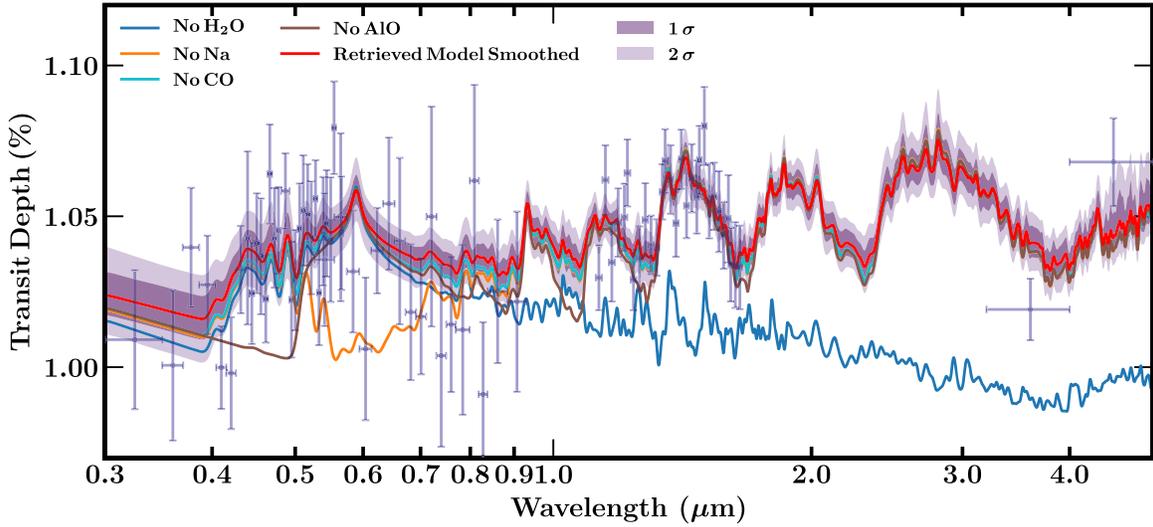


Figure 14. Retrieval of HAT-P-41b using a simplified model compared with the fiducial parameter set (see Section 8.1). Observations are shown using blue markers. Retrieved median spectrum is shown in red, while the 1σ and 2σ regions are shown using shaded purple areas. Forward models using the retrieved median parameters show the contributions to the spectra due to individual chemical species. Forward models shown exclude absorption due to H_2O (blue), Na (orange), CO (cyan), and AIO (brown).

Table 13
Retrieved Models

#	Model	$\log_{10}(X_{\text{H}_2\text{O}})$	$\log_{10}(X_{\text{Na}})$	$\log_{10}(X_{\text{AIO}})$	$\log_{10} Z/Z_{\odot}$	STIS _{Shift}	WFC3 _{Shift}	$\ln(Z)$	$\bar{\chi}^2$	DS
1	Full Model	$-1.65^{+0.39}_{-0.55}$	$-3.09^{+1.03}_{-1.83}$	$-6.44^{+0.66}_{-0.91}$	$1.72^{+0.39}_{-0.55}$	N/A	N/A	559.1	0.93	Ref.
2	No H_2O	N/A	$-2.41^{+0.99}_{-2.99}$	$-5.71^{+0.99}_{-1.39}$	N/A	N/A	N/A	548.9	1.37	4.89
3	No Na	$-1.62^{+0.42}_{-0.67}$	N/A	$-6.90^{+0.84}_{-1.05}$	N/A	N/A	N/A	558.0	0.95	2.09
4	No AIO	$-1.49^{+0.35}_{-0.70}$	$-4.32^{+1.88}_{-4.31}$	N/A	N/A	N/A	N/A	557.0	1.03	2.58
5	No TiO	$-1.70^{+0.41}_{-0.56}$	$-2.97^{+0.95}_{-1.25}$	$-6.39^{+0.66}_{-0.88}$	N/A	N/A	N/A	559.7	0.92	N/A
6	No Metal Oxides	$-1.52^{+0.38}_{-0.91}$	$-3.59^{+1.28}_{-1.47}$	N/A	N/A	N/A	N/A	554.2	1.21	3.59
7	Simpler Model	$-1.65^{+0.40}_{-0.63}$	$-2.60^{+0.94}_{-1.10}$	$-5.81^{+0.51}_{-0.66}$	$1.72^{+0.40}_{-0.63}$	N/A	N/A	560.0	0.89	N/A
8	Gaussian Shifts G430L, G750L, and WFC3	$-1.91^{+0.53}_{-0.68}$	$-2.38^{+0.81}_{-1.33}$	$-6.64^{+0.70}_{-0.96}$	$1.46^{+0.53}_{-0.68}$	G430L -51^{+60}_{-62} G750L 79^{+59}_{-56}	-91^{+59}_{-56}	562.0	0.90	N/A
9	Uniform Shifts G430L, G750L, and WFC3	$-2.96^{+0.98}_{-0.88}$	$-2.43^{+0.84}_{-1.34}$	$-7.05^{+0.75}_{-0.94}$	$0.40^{+0.98}_{-0.88}$	G430L 1^{+144}_{-156} G750L 175^{+150}_{-160}	-189^{+90}_{-94}	561.8	0.88	N/A
10	Uniform Shifts STIS and WFC3	$-3.34^{+1.00}_{-0.86}$	$-3.43^{+1.35}_{-2.19}$	$-6.98^{+0.77}_{-0.78}$	$0.03^{+1.00}_{-0.86}$	89^{+166}_{-156}	-203^{+97}_{-97}	560.7	0.89	N/A

Notes. The metallicity is approximated from water abundance (see Section 8.1 for details). N/A means that the parameter was not considered in the retrieval or that it is impossible to estimate the detection significance (DS) because the model has a larger amount of evidence than the reference model (model 1). Shifts are given in ppm.

minimal. CO is responsible for small changes in the optical and infrared that are well within the error bar of the observations.

Our retrieval analysis of the broadband transmission spectrum of HAT-P-41b provides excellent fits to the data; using our fiducial model (Model 1), we obtain a best-fit $\bar{\chi}^2$ of 0.93 and $\ln(Z) = 559.1$. We note that we do not require additional continuum opacity sources (e.g., H^-) in our models in order to explain the data, as recently claimed by Lewis et al. (2020).

8.2. Possible Offsets in the Data

Last, we consider the presence of offsets in the data and their effect on the retrieved atmospheric properties. We consider the three scenarios from Section 5.1.2. We note that these retrieved offsets are relative to the atmospheric model and that the Spitzer observations remain unchanged in all scenarios. We consider both Gaussian and uniform priors, as seen in Table 8.

We present the results of considering the presence of three offsets with Gaussian priors informed by the analysis of the white light transit curves (Model 8; Scenario 1 from Section 5.1.2). These priors are shown in Table 8. The retrieved shifts are -52^{+61}_{-63} ppm for G430L, 80^{+59}_{-56} ppm for G750L, and -91^{+48}_{-50} ppm for WFC3. Similar to PLATON, the retrieval generally prefers to increase the G750L depths, primarily motivated by aligning the transit depths in the overlapping wavelength region between G430L and G750L. The retrieval also prefers to decrease WFC3 depths in order to better capture the Spitzer $3.6\ \mu\text{m}$ depth. The retrieved abundances are $\log_{10}(X_{\text{H}_2\text{O}}) = -1.91^{+0.53}_{-0.68}$, $\log_{10}(X_{\text{Na}}) = -2.38^{+0.81}_{-1.33}$, and $\log_{10}(X_{\text{AIO}}) = -6.64^{+0.70}_{-0.96}$. Although the retrieved H_2O abundance corresponds to a lower metallicity estimate, the derived range $\log_{10} Z/Z_{\odot} = 1.46^{+0.53}_{-0.68}$ is consistent with the fiducial model and describes a metal-rich atmosphere. The median metallicity is superstellar

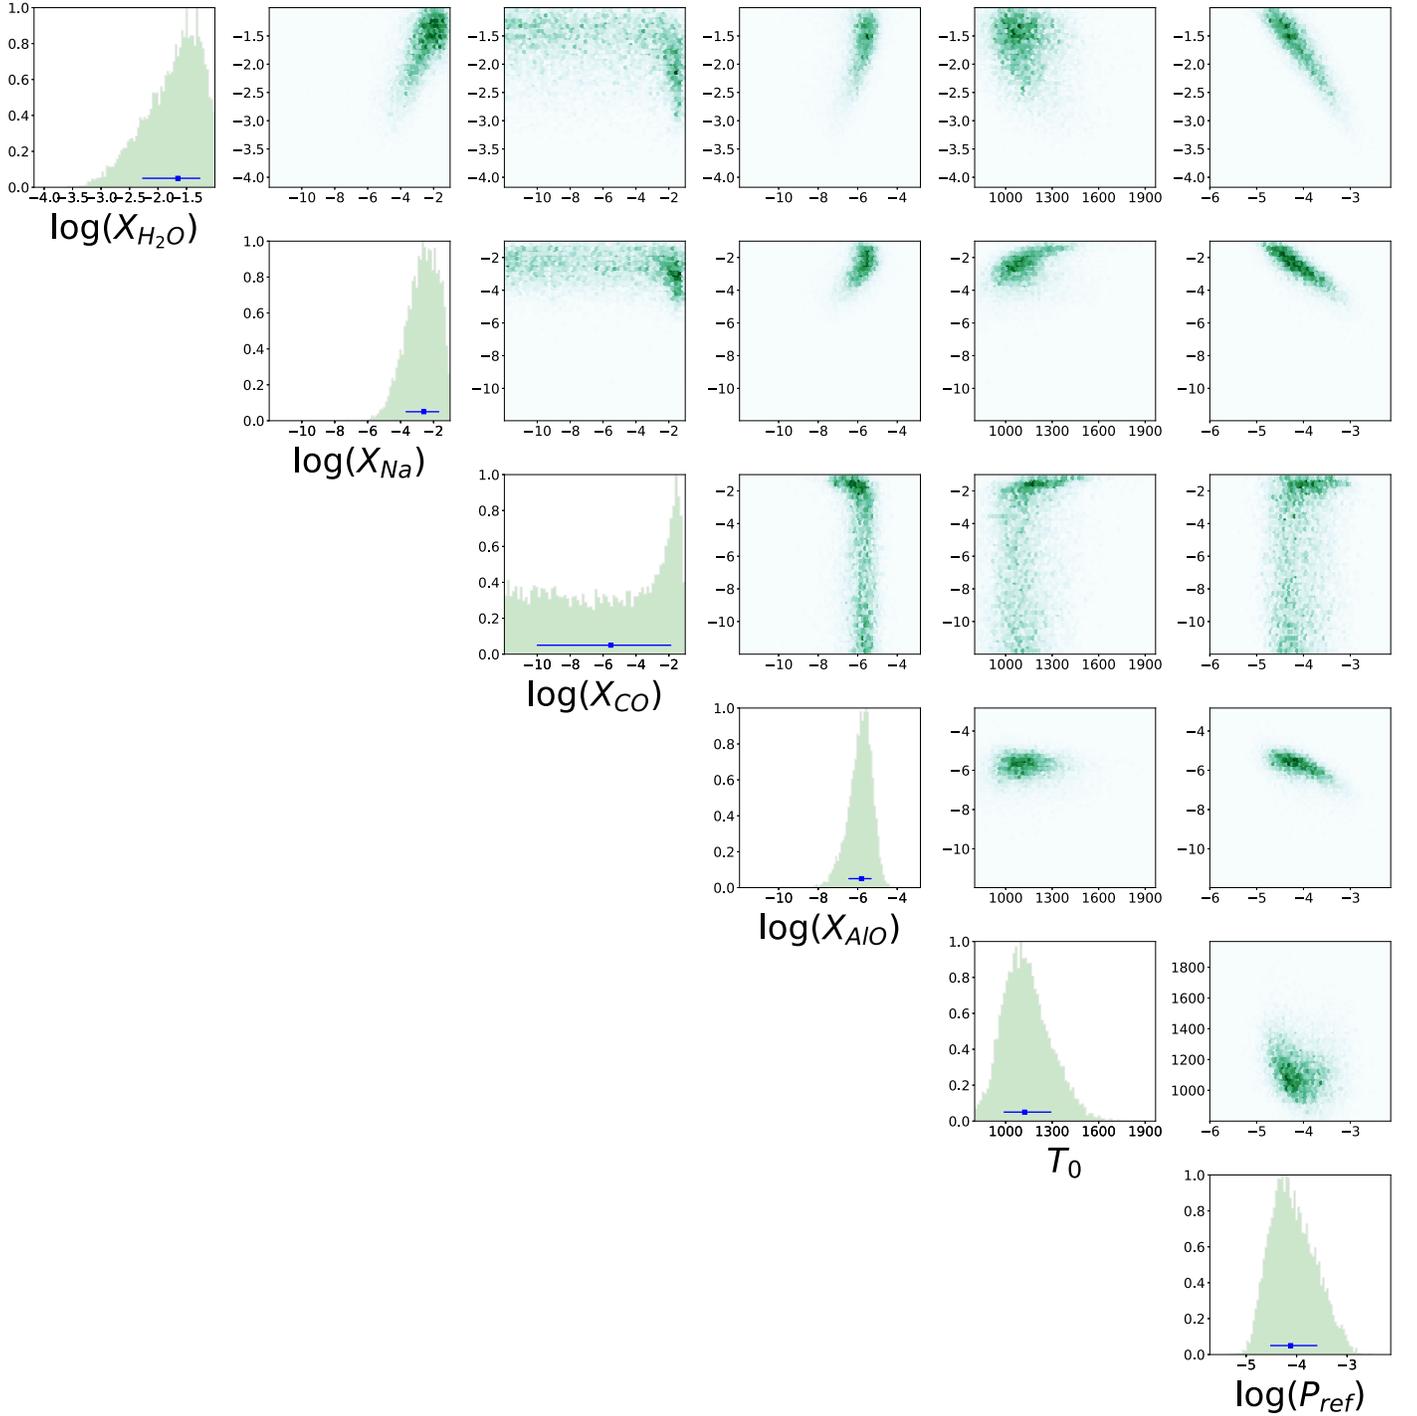


Figure 15. Full posterior distributions for the simpler model, Model 7 in Table 13.

($\log_{10} Z/Z_{\text{star}} = 1.09^{+0.53}_{-0.68}$), though it is consistent with stellar metallicity to within 2σ .

Second, we present the results for the case with three uniform shifts between HST-STIS G430L, HST-STIS G750L, and HST-WFC3 observations (Scenario 2 from Section 5.1.2). The retrieved G430L shift is consistent with 0 (1^{+144}_{-156} ppm), but the retrieval prefers a large positive G750L offset (176^{+151}_{-160} ppm) and a large negative WFC3 offset (-189^{+91}_{-94} ppm). In addition to aligning the overlapping G430L-G750L region, it is possible that this large G750L shift is due to the model forcing the data to match features it finds easier to explain. This

uncertainty is the danger in using uniform prior offsets, especially in an already-flexible free-chemistry retrieval. The retrieved abundances are shown in Table 13 as Model 9, and are $\log_{10}(X_{\text{H}_2\text{O}}) = -2.96^{+0.98}_{-0.88}$, $\log_{10}(X_{\text{Na}}) = -2.43^{+0.84}_{-1.34}$, and $\log_{10}(X_{\text{AIO}}) = -7.05^{+0.75}_{-0.94}$. While the retrieved abundances for these three species are consistent within 1σ with the full unshifted model, the retrieved H_2O abundance corresponds to a lower metallicity estimate consistent with solar, subsolar, and substellar values $\log_{10} Z/Z_{\odot} = 0.40^{+0.98}_{-0.88}$.

Last, we present the results accounting for offsets in the STIS and WFC3 observations using a uniform prior, while

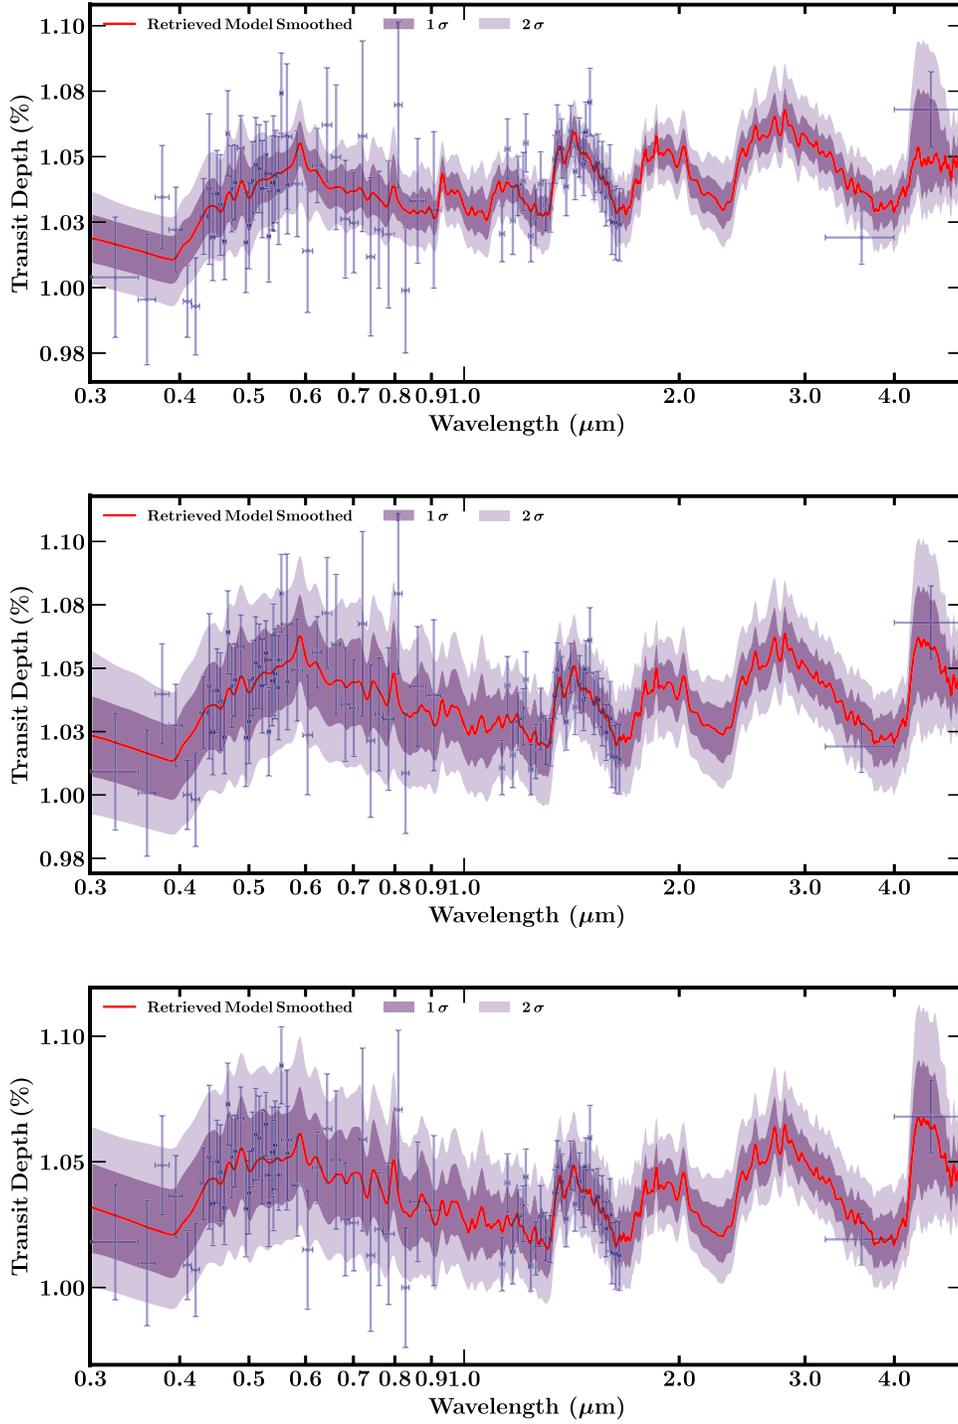


Figure 16. Retrieved spectrum of HAT-P-41b allowing for offsets in the STIS and WFC3 data sets. Observations are shown using blue markers and are shifted according to the models’ retrieved median shifts. Retrieved median spectrum is shown in red, while the 1σ and 2σ regions are shown using shaded purple areas. Top: three shifts with Gaussian priors (Model 8) and retrieved median offsets of ~ -50 ppm for STIS G430L, ~ 80 ppm for G750L, and ~ -90 for WFC3. Middle: three shifts with uniform priors (Model 9) and retrieved median offsets of ~ 0 ppm for STIS G430L, ~ 180 ppm for G750L, and ~ -190 for WFC3. Bottom: two shifts with uniform priors (Model 10) and retrieved median offsets of ~ 90 ppm for STIS and ~ -200 ppm for WFC3.

keeping the Spitzer observations unshifted (Scenario 3 from Section 5.1.2). The retrieval results in a shift in the STIS data of 90^{+167}_{-157} ppm and a shift in the WFC3 data of -204^{+97}_{-98} ppm. While the retrieved value for the STIS observations is consistent with no shift, the WFC3 observations preferentially retrieve a negative offset. The derived abundances, shown as Model 10 in Table 13, are $\log_{10}(X_{\text{H}_2\text{O}}) = -3.34^{+1.00}_{-0.86}$,

$\log_{10}(X_{\text{Na}}) = -3.43^{+1.35}_{-2.19}$, $\log_{10}(X_{\text{AlO}}) = -6.98^{+0.77}_{-0.78}$. The H_2O abundance, like Model 9, corresponds to a metallicity consistent with solar and subsolar values: $\log_{10} Z/Z_{\odot} = 0.03^{+1.00}_{-0.86}$.

Figure 16 shows the retrieved median models and confidence contours along with their respectively shifted observations for the cases described in this section (Models 8, 9, and 10).

The models considering instrumental shifts are all preferred over the fiducial model at above the 2σ level. The model with Gaussian priors has a preference at the 2.9σ level, followed by the model with three uniform shifts at a 2.8σ level. The model with two uniform shifts is preferred over the fiducial model at 2.3σ . We note that, while both models with three offsets are similarly preferred over our fiducial model, the associated metallicity ranges are different. The model with three uniform shifts retrieves an H_2O abundance corresponding to a metallicity estimate consistent with substellar and stellar values. On the other hand, the model with Gaussian priors retrieves an associated metallicity range that is mostly super-stellar and in agreement with the fiducial model. These results highlight the sensitivity of the inferred metallicity ranges to possible large offsets between instruments. Model comparisons suggest a preference for the models considering offsets, though it is inconclusive between these models. We favor the more physically plausible Gaussian prior model (i.e., Model 8) as the reference for our discussion (Section 9).

9. Discussion

9.1. Comparison between Retrieval Methods

9.1.1. Results Comparison

In this section, we compare the results from the preferred PLATON and AURA models. These include the fiducial model with partial clouds and Gaussian instrumental offsets for PLATON (Section 6.2.3) and the Gaussian instrumental offset model for AURA (Model 8; Section 8.2).

The similarities reveal the most robust conclusions of our analysis, since they are retrieved despite the many different assumptions that went into each method. Notably, both retrievals robustly find a metal-rich atmosphere with metallicity (defined as O/H) inconsistent with the solar metallicity at $>2\sigma$. Both methods find a decisive ($>4.8\sigma$) water vapor detection, and at least a moderate detection ($>2.7\sigma$) of a non-haze gas absorption feature in the optical. Further, both PLATON and AURA retrievals are consistent with a mostly clear atmosphere, with neither finding strong evidence of haze or uniform, high-altitude gray clouds.

Though the atmospheric properties derived from PLATON and AURA are similar, there are noteworthy differences. AURA infers a cooler limb temperature at 100 mbar (1320^{+270}_{-200} K compared to 1710^{+100}_{-80} K for PLATON) as well as a lower metallicity of $\log_{10} Z/Z_{\odot} = 1.46^{+0.53}_{-0.68}$ compared to $\log_{10} Z/Z_{\odot} = 2.33^{+0.23}_{-0.25}$ for PLATON, a difference of 1.3σ . This translates to $29^{+69}_{-23} \times Z_{\odot}$ for AURA and $214^{+149}_{-88} \times Z_{\odot}$. The optical absorber also differs: AURA determines the best description of the STIS feature to be absorption from sodium and AIO, whereas PLATON prefers some combination of TiO and VO absorption. Finally, AURA makes no claim on the C/O ratio, as it is a free retrieval framework and no C-bearing species are detected or meaningfully constrained. On the other hand, the chemical equilibrium assumption allows PLATON to find a 3σ upper limit on the C/O ratio of $\text{C}/\text{O} < 0.83$.

To further contextualize the results, we added the functionality to retrieve the abundance profiles of relevant molecules in PLATON. We show abundance profiles for six spectroscopically relevant species from AURA, which are also included in PLATON— H_2O , CO, CO_2 , Na, TiO, and VO—in Figure 17. We emphasize that enforcing chemical equilibrium

narrows the abundance constraints, and we are not reporting these abundances. Instead, they should be interpreted as the expected abundance profiles under the conditions of stable chemical equilibrium for the reported temperature, metallicity, and C/O ratio. As an example, we find no observational constraint on CO, but its abundance is well-defined under chemical equilibrium for the temperatures and metallicities that we do observationally constrain via the water feature. Still, these profiles provide a useful baseline for comparison to free-chemistry retrieval abundances.

The abundance profiles for the optical-wavelength absorbers reflect the disagreement on the primary gas absorber: AURA prefers Na, and so it retrieves $\sim 10\times$ more Na than PLATON and significantly less TiO and VO. Note that the decreasing TiO and VO abundance with increasing pressure for PLATON is due to those molecules condensing out of the atmosphere. PLATON's inferred water abundance is typically a few times greater than AURA's, reflecting the difference in inferred metallicities. CO_2 and CO are unconstrained by AURA, while PLATON finds a high abundance of CO_2 is consistent with the Spitzer observations. This difference is expected, given that PLATON finds weak evidence of CO_2 while AURA found none. Interestingly, this may relate to AURA's only chemical constraint, which is that CO_2 must be less abundant than CO and H_2O , due to the inferred temperatures (Section 5.2).

In total, AURA finds a cooler atmosphere with less oxygen but a large sodium enrichment to explain the optical absorption, while PLATON finds a hotter atmosphere with more oxygen and TiO/VO absorption in the optical.

9.1.2. Impact of Retrieval Model Assumptions

The differences between a free-chemistry retrieval (AURA) and one constrained by chemical equilibrium (PLATON) are the natural result of the different assumptions made by each method. We therefore consider the PLATON and AURA retrievals to be two orthogonal analyses. We examine the impact of the differences by first explicitly listing the notable assumptions in each method, and then by providing the rationale for which assumptions are driving the differences.

The relevant methodological differences for PLATON as compared to AURA are: (1) the assumption of chemical equilibrium, (2) fixing the elemental ratios between all metals other than carbon to their solar values, (3) assuming an isothermal profile for the atmosphere, (4) not including opacity from AIO, and (5) not including the Allard et al. (2019) H_2 -broadened Na line profile.

We find that the Allard et al. (2019) H_2 -broadened Na line profile is the key driver in the differences between the retrievals, and flexible element abundances and chemical equilibrium also play roles. AURA is the more flexible retrieval, so we first describe its solution before addressing why PLATON differs.

AURA's lower-temperature solution is preferred for being able to explain the H_2O feature in the WFC3 spectrum, while also explaining the STIS data with H_2 -broadened Na absorption and capturing the Spitzer data. AURA is able to provide a fit to the STIS data by independently increasing the Na abundance and by also invoking AIO at relatively low temperatures. At this lower temperature ($T \sim 1300$ K), the amount of oxygen necessary for the water abundance and scale height to explain the observed water feature is about $29 \times Z_{\odot}$, with a mean

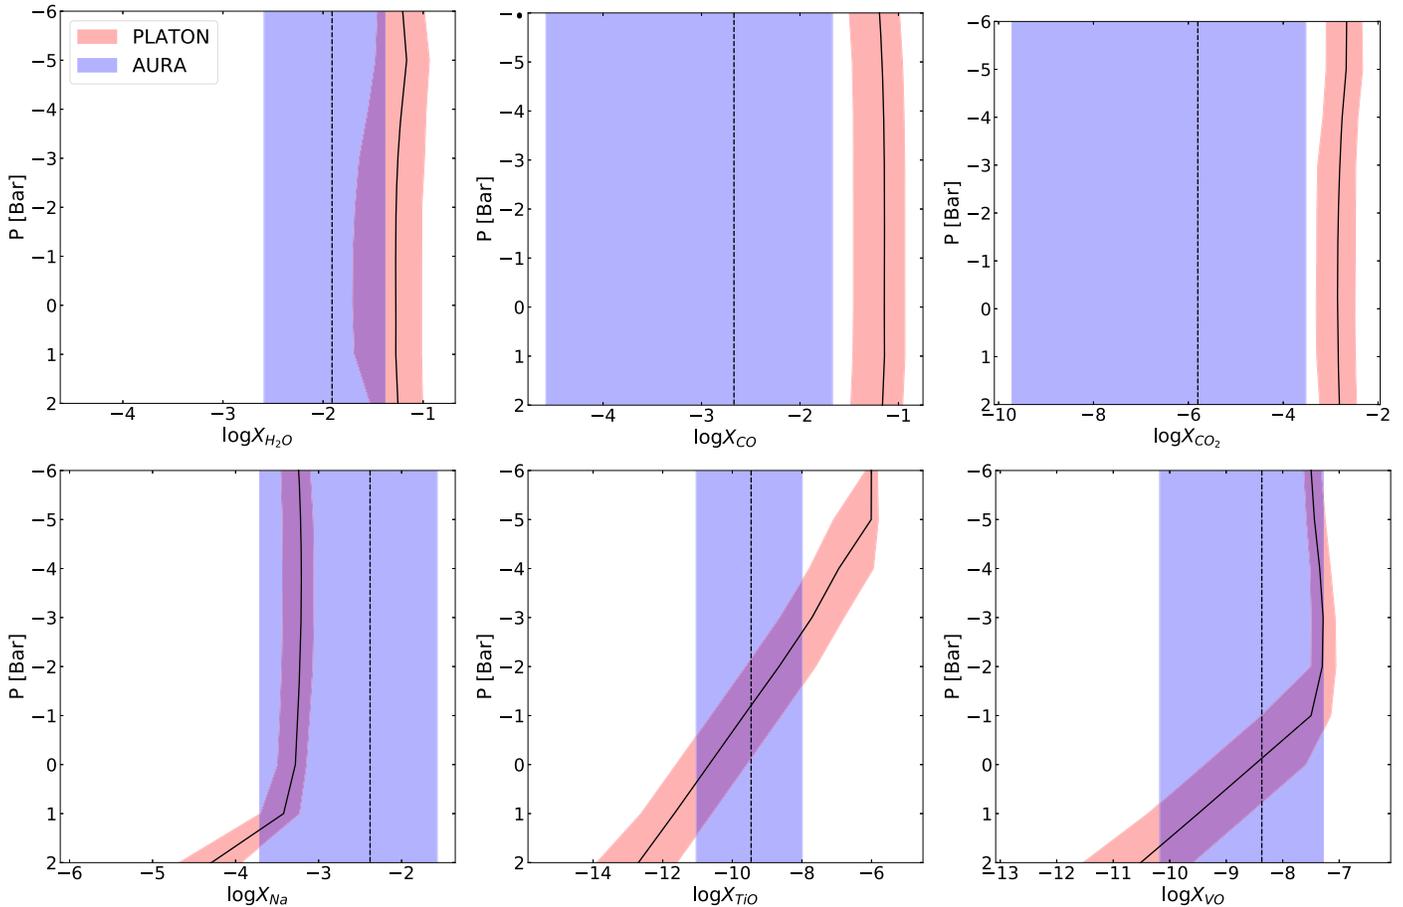


Figure 17. Abundance profiles for PLATON (red) and AURA (blue) for four relevant gaseous species. PLATON abundance profile distributions are derived by sampling the posterior 200 times, calculating the abundance profiles for each species for each sample, and finding the median value (solid black line) with 1σ uncertainties at each pressure layer. AURA assumes abundances to be constant with pressure. The median retrieved value (dashed black line) and 1σ uncertainty range are shown.

molecular weight of about 2.7 AMU and a scale height of about 440 km.

Since PLATON has not yet incorporated the H_2 -broadened Na line profile, the low-temperature solution is a relatively poor fit to the STIS data. Instead, TiO/VO are needed to explain the STIS absorption feature, and these are only abundant enough in chemical equilibrium (with fixed metal ratios) at around 1650 K. At this higher temperature, a higher mean molecular weight is required for the same scale height, which must be small enough to explain the molecular feature sizes as well as the dominance of TiO/VO absorption over Rayleigh scattering. The atmospheric metallicity necessary to achieve the higher mean molecular weight is the much higher $\sim 200 \times Z_{\odot}$. Therefore, the differences make sense in light of the stricter assumptions.

To provide more support to this idea, we compare our results with those of a third retrieval method, ATMO (Amundsen et al. 2014; Tremblin et al. 2015, 2016, 2017; Sing et al. 2016), which acts as a middle ground between PLATON and AURA. ATMO’s spectral retrievals can further help us to gain insight into the effect of retrieval assumptions, as it includes the Allard et al. (2007) pressure-broadened sodium line but also has the added flexibility of performing a free-element equilibrium-chemistry retrieval. With this assumption for the chemistry, the elemental abundances for each model are freely fit and calculated in equilibrium on the fly. Four elements were

selected to vary independently, as they are major species that are also likely to be sensitive to spectral features in the data, while the rest were varied by a trace metallicity parameter ($[Z_{\text{trace}}/Z_{\odot}]$). By separately varying the carbon, oxygen, sodium, and vanadium elemental abundances ($[C/C_{\odot}]$, $[O/O_{\odot}]$, $[Na/Na_{\odot}]$, $[V/V_{\odot}]$), we allow for nonsolar compositions—but with chemical equilibrium imposed such that each model fit has a chemically plausible mix of molecules, given the retrieved temperatures, pressures, and underlying elemental abundances.

The resulting retrieved atmospheric parameters describe an atmosphere that is most consistent with the one described by AURA. ATMO prefers a temperature of 1190_{-120}^{+170} K and a metallicity (as defined by the oxygen abundance) of $\log_{10} O/O_{\odot} = \log_{10} Z/Z_{\odot} = 1.53_{-0.67}^{+0.55}$, in excellent agreement with AURA’s values, and consistent with PLATON’s metallicity to 1.3σ , though the retrieved temperatures differ significantly. Like AURA, ATMO finds an enhanced sodium abundance, though uncertainties are large ($\log_{10} Na/Na_{\odot} = 1.40_{-1.80}^{+0.75}$). This supports the idea that the inclusion of H_2 -broadened sodium line profiles and the flexibility of nonsolar metal ratios—and not necessarily the equilibrium chemistry constraint—allow for the low-temperature, lower oxygen abundance solution found by AURA. The metallicities on all three retrievals indicate a metal-rich atmosphere and agree at the $\sim 1.3\sigma$ level.

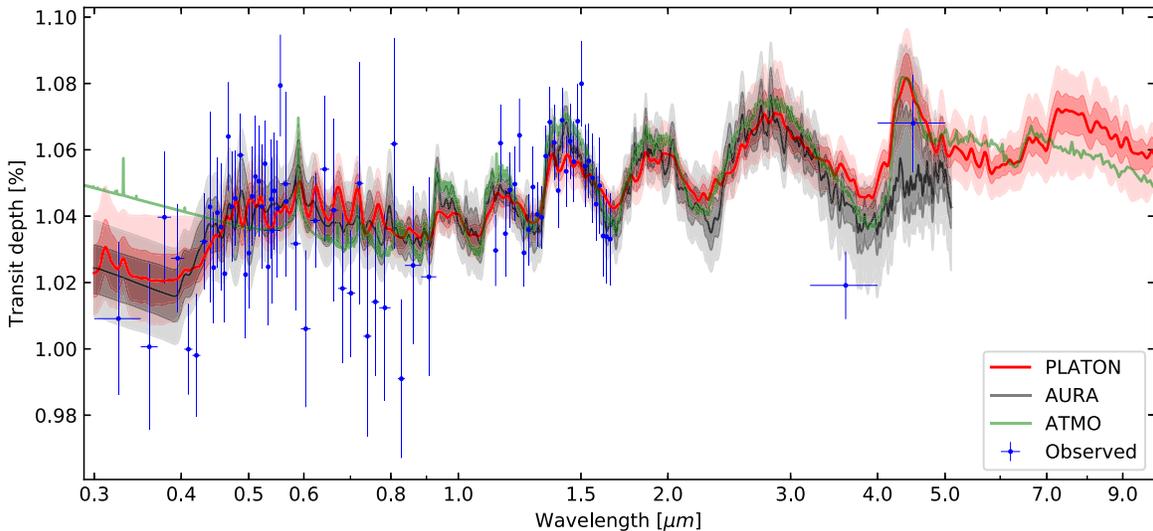


Figure 18. Comparison of the median retrieved model for each retrieval method’s fiducial model. The 1σ and 2σ uncertainty contours are included for PLATON and AURA. PLATON and AURA are smoothed with a Gaussian filter with $\sigma = 15$ for clarity. The chemical equilibrium assumption used by PLATON and ATMO allows for meaningful predictions at unobserved wavelengths, and so those models are shown out to $10\ \mu\text{m}$.

(The data used to create this figure are available.)

Like PLATON (Section 7), ATMO also finds a subsolar C/O ratio ($C/O = 0.17^{+0.53}_{-0.16}$ consistent with stellar ($C/O = 0.19$), though carbon is not well-constrained, so the uncertainties are large. The 3σ upper limit of 0.94 is in good agreement with PLATON’s 0.83 upper limit. However, unlike PLATON or AURA, ATMO finds no evidence of optical absorbers beyond Na; instead, it prefers a haze and Na to explain the STIS optical data.

Figure 18 elucidates the differences in retrievals by showing the median retrieved fiducial model for PLATON (red), AURA (black), and ATMO (green) from 0.3 to $10\ \mu\text{m}$. The 1 and 2σ uncertainty contours are shown for PLATON and AURA, both of which are smoothed with a Gaussian filter with $\sigma = 15$ for clarity. The AURA predictions are only shown up to $5\ \mu\text{m}$ —as a free-chemistry retrieval, AURA retrieving on the 0.3– $5\ \mu\text{m}$ data does not place meaningful constraints on multiple molecules with significant opacity in the 5– $10\ \mu\text{m}$ range. Therefore, a prediction is not warranted.

While there are subtle differences, such as PLATON and ATMO’s preference for CO_2 at $4.5\ \mu\text{m}$ and sodium’s prominence at $0.6\ \mu\text{m}$ in the AURA and ATMO retrievals, the most obvious difference is below $0.5\ \mu\text{m}$, where ATMO prefers a haze instead of a metallic oxide feature. Though ATMO does not include AIO as an opacity source, this difference is likely due to different condensation schemes. PLATON uses GGchem’s prescription (Woitke et al. 2018) such that species condense out when it is energetically favorable. AURA is a free-chemistry retrieval, so there are no restrictions on oxides being in the gaseous phase. ATMO, however, includes rainout chemistry (Goyal et al. 2019), such that if a species condenses at a higher pressure, that then depletes the element above that layer. It is plausible that, although PLATON’s condensation scheme allows TiO/VO to be in the gas phase around 1700 K, ATMO’s scheme does not, making the metallic oxide feature difficult to capture.

In total, we tentatively favor AURA’s derived atmospheric parameters over PLATON’s, for two main reasons. First, the inclusion of the most up-to-date sodium line profiles and AIO

opacity impact the retrieval. Second, constraints from interior modeling (Section 9.2), though not necessarily decisive, are consistent with AURA and in tension with PLATON. Overall, this paints a picture of an atmosphere with a supersolar—but not necessarily superstellar—metallicity, sodium enrichment, possible disequilibrium metallic oxides (e.g., circulated from dayside, dredged up due to vertical mixing), and a planet with a well-mixed interior and a limb temperature lower than the equilibrium temperature.

9.2. Comparison to Interior Modeling Metallicity Constraints

Though they both describe metal-rich atmospheres, the 1σ retrieved atmospheric metallicities ranges from AURA and PLATON are inconsistent ($\log_{10}Z/Z_{\odot} = 0.78$ – 1.99 and $\log_{10}Z/Z_{\odot} = 2.08$ – 2.56 , respectively). Further, it is questionable whether such supersolar metallicities—especially those retrieved by PLATON—are physically reasonable. We check the viability of these values by comparing them to atmospheric metallicity constraints from interior structure models.

Thorngren & Fortney (2019) demonstrated how interior models can constrain atmospheric metallicity. Essentially, this is a three-step process: (1) determine what range of bulk metallicities are necessary for structure models to explain the observed radius, taking into account the planet’s mass, age, heating efficiency, and parameter uncertainties; (2) set the maximum bulk metallicity to be the 3σ upper limit of the derived posterior distribution; and (3) set the maximum atmospheric metallicity to be equal to the maximum bulk metallicity.

The third step assumes that the atmospheric metallicity cannot be greater than the core’s metallicity for significant timescales, due to convection or Rayleigh–Taylor instability. Thorngren & Fortney define metallicity as the ratio of all metals to hydrogen compared to the ratio in the Sun’s photosphere. This is a good proxy for O/H, and so it is a valid comparison to the retrieved atmospheric metallicities. For more details on the derivation, see Thorngren & Fortney (2019).

Using stellar parameters from Hartman et al. (2012; Table 3), the interior structure model fit yields a bulk metal abundance ratio of $Z/Z_{\odot} = 33.7 \pm 9.1$, corresponding to a maximum atmospheric metallicity of $50 \times Z_{\odot}$ (D. Thorngren 2020, private communication). There is no significant uncertainty on this number, as it is the 3σ upper limit of the distribution. This is consistent with the metallicity from the AURA retrieval, but it is in tension with PLATON’s retrieved metallicity— $50 \times Z_{\odot}$ falls outside PLATON’s 1σ range (but within 2σ , as the metallicity distribution is asymmetric and PLATON’s 2σ lower limit is $37 \times Z_{\odot}$). This could indicate that the “true” atmospheric metallicity falls in the lower range of PLATON’s retrieved metallicity, or it could be interpreted as slight evidence in support of AURA over PLATON. Either way, the atmospheric metallicity approaching the bulk metallicity indicates a well-mixed interior. Such vertical mixing could allow for micron-sized particles to stay afloat in the atmosphere, potentially facilitating gaseous metal oxide survival and Mie scattering.

9.3. Implications for Planet Formation

The atmospheric metallicity we retrieve for HAT-P-41b provides important constraints on the formation and migration history of the planet. At the outset, the supersolar metallicity (O/H) of $\sim 30\text{--}200 \times Z_{\odot}$ requires substantial accretion of solids, beyond several Earth masses of H_2O ice, during the planet’s evolutionary history. It is unlikely that such a large amount of volatile accretion is possible at the planet’s current orbit. Therefore, the planet is unlikely to have formed *in situ* (Batygin et al. 2016). Instead, it probably formed far out beyond the H_2O snow line and migrated inward. The formation location and migration path of a giant planet can significantly affect its chemical composition. Beyond the H_2O snow line, the gas in the protoplanetary disk is depleted in oxygen whereas the solids are enriched in oxygen (Öberg et al. 2011). Therefore, planets with high enrichment of oxygen require predominant accretion of H_2O -rich planetesimals while forming and migrating through the protoplanetary disk.

The high metallicity (specifically O/H) of HAT-P-41b, therefore, supports the migration of the planet through the disk via viscous torques (Madhusudhan et al. 2014a). This is in contrast to other hot Jupiters with low O/H abundances, which have been suggested to be caused by insufficient solid accretion, e.g., via disk-free migration (Madhusudhan et al. 2014a) or formation via pebble accretions whereby the oxygen-rich solids are locked in the core (Madhusudhan et al. 2017). The fact that HAT-P-41b’s orbit is moderately misaligned to the host star’s rotation axis is also in tension with the disk migration hypothesis, since spin–orbit misalignments are considered to be evidence of disk-free migration and planet–planet interactions (Winn et al. 2010). In principle, instead of disk migration, super-solar elemental abundances could be caused by accreting gas whose metallicity has been enhanced due to pebble drift (Öberg & Bergin 2016; Booth et al. 2017). However, while pebble drift can cause metal enhancements up to $\sim 10 \times Z_{\odot}$, much larger enhancements (as constrained in the present case) are unlikely to be explained by this process. More importantly, such enhancements due to pebble drift are also expected to cause high C/O ratios (~ 1), which may be at odds with the high H_2O abundance and the low C/O ratio retrieved for the planet.

Overall, the most plausible explanation for the potentially high atmospheric metallicity inferred for HAT-P-41b is

formation outside the H_2O snowline and migration inward while accreting substantial mass in planetesimals. If confirmed, this would be a departure from other hot Jupiters observed hitherto, which have generally shown low H_2O abundances, indicative of the low accretion efficiency of H_2O -rich ices that is possible for disk-free migration mechanisms (Madhusudhan et al. 2014a; Pinhas et al. 2019; Welbanks et al. 2019). Such an abundance is also a substantial departure from expectations based on giant planets in our Solar System. The metallicity of Jupiter in multiple elements is $\sim 1\text{--}5 \times Z_{\odot}$ (Atreya et al. 2016; Li et al. 2020). With the mass of HAT-P-41 b being similar to that of Jupiter, its higher metallicity would indicate an even higher amount of solids accreted than that of Jupiter.

10. Summary

We have conducted a comprehensive, multipronged Bayesian retrieval analysis of the $0.3\text{--}5 \mu\text{m}$ transit spectrum of HAT-P-41b derived from HST STIS (previously unpublished; Section 4.1), HST WFC3 (reanalysis; Section 4.2), and Spitzer (independent analysis; Section 4.3) transit observations. We determined the host star has, at most, a low level of stellar activity ($\log L_X/L_{\text{bol}} < -5.2$) using both visible and X-ray photometric monitoring observations (Section 3.1).

We performed two complementary retrieval analyses: a relatively strict PLATON analysis (Sections 5.1, 7) assuming chemical equilibrium and solar metal ratios (except carbon), and a more flexible AURA free-chemistry retrieval (Sections 5.2, 8.1). Both methods’ fiducial models are excellent fits to the entire transit spectrum. We further tested an array of more complicated models (Sections 6 and 8.2), including instrumental transit depth biases (offsets), parametric Rayleigh scattering, partial cloud coverage, Mie scattering (PLATON only), and stellar activity (PLATON only). We find the conclusions to be insensitive to model choice within a paradigm.

Despite PLATON and AURA’s differing model assumptions, priors, and even opacity sources, we find several shared conclusions between the two methods (Section 9.1). Both PLATON and AURA retrieve a high atmospheric metallicity (O/H) that is inconsistent with Z_{\odot} to greater than 2σ ($\log_{10} Z/Z_{\odot} = 1.46^{+0.53}_{-0.68}$ compared to $\log_{10} Z/Z_{\odot} = 2.33^{+0.23}_{-0.25}$, respectively). They also both are consistent with a haze-free and cloud-free atmosphere, and both find a decisive water vapor detection and at least suggestive evidence of an optical absorption feature. We further confirm the result by performing a middle-ground retrieval, ATMO, and find results generally consistent with AURA’s (Section 9.1.2). We determine the inclusion of H_2 -broadened sodium opacity impacts the retrieved metallicities. While we consider AURA to be more physically plausible, due to its consistency with interior modeling constraints and inclusion of H_2 -broadened sodium opacity, we present the results from both PLATON and AURA as assumption-dependent orthogonal analyses. Overall, our study emphasizes the importance of comparative retrievals with different forward modeling, prior, and model selection assumptions in order to best contextualize presented results.

11. Postscript

After our retrieval analysis was complete, we were made aware of a recently accepted paper that used a model-grid fit to claim that a combined transit spectrum using HST’s WFC3

UV/Visible channel (UVIS) and independently derived Spitzer depths is best fit by a high-metallicity atmosphere (Wakeford et al. 2020). While we were aware that a manuscript was being submitted comparing STIS and UVIS data, we were unaware of the atmospheric metallicity claim. Because we were unaware of that result during our analysis, our conclusions were not biased by the results claimed in that paper.

The authors are grateful to Michael Zhang and Yayaati Chachan for many helpful discussions about PLATON. We thank Daniel Thornngren for providing the interior modeling metallicity constraint, as well as Daniel Kitzmann for providing the refractive indices of the Mie scattering condensates. We thank Eliza Kempton for productive conversations. This work is based on observations from the Hubble Space Telescope, operated by AURA, Inc. on behalf of NASA/ESA. This work also includes observations made with the Spitzer Space Telescope, operated by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA. K. B. S. acknowledges support from CRESST, as well as funding from HST grant HST-GO-14767. A. M. M. acknowledges support from HST grant HST-GO-14260, and the GSFC Sellers Exoplanet Environments Collaboration (SEEC), which is funded in part by the NASA Planetary Science Division’s Internal Scientist Funding Model. J. S. F. acknowledges support from the Spanish State Research Agency project AYA2016-79425-C3-2-P. This work has made use of data from the European Space Agency (ESA) mission

Gaia (<https://www.cosmos.esa.int/gaia>), processed by the Gaia Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement.

Software: IRAF (Tody 1986, 1993), SciPy (Jones et al. 2001), Matplotlib (Hunter 2007), nestle (<https://github.com/kbarbary/nestle>), BATMAN (Kreidberg 2015), Kapetyn (Terlouw & Vogelaar 2015), Corner.py (Foreman-Mackey 2016), AstroPy (Astropy Collaboration et al. 2018), PLATON (Zhang et al. 2019), Dynesty (Higson et al. 2019), NumPy (Harris et al. 2020), Pandas (The Pandas development Team 2020).

Appendix

HST Spectrophotometric Light-curve Fits

This section shows the spectrophotometric light curves for each HST grism observation. For each spectral bin, the normalized, detrended light curve data is plotted over the best-fit transit model. The spectral light curves are offset by a constant and alternate in transparency for clarity. Figures 19 and 20 illustrate the detrended light curves for HST STIS G430L (visits 83 and 84, respectively). Figure 21 illustrates the detrended light curves for HST STIS G750L (visit 85). Finally, Figure 22 shows the spectral light curves for the single HST WFC3 visit.

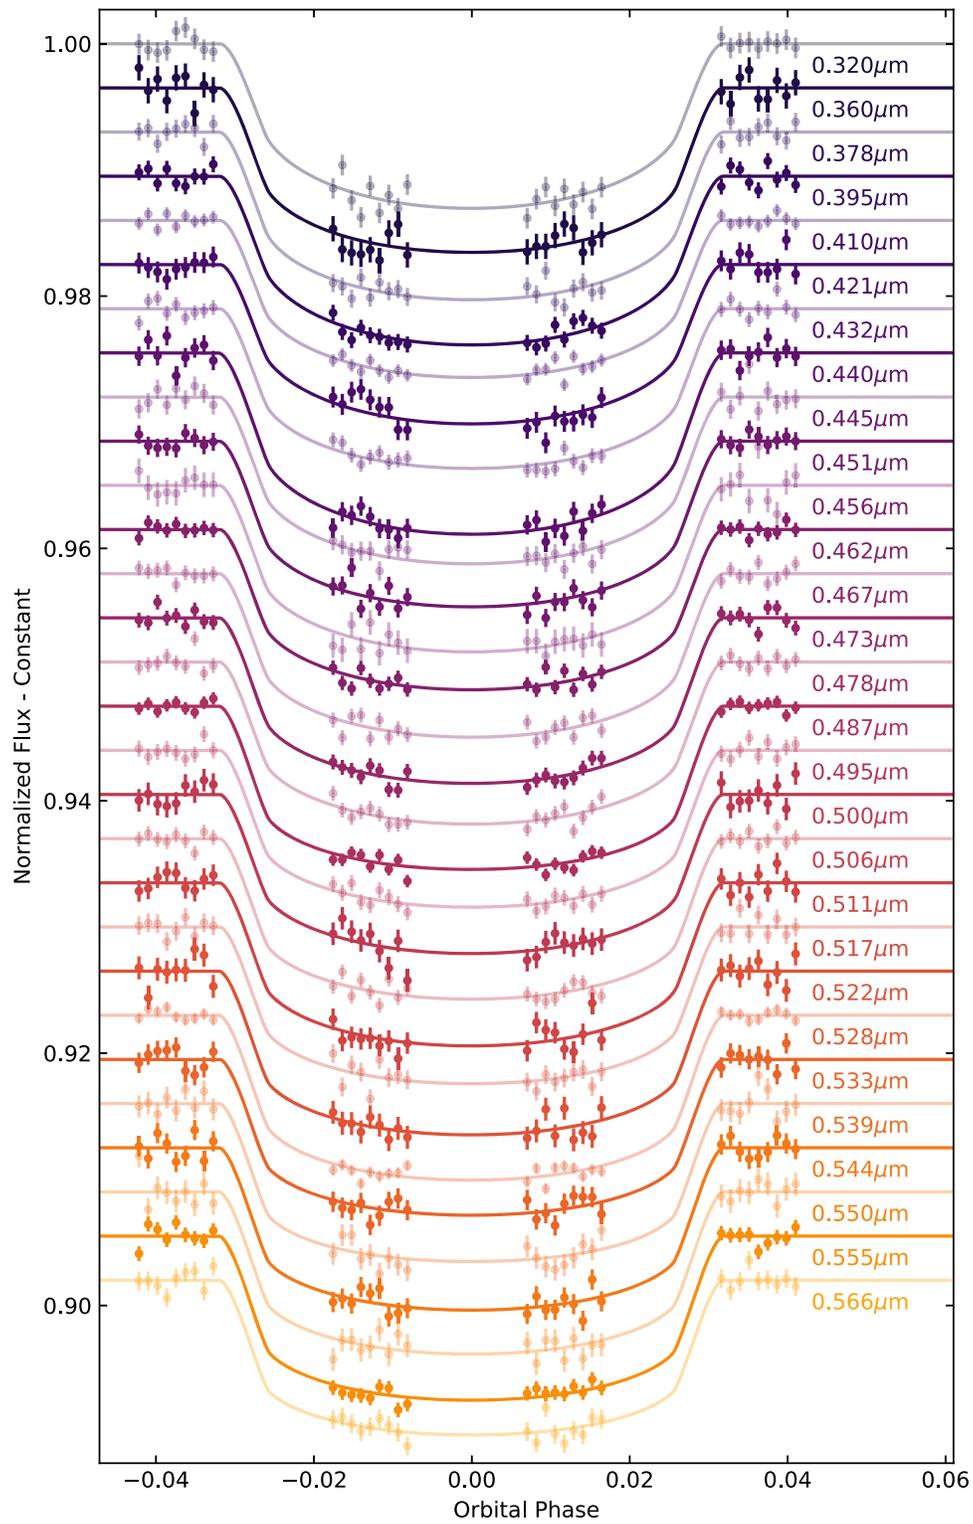


Figure 19. Spectral light curves for STIS G430L, visit 83.
(The data used to create this figure are available.)

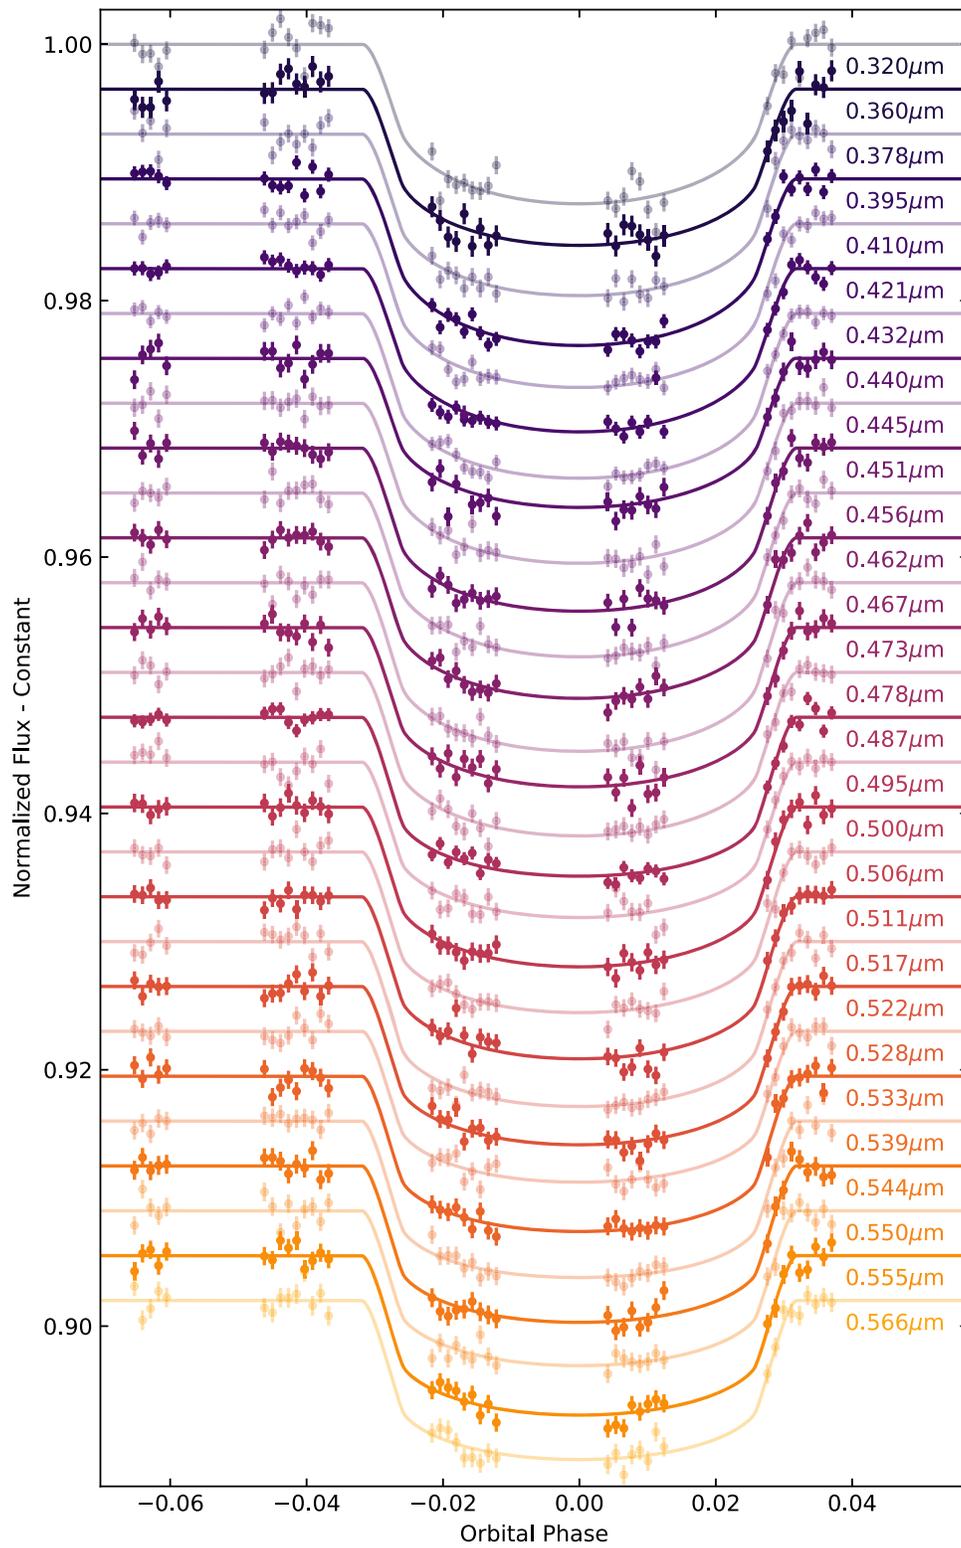


Figure 20. Spectral light curves for STIS G430L, visit 84.
(The data used to create this figure are available.)

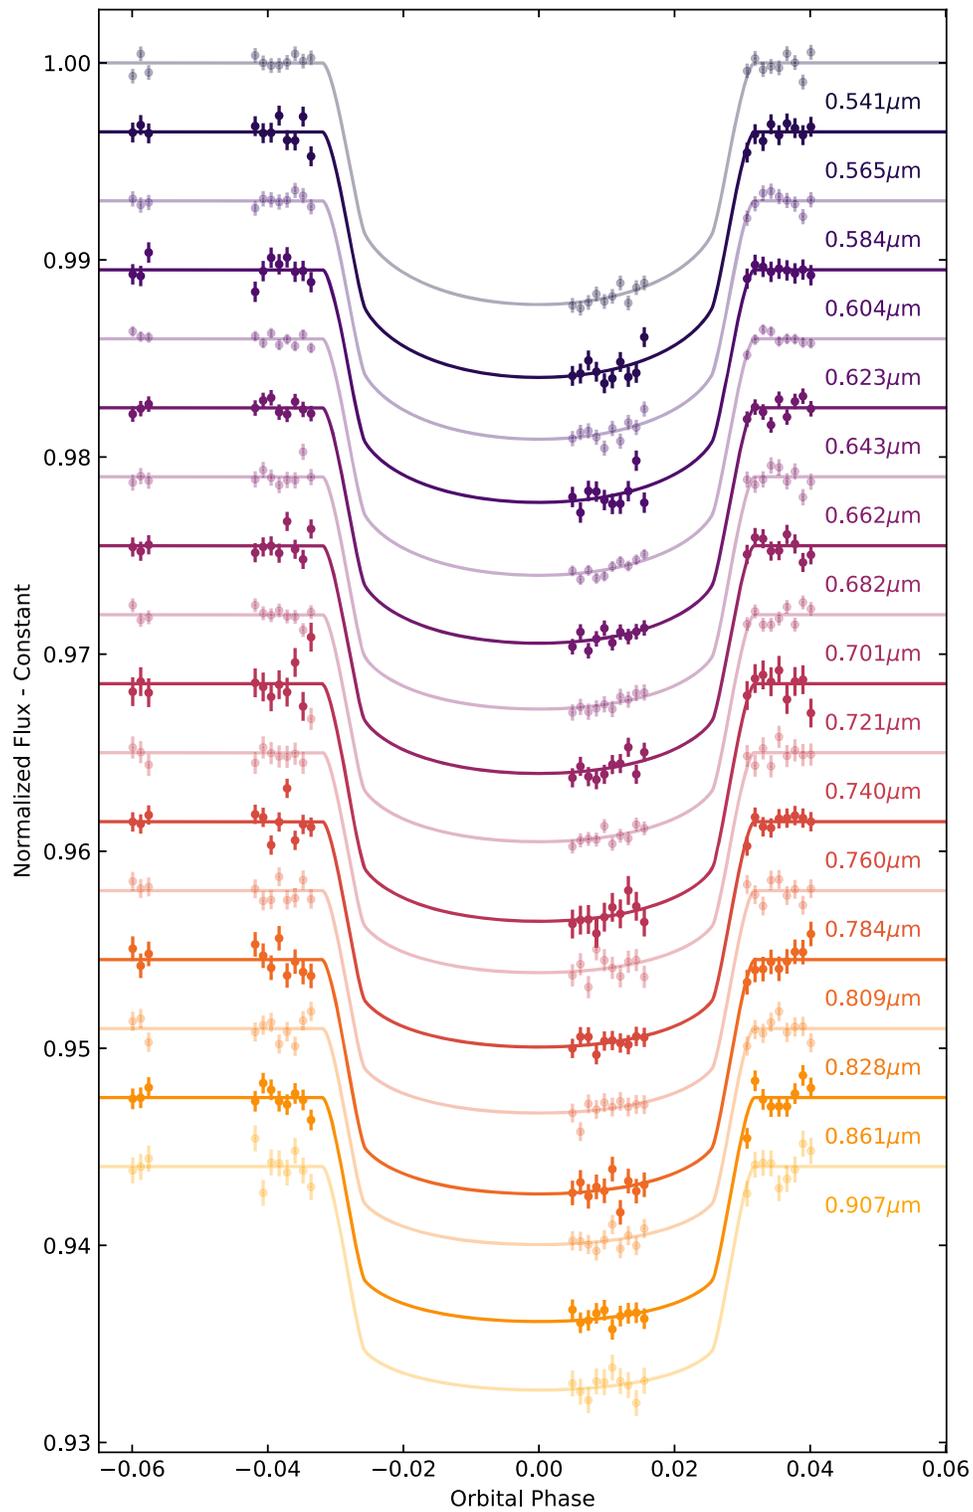


Figure 21. Spectral light curves for STIS G750L, visit 85.
(The data used to create this figure are available.)

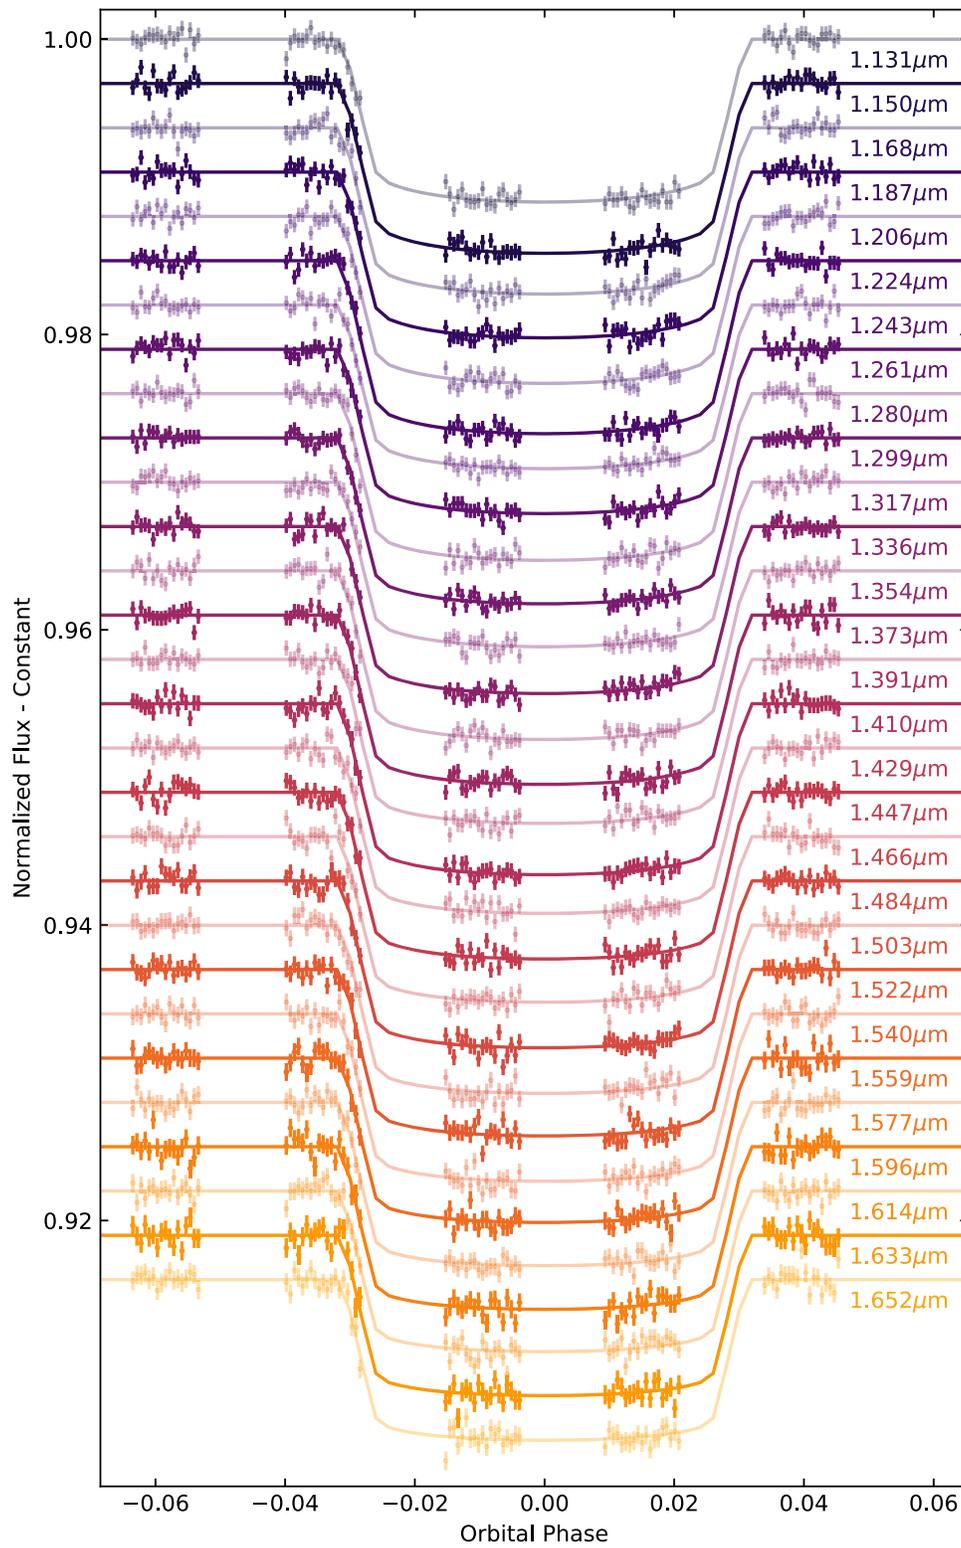


Figure 22. Spectral light curves for single WFC3 visit.
(The data used to create this figure are available.)

ORCID iDs

Kyle B. Sheppard <https://orcid.org/0000-0003-4552-9541>
 Luis Welbanks <https://orcid.org/0000-0003-0156-4564>
 Avi M. Mandell <https://orcid.org/0000-0002-8119-3355>

Nikku Madhusudhan <https://orcid.org/0000-0002-4869-000X>

Nikolay Nikolov <https://orcid.org/0000-0002-6500-3574>

Drake Deming <https://orcid.org/0000-0001-5727-4094>

Gregory W. Henry <https://orcid.org/0000-0003-4155-8513>

David K. Sing  <https://orcid.org/0000-0001-6050-7645>
 Mercedes López-Morales  <https://orcid.org/0000-0003-3204-8183>
 Jegug Ih  <https://orcid.org/0000-0003-2775-653X>
 Jorge Sanz-Forcada  <https://orcid.org/0000-0002-1600-7835>
 Panayotis Lavvas  <https://orcid.org/0000-0002-5360-3660>
 Thomas M. Evans  <https://orcid.org/0000-0001-5442-1300>
 Antonio García Muñoz  <https://orcid.org/0000-0003-1756-4825>
 Leonardo A. dos Santos  <https://orcid.org/0000-0002-2248-3838>

References

- Akaike, H. 1974, *ITAC*, **19**, 716
 Alexoudi, X., Mallonn, M., von Essen, C., et al. 2018, *A&A*, **620**, A142
 Allan, D. W. 1966, *IEEEP*, **54**, 221
 Allard, N. F., Kielkopf, J. F., & Allard, F. 2007, *EPJD*, **44**, 507
 Allard, N. F., Spiegelman, F., & Kielkopf, J. F. 2016, *A&A*, **589**, A21
 Allard, N. F., Spiegelman, F., Leininger, T., & Molliere, P. 2019, *A&A*, **628**, A120
 Amundsen, D. S., Baraffe, I., Tremblin, P., et al. 2014, *A&A*, **564**, A59
 Andrae, R., Fousneau, M., Creevey, O., et al. 2018, *A&A*, **616**, A8
 Andrae, R., Schulze-Hartung, T., & Melchior, P. 2010, arXiv:1012.3754
 Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, *ARA&A*, **47**, 481
 Astronomy Collaboration, Price-Whelan, A. M., SipHocz, B. M., et al. 2018, *AJ*, **156**, 123
 Atreya, S. K., Crida, A., Guillot, T., et al. 2016, arXiv:1606.04510
 Bailer-Jones, C. A. L., Rybizki, J., Fousneau, M., Mantelet, G., & Andrae, R. 2018, *AJ*, **156**, 58
 Barber, R. J., Strange, J. K., Hill, C., et al. 2014, *MNRAS*, **437**, 1828
 Barstow, J. K., Aigrain, S., Irwin, P. G. J., & Sing, D. K. 2017, *ApJ*, **834**, 50
 Barstow, J. K., Changeat, Q., Garland, R., et al. 2020, *MNRAS*, **493**, 4884
 Batygin, K., Bodenheimer, P. H., & Laughlin, G. P. 2016, *ApJ*, **829**, 114
 Benneke, B., Knutson, H. A., Lothringer, J., et al. 2019, *NatAs*, **3**, 813
 Benneke, B., & Seager, S. 2012, *ApJ*, **753**, 100
 Benneke, B., & Seager, S. 2013, *ApJ*, **778**, 153
 Berta, Z. K., Charbonneau, D., Désert, J.-M., et al. 2012, *ApJ*, **747**, 35
 Booth, R. A., Clarke, C. J., Madhusudhan, N., & Ilee, J. D. 2017, *MNRAS*, **469**, 3994
 Brewer, J. M., & Fischer, D. A. 2018a, *ApJS*, **237**, 38
 Brewer, J. M., & Fischer, D. A. 2018b, *yCat*, **2237**, 0
 Castelli, F., & Kurucz, R. L. 2003, in IAU Symp. 210, Modelling of Stellar Atmospheres, ed. N. Piskunov, W. W. Weiss, & D. F. Gray, **A20**
 Chachan, Y., Knutson, H. A., Gao, P., et al. 2019, *AJ*, **158**, 244
 Charbonneau, D., Brown, T. M., Noyes, R. W., & Gilliland, R. L. 2002, *ApJ*, **568**, 377
 Claret, A., Hauschildt, P. H., & Witte, S. 2013, *A&A*, **552**, A16
 Cubillos, P., Harrington, J., Loredo, T. J., et al. 2017, *AJ*, **153**, 3
 Deming, D., Knutson, H., Kammer, J., et al. 2015, *ApJ*, **805**, 132
 Deming, D., Seager, S., Richardson, L. J., & Harrington, J. 2005, *Natur*, **434**, 740
 Deming, D., Wilkins, A., McCullough, P., et al. 2013, *ApJ*, **774**, 95
 Eaton, J. A., Henry, G. W., & Fekel, F. C. 2003, *ASSTL*, **288**, 189
 Evans, D. F., Southworth, J., Maxted, P. F. L., et al. 2016, *A&A*, **589**, A58
 Evans, D. F., Southworth, J., Smalley, B., et al. 2018, *A&A*, **610**, A20
 Fisher, C., & Heng, K. 2018, *MNRAS*, **481**, 4698
 Foreman-Mackey, D. 2016, *JOSS*, **1**, 24
 Fortney, J. J. 2005, *MNRAS*, **364**, 649
 Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1
 Gaia Collaboration, Prusti, T., de Bruijn, J. H. J., et al. 2016, *A&A*, **595**, A1
 Gandhi, S., & Madhusudhan, N. 2017, *MNRAS*, **472**, 2334
 Garhart, E., Deming, D., Mandell, A., et al. 2020, *AJ*, **159**, 137
 Gelman, A., & Rubin, D. B. 1992, *StatSci*, **7**, 457
 Gibson, N. P. 2014, *MNRAS*, **445**, 3401
 Goyal, J. M., Mayne, N., Sing, D. K., et al. 2019, *MNRAS*, **486**, 783
 Griffith, C. A. 2014, *RSPTA*, **372**, 20130086
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
 Hartman, J. D., Bakos, G. Á., Békly, B., et al. 2012, *AJ*, **144**, 139
 Haynes, K., Mandell, A. M., Madhusudhan, N., Deming, D., & Knutson, H. 2015, *ApJ*, **806**, 146
 Henry, G. W. 1999, *PASP*, **111**, 845
 Higgs, E., Handley, W., Hobson, M., & Lasenby, A. 2019, *S&C*, **29**, 891
 Huitson, C. M., Sing, D. K., Pont, F., et al. 2013, *MNRAS*, **434**, 3252
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Isaacson, H., & Fischer, D. 2010, *ApJ*, **725**, 875
 Johnson, M. C., Cochran, W. D., Addison, B. C., Tinney, C. G., & Wright, D. J. 2017, *AJ*, **154**, 137
 Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>
 Kempton, E. M. R., Lupu, R., Owusu-Asare, A., Slough, P., & Cale, B. 2017, *PASP*, **129**, 044402
 Kilpatrick, B. M., Cubillos, P. E., Stevenson, K. B., et al. 2018, *AJ*, **156**, 103
 Kitzmann, D., & Heng, K. 2018, *MNRAS*, **475**, 94
 Kreidberg, L. 2015, *PASP*, **127**, 1161
 Kreidberg, L. 2018, in Handbook of Exoplanets, ed. H. Deeg & J. Belmonte (Cham: Springer), **100**
 Kreidberg, L., Bean, J. L., Désert, J.-M., et al. 2014, *ApJL*, **793**, L27
 Kreidberg, L., Line, M. R., Bean, J. L., et al. 2015, *ApJ*, **814**, 66
 Lee, J. M., Fletcher, L. N., & Irwin, P. G. J. 2012, *MNRAS*, **420**, 170
 Lewis, N. K., Wakeford, H. R., MacDonald, R. J., et al. 2020, *ApJL*, **902**, L19
 Li, C., Ingersoll, A., Bolton, S., et al. 2020, *NatAs*, **4**, 609
 Line, M. R., & Parmentier, V. 2016, *ApJ*, **820**, 78
 Line, M. R., Wolf, A. S., Zhang, X., et al. 2013, *ApJ*, **775**, 137
 MacDonald, R. J., & Madhusudhan, N. 2017, *MNRAS*, **469**, 1979
 Madhusudhan, N. 2012, *ApJ*, **758**, 36
 Madhusudhan, N. 2018, in Handbook of Exoplanets, ed. H. Deeg & J. Belmonte (Cham: Springer), **104**
 Madhusudhan, N., Amin, M. A., & Kennedy, G. M. 2014a, *ApJL*, **794**, L12
 Madhusudhan, N., Bitsch, B., Johansen, A., & Eriksson, L. 2017, *MNRAS*, **469**, 4102
 Madhusudhan, N., Crouzet, N., McCullough, P. R., Deming, D., & Hedges, C. 2014b, *ApJL*, **791**, L9
 Madhusudhan, N., & Seager, S. 2009, *ApJ*, **707**, 24
 Madhusudhan, N., & Seager, S. 2010, *ApJ*, **725**, 261
 Magic, Z., Chiavassa, A., Collet, R., & Asplund, M. 2015, *A&A*, **573**, A90
 Mandel, K., & Agol, E. 2002, *ApJL*, **580**, L171
 Mandell, A. M., Haynes, K., Sinukoff, E., et al. 2013, *ApJ*, **779**, 128
 McKemmish, L. K., Yurchenko, S. N., & Tennyson, J. 2016, *MNRAS*, **463**, 771
 Nikolov, N., Sing, D. K., Burrows, A. S., et al. 2015, *MNRAS*, **447**, 463
 Nikolov, N., Sing, D. K., Pont, F., et al. 2014, *MNRAS*, **437**, 46
 Noyes, R. W., Hartmann, L. W., Baliunas, S. L., Duncan, D. K., & Vaughan, A. H. 1984, *ApJ*, **279**, 763
 Öberg, K. I., & Bergin, E. A. 2016, *ApJL*, **831**, L19
 Öberg, K. I., Murray-Clay, R., & Bergin, E. A. 2011, *ApJL*, **743**, L16
 Patrascu, A. T., Yurchenko, S. N., & Tennyson, J. 2015, *MNRAS*, **449**, 3613
 Pinhas, A., Madhusudhan, N., Gandhi, S., & MacDonald, R. 2019, *MNRAS*, **482**, 1485
 Pinhas, A., Rackham, B. V., Madhusudhan, N., & Apai, D. 2018, *MNRAS*, **480**, 5314
 Pont, F., Sing, D. K., Gibson, N. P., et al. 2013, *MNRAS*, **432**, 2917
 Pont, F., Zucker, S., & Queloz, D. 2006, *MNRAS*, **373**, 231
 Rackham, B. V., Apai, D., & Giampapa, M. S. 2019, *AJ*, **157**, 96
 Richard, C., Gordon, I. E., Rothman, L. S., et al. 2012, *JQSRT*, **113**, 1276
 Rothman, L. S., Gordon, I. E., Barber, R. J., et al. 2010, *JQSRT*, **111**, 2139
 Schwarz, G. 1978, *AnSta*, **6**, 461
 Schwenke, D. W. 1998, *FaDi*, **109**, 321
 Shapiro, S. S., & Wilk, M. B. 1965, *Biometrika*, **52**, 591
 Sheppard, K. B., Mandell, A. M., Tamburo, P., et al. 2017, *ApJL*, **850**, L32
 Sing, D. K., Fortney, J. J., Nikolov, N., et al. 2016, *Natur*, **529**, 59
 Sing, D. K., Wakeford, H. R., Showman, A. P., et al. 2015, *MNRAS*, **446**, 2428
 Skilling, J. 2004, in AIP Conf. Ser. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. R. Fischer, R. Preuss, & U. V. Toussaint (Melville, NY: AIP), **395**
 Soubiran, C., Le Campion, J.-F., Brouillet, N., & Chemin, L. 2016, *A&A*, **591**, A118
 Stassun, K. G. 2019, *yCat*, **4038**, 0
 Stassun, K. G., Collins, K. A., & Gaudi, B. S. 2017, *AJ*, **153**, 136
 Stassun, K. G., Oelkers, R. J., Paegert, M., et al. 2019, *AJ*, **158**, 138
 Terlouw, J. P., & Vogelaar, M. G. R. 2015, Kapteyn Package, v2.3b6, Kapteyn Astronomical Institute, Groningen, <http://www.astro.rug.nl/software/kapteyn>
 The Pandas development Team 2020, pandas-dev/pandas: Pandas, Version 1.1.5, Zenodo, doi:10.5281/zenodo.3509134
 Thorngren, D., & Fortney, J. J. 2019, *ApJL*, **874**, L31

- Tody, D. 1986, *Proc. SPIE*, **627**, 733
- Tody, D. 1993, in ASP Conf. Ser. 52, *Astronomical Data Analysis Software and Systems II*, ed. R. J. Hanisch, R. J. V. Brissenden, & J. Barnes (San Francisco, CA: ASP), 173
- Tremblin, P., Amundsen, D. S., Chabrier, G., et al. 2016, *ApJL*, **817**, L19
- Tremblin, P., Amundsen, D. S., Mourier, P., et al. 2015, *ApJL*, **804**, L17
- Tremblin, P., Chabrier, G., Mayne, N. J., et al. 2017, *ApJ*, **841**, 30
- Trotta, R. 2008, *ConPh*, **49**, 71
- Tsiaras, A., Waldmann, I. P., Zingales, T., et al. 2018, *AJ*, **155**, 156
- Wakeford, H. R., & Sing, D. K. 2015, *A&A*, **573**, A122
- Wakeford, H. R., Sing, D. K., Deming, D., et al. 2013, *MNRAS*, **435**, 3481
- Wakeford, H. R., Sing, D. K., Deming, D., et al. 2018, *AJ*, **155**, 29
- Wakeford, H. R., Sing, D. K., Evans, T., Deming, D., & Mandell, A. 2016, *ApJ*, **819**, 10
- Wakeford, H. R., Sing, D. K., Kataria, T., et al. 2017, *Sci*, **356**, 628
- Wakeford, H. R., Sing, D. K., Stevenson, K. B., et al. 2020, *AJ*, **159**, 204
- Waldmann, I. P., Rocchetto, M., Tinetti, G., et al. 2015a, *ApJ*, **813**, 13
- Waldmann, I. P., Tinetti, G., Rocchetto, M., et al. 2015b, *ApJ*, **802**, 107
- Welbanks, L., & Madhusudhan, N. 2019, *AJ*, **157**, 206
- Welbanks, L., Madhusudhan, N., Allard, N. F., et al. 2019, *ApJL*, **887**, L20
- Winn, J. N., Fabrycky, D., Albrecht, S., & Johnson, J. A. 2010, *ApJL*, **718**, L145
- Woitke, P., Helling, C., Hunter, G. H., et al. 2018, *A&A*, **614**, A1
- Wöllert, M., & Brandner, W. 2015, *A&A*, **579**, A129
- Wright, N. J., Drake, J. J., Mamajek, E. E., & Henry, G. W. 2011, *ApJ*, **743**, 48
- Yurchenko, S. N., Barber, R. J., & Tennyson, J. 2011, *MNRAS*, **413**, 1828
- Yurchenko, S. N., & Tennyson, J. 2014, *MNRAS*, **440**, 1649
- Zahnle, K., Marley, M. S., Freedman, R. S., Lodders, K., & Fortney, J. J. 2009, *ApJL*, **701**, L20
- Zhang, M., Chachan, Y., Kempton, E. M. R., & Knutson, H. A. 2019, *PASP*, **131**, 034501
- Zhang, M., Chachan, Y., Kempton, E. M. R., Knutson, H. A., & Chang, W. H. 2020, *ApJ*, **899**, 27