



No Escaping Helium from 55 Cnc e*

Michael Zhang¹ , Heather A. Knutson² , Lile Wang³ , Fei Dai² , Antonija Oklopčič⁴ , and Renyu Hu⁵ 

¹Department of Astronomy, California Institute of Technology, Pasadena, CA 91125, USA; mzzhang2014@gmail.com

²Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, CA 91125 USA

³Flatiron Institute, 162 5th Ave., New York, NY 10010, USA

⁴Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

⁵Jet Propulsion Laboratory, 4800 Oak Grove Dr., Pasadena, CA 91109, USA

Received 2020 September 9; revised 2021 February 1; accepted 2021 February 3; published 2021 March 12

Abstract

We search for escaping helium from the hot super-Earth 55 Cnc e by taking high-resolution spectra of the 1083 nm line during two transits using Keck/NIRSPEC. We detect no helium absorption down to a 90% upper limit of 250 ppm in excess absorption or 0.27 mÅ in equivalent width. This corresponds to a mass-loss rate of less than $\sim 10^9$ g s⁻¹ assuming a Parker wind model with a plausible exosphere temperature of 5000–6000 K, although the precise constraint is heavily dependent on model assumptions. We consider both hydrogen- and helium-dominated atmospheric compositions and find similar bounds on the mass-loss rate in both scenarios. Our hydrodynamical models indicate that if a lightweight atmosphere exists on 55 Cnc e, our observations would have easily detected it. Together with the nondetection of Ly α absorption by Ehrenreich et al., our helium nondetection indicates that 55 Cnc e either never accreted a primordial atmosphere in the first place or lost its primordial atmosphere shortly after the dissipation of the gas disk.

Unified Astronomy Thesaurus concepts: [Exoplanet atmospheres \(487\)](#); [Exoplanet atmospheric composition \(2021\)](#); [High resolution spectroscopy \(2096\)](#); [Super Earths \(1655\)](#)

1. Introduction

The observed radius distribution of sub-Neptune-sized planets is bimodal, with peaks at $<1.5 R_{\oplus}$ and $2\text{--}3 R_{\oplus}$ (Fulton et al. 2017). This bimodality can be explained if the observed population of sub-Neptune-sized planets formed with several M_{\oplus} rocky cores and hydrogen-rich atmospheres, which were then stripped away from the most highly irradiated planets (i.e., Lopez & Fortney 2013; Ginzburg et al. 2018; see Owen 2019 for a literature review). The high inferred core densities in these models argue strongly for formation inside the ice line, and the semimajor-axis distribution of this population of planets is well matched by in situ formation models (Lee & Chiang 2017). While the prevailing evidence at the moment appears to favor in situ (or at least nearby) formation, the arguments proposed to date are by no means definitive, as they rely on indirect model-based inferences. If the mass-loss models used to simulate the observed radius distribution are incomplete or rely on incorrect assumptions, our conclusions about these planets may be incorrect as well. It is therefore important to have observational data to nail down theoretical models.

In this study, we focus on one of the most observationally favorable transiting super-Earths currently known: 55 Cnc e. 55 Cnc is a binary system with a K0 main-sequence star and an M-dwarf companion, separated by 1000 au. The primary is a bright ($V = 5.95$) star whose activity and rotation rate indicate that it is very old, probably around 10 Gyr. It has at least five planets of various sizes and orbital distances, including e, a superheated $R = 1.88 \pm 0.03 R_{\oplus}$ (Bourrier et al. 2018) super-Earth with a period of 0.74 days and the only planet known to transit. The mean density of e as derived from radial velocities

and transit depths is 6.7 g cm^{-3} , suggesting a rocky interior with an atmosphere contributing up to a few percent of the planet radius (Bourrier et al. 2018), although it is also consistent with a small iron core with a silicate mantle and no water or gas layer. At first glance 55 Cnc e seems unlikely to host a primordial hydrogen- and helium-rich atmosphere, as it has a relatively high equilibrium temperature and correspondingly high predicted escape rate (Valencia et al. 2010). However, there are large theoretical uncertainties in mass-loss models, driven in part by the uncertain X-ray and extreme ultraviolet (XUV) spectrum of the star (Owen 2019). It has also been suggested that this planet might have a lava ocean on its dayside that could outgas enough trapped hydrogen to form a thick secondary atmosphere (Chachan & Stevenson 2018).

Current observational constraints on 55 Cnc e's atmosphere are also conflicting. Some observations, such as the nondetection of hydrogen Ly α absorption by Ehrenreich et al. (2012) or the variations in infrared emission reported by Demory et al. (2016b), seem to indicate that it is unlikely to host a substantial hydrogen-rich atmosphere. Other observations suggest the opposite, including the planet's small measured day–night temperature gradient (Demory et al. 2016a), the tentative evidence for sodium and calcium absorption (Ridden-Harper et al. 2016), and the detection of a strong absorption feature in the planet's $1.1\text{--}1.7 \mu\text{m}$ transmission spectrum from Hubble Space Telescope (HST) WFC3, which has been attributed to HCN (Tsiaras et al. 2016). Regardless of its source, the large amplitude of the WFC3 absorption feature can only be matched by a hydrogen- or helium-dominated atmosphere, because an atmosphere dominated by heavier elements would have a much smaller scale height and correspondingly weaker absorption features during transit. If 55 Cnc e does have a significant low mean molecular weight atmosphere, its ability to retain this atmosphere in the face of ongoing mass loss would place strict constraints on the magnitude of relevant mass-loss processes, with correspondingly wide-

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

reaching implications for our understanding of the overall population of short-period super-Earths and sub-Neptunes.

In this paper, we use the helium 1083 nm metastable triplet (Oklopčić & Hirata 2018) to search for evidence of helium outflow from 55 Cnc e. Unlike previous Ly α observations, this helium triplet is easily accessible using ground-based telescopes and has been used to detect the extended atmospheres of multiple exoplanets. Most detections have been around Jupiter-radius planets such as WASP-107b (i.e., Spake et al. 2018), with the 4 R_{\oplus} GJ 3470b being the smallest planet with a detection to date (Ninan et al. 2020; Palle et al. 2020). For 55 Cnc e, our observations of the 1083 nm helium line supplement the Ehrenreich et al. (2012) Ly α observations in shaping our understanding of the exosphere. The Ehrenreich et al. (2012) observations took place at a stellar activity minimum, when the star’s X-ray flux was 2–3 times lower than during our observations (see Section 6); this decreased X-ray flux might have suppressed the mass-loss rate during these observations. In addition, some studies have proposed that mass loss over many gigayears can preferentially remove hydrogen from the atmosphere of a small planet while leaving helium behind (Hu et al. 2015; Malsky & Rogers 2020). A helium-dominated atmosphere would be consistent with the molecular weight inferred from the HST transit spectrum (Tsiaras et al. 2016). Such an atmosphere might be undetectable in Ly α but easy to detect using metastable helium. Finally, because Ly α is such a strong line, the observed signal is dependent on the behavior of the diffuse exosphere far from the planet. The atmospheric absorption signal in the helium 1083 nm triplet is typically much weaker than the absorption signal in the Ly α line, but IR measurements are significantly more precise than UV measurements owing to the much higher photon flux, and (unlike Ly α) we are able to observe the cores of the lines. This means that helium observations are sensitive to gas at smaller radii and with lower outflow velocities than Ly α observations, making it easier to compare to mass-loss models. It is for these three reasons—the differing stellar XUV irradiation, the possibility of a helium-dominated atmosphere, and the complementary sensitivities—that observations of the helium line are meaningful even for small planets with existing Ly α nondetections.

We describe our observations in Section 2, the data reduction pipeline in Section 3, analysis in Section 4, outflow models in Section 5, and scientific implications in Section 6.

2. Observations

We observed two transits of 55 Cnc e in Y band using the upgraded NIRSPEC instrument on Keck (Martin et al. 2018): one on 2019 December 4 at 13:56 UTC (barycentric), and one on 2019 December 18 at 13:48 UTC (barycentric). These observations used the $0''.288 \times 12''$ slit, giving NIRSPEC a resolution of 37,500, with an FWHM sampling of 3 pixels. All observations were performed with 60 s exposure times in an ABBA nod pattern to facilitate background subtraction. Details of the observations are given in Table 1.

Observation conditions were stable with good seeing during the first night. The seeing was worse and more variable during the second night, decreasing the per-pixel count rate on the detector and making it possible to use a more efficient observing strategy of 20 s subexposures with three co-adds instead of 15 s subexposures with four co-adds. The marginally increased observing efficiency could not fully compensate for the higher seeing, leading the second night to have a marginally worse signal-to-noise ratio (S/N) per spectrum.

Table 1
Keck/NIRSPEC Observations

Parameter	12/04/19 (UTC)	12/18/19 (UTC)
Obs. duration	3.2 hr	3.7 hr
Number of obs.	107	136
Obs. before transit	23	40
Obs. during transit	51	56
Obs. after transit	33	40
Air mass before transit	1.1	1.04
Air mass during transit	1.02	1.01
Air mass after transit	1.03	1.11
Obs. efficiency	56%	61%
Avg. S/N ^a	660	640

Note.

^a Average S/N is calculated per spectral pixel per exposure, based on the estimated error from optimal extraction.

For unknown reasons and at unpredictable times, the telescope would fail to nod. We encountered this issue on each of our four half-nights with NIRSPEC—two in 2019 April, and two in 2019 December. Sometimes it would not move at all; at other times, it would jump out of the slit, either by nodding too far or by adding a perpendicular component to the nod. When this happened, we would wait for the exposure to finish (NIRSPEC does not allow us to stop during the middle of an exposure), recenter the star, and restart the nod sequence. We then discard the exposure, together with the previous exposure if it is necessary for background subtraction. On December 4, nodding failed once during the transit, and we lost four exposures (because we did not notice the failure immediately), corresponding to 7 minutes of observation time. On December 18, nodding failed twice during the transit and once after the transit, resulting in a loss of six exposures. Other than these nodding failures, our observing sequence is continuous on both nights.

3. Data Reduction

We calibrated the raw images and extracted 1D spectra for each order using a custom Python pipeline designed for the upgraded NIRSPEC. We describe each step of this process below.

3.1. Cross-talk Removal

In its current configuration NIRSPEC is divided into 32 readout channels, with 64 rows per channel. The rows of adjacent channels are read simultaneously, but in reverse order: row 0 of channel 0 should be read at the same time as row 63 of channel 1. The simultaneous readout causes cross-talk signals between each pair of channels. We follow the approach described in George et al. (2020) and initially assume that the cross-talk signal $X_{k,j,i}$ in NIRSPEC is proportional to the derivative of the signal in the source channel:

$$X_{k,j,i} = a_{k,j}(S_{k,i} - S_{k,i-1}), \quad (1)$$

where channel k is the source of the cross talk, channel j is the destination, i is the row number, and a is a 32×32 matrix of scaling factors whose elements are of the order 400 ppm. The cross talk introduced is therefore on the order of a few hundred electrons for signals on the order of 1,000,000 electrons, but it varies widely for different k, j pairs.




Figure 1. Portion of a raw A-B frame, showing the order containing the helium line. Top: no calibration corrections. Bottom: cross talk subtracted. Notice that the subtle horizontal ripple pattern is gone, as are the two sharp horizontal lines.

Unfortunately, there are no calibration data for NIRSPEC that can help to characterize the cross talk, as we require a high flux derivative along the spatial axis. We therefore use archival NIRSPEC observations of 55 Cnc in *L* band from 2019 April 2 and 16 to measure the cross-talk matrix a . On each night, we took ~ 50 minute-long exposures, each with very high flux ($\sim 300,000$ electrons per spectral pixel). To measure the cross-talk matrix a , we masked out pixels with a count greater than 70 ADU (200 electrons) and used a linear least-squares fit to estimate a and the uncertainty on a . The uncertainty is high where the source channel is unilluminated and/or the target channel has few unilluminated pixels and low where the source channel is illuminated but the target channel is not. We perform a weighted average over the estimates of a from each exposure on each night to obtain our best estimate for the matrix.

Relying on science data to measure the cross-talk matrix is less than ideal because channels that have no trace do not produce a measurable cross-talk signal in other channels and the trace dominates the cross talk from other channels in channels where the trace is nearly horizontal and close to the middle. Nevertheless, a visual inspection indicates that this imperfect approach to cross-talk subtraction is sufficient to remove the cross-talk signal from our data (Figure 1).

We find that for NIRSPEC, Equation (1) is less accurate when j and i are of different parity. By carefully examining the cross-talk patterns, we found that the cross-talk signal is proportional not to $S_{k,i} - S_{k,i-1}$ but to $S_{k,i+2} - S_{k,i+1}$. In physical terms, this means that the second rows of odd channels (not the zeroth rows) are read at the same time as the 63rd rows of even channels. Similarly, the first row corresponds to the 62nd, the second row to the 61st, and so on, until the 63rd rows of odd channels are read simultaneously with the second rows of even channels. Therefore, rows 0 and 1 of every channel are being read out when their sister channels of opposite parity are not being read out.

We believe that this mismatch in the readout pattern causes another detector artifact: the anomalous rows, seen in Figure 1 as horizontal lines. Anomalous rows occur when the row number modulo 128 is equal to 0, 1, or 64. The first two correspond to rows 0 and 1 of even channels; the last corresponds to row 0 of odd channels. We speculate that the anomalous rows could be caused by cross talk: when row 0 or 1 is being read out, the corresponding readout line in opposite-parity channels could be carrying signals of much higher amplitude than image data, causing much higher cross talk than normal.

In the absence of photons, the anomalous rows are mostly, but not entirely, consistent across the entire image. Thus, all 32 rows divisible by 128 are similar, all rows whose remainder is 1 when divided by 128 are similar, and all rows whose remainder is 64 when divided by 128 are similar. We create a template of these three different categories of anomalous rows by identifying regions of the detector that received few photons, either because they are in between orders or because they are within an order but far from the trace. At each column, the template is equal to the median of the rows fitting this criterion. The template is subtracted from all relevant rows. This method works very well, but not perfectly—a faint hint of under-subtraction can be seen in Figure 1 (bottom) for row 128.

3.2. Image Calibration

On the first (second) night, we took 80 (63) flat fields, each with an exposure time of 4.4 s per co-add and 20 co-adds. In total, we collected 750 million (590 million) electrons per pixel, ensuring that the photon error in the flat fields is far below the photon noise in the observational data. To calibrate the flats, we took 19 (3) darks, each with an exposure time of 4.4 s per co-add and 20 co-adds. We create a master dark by median-stacking the individual dark frames. In each individual dark, pixels that deviate from the image-wide mean by more than 5σ are marked as bad pixels. Pixels marked as bad in more than half of the individual darks are marked as bad in the master dark.

We create a master flat by taking the median of the individual flats. Prior to combining, we subtract the cross talk from each flat frame using the algorithm described in the previous subsection, subtract the master dark, and divide by the median flux. We identify the order containing the helium line and mask out everything else, as there is no need to extract spectra from other orders. The relevant order is fitted with a polynomial that is fifth order in x and third order in y and then divided by the polynomial, in order to bring the values of all good pixels close to 1. We then create a mask of bad pixels, including both the bad pixels identified in the master dark and the pixels with a flat value lower than 0.5 or higher than 1.5.

The final master flat and the bad pixel mask are saved in a FITS file.

Finally, we correct the raw science frames using the master flat. For each $A_1B_1B_2A_2$ nod, we create four difference images: $A_1 - B_1$, $B_1 - A_1$, $B_2 - A_2$, and $A_2 - B_2$. Each difference image is cross talk subtracted, multiplied by the gain of $g = 2.85 e \text{ ADU}^{-1}$, and divided by the master flat F . We next construct a variance image for each differenced frame, which indicates the uncertainty in the measured flux at each individual pixel location. This variance image includes the photon noise from the star, the sky background, and the detector read noise. The background b , read noise N_R , and total variance V are computed as follows:

$$b = g \frac{A + B - |A - B|}{F} \quad (2)$$

$$N_R = \sqrt{2} \sqrt{\frac{N_{\text{coadds}}}{N_{\text{reads}}}} N_{R,0} \quad (3)$$

$$V = g \frac{|A| + |B|}{F^2} + N_R^2, \quad (4)$$

where $N_{R,0} = 44e^-$, $N_{\text{reads}} = 4$ for all of our observations, $g = 2.85 e^- \text{ ADU}^{-1}$ is the gain, and $N_{\text{coadds}} = 4$ on our first night and $N_{\text{coadds}} = 3$ on most of our second night. The iterative bad pixel algorithm in REDSPEC 3.0⁶ (Kim et al. 2015), which we ported to Python, is used to identify and repair bad pixels. This algorithm identifies bad pixels with a variant of local sigma clipping. In total, it typically identifies and repairs 1000 hot pixels and 1000 cold pixels in a given image. Since the identification is limited to the ~ 100 pixel tall order containing the helium line, 2000 bad pixels represents 1% of all pixels. The bad pixels identified by the algorithm are combined with the bad pixels identified in the master dark and master flat to create a master bad pixel mask. The difference image, variance image, background image, and bad pixel mask are all saved in a FITS file.

3.3. Optimal Extraction

After calibrating the images, we extract the 1D spectrum from each 2D spectral trace. We first determine the position of the trace in each image. For every column (corresponding to one wavelength), we fit a Gaussian to the pixel values to estimate the trace position. We then fit a fifth-order polynomial to the trace locations as a function of column number. The residuals in this final fit are typically smaller than 0.01 pixels. Accuracy is not paramount because we only use the trace to identify regions of the image very far from the trace, in order to mask them out.

After determining the position of the trace, we perform optimal extraction using a variant of the method described in Horne (1986). The original Horne (1986) algorithm assumes that the wavelength axis is aligned with the columns and the spatial axis is aligned with the rows. Unfortunately, on NIRSPEC neither the wavelength nor spatial axes are aligned with either the rows or the columns. The axes are also not perpendicular to each other, and neither axis is straight, as can be seen in Figure 1. The typical way around this problem is to rectify the order spatially and spectrally by interpolation onto a

rectilinear grid. This works well enough for low-S/N data, but for our exceptionally high S/N data it introduces aliasing artifacts at the 0.1%–1% level.

We instead use a rolling window approach where we assume that each column corresponds to one wavelength—a reasonable assumption, given that most of the flux is concentrated within 5 pixels of the center of the trace. Regions more than 15 pixels from the center of the trace are masked out to avoid interference from neighboring orders or cosmic rays. For each column, we take a window 81 pixels wide and 72 pixels tall and fit the profile for every row. We chose 81 pixels as the width because we want the window to be small enough for the trace to deviate vertically by less than 1 pixel but big enough to fit the profile accurately. We chose a 72-pixel window in the y (cross-dispersion) direction because the trace is 50 pixels lower on the left end of the detector than on the right, and we needed a window large enough to encompass not just the center of the trace but also the wings across the full width of the image.

In Horne (1986), the spatial profile is estimated by dividing each column by its sum. This runs into problems in regions of high telluric absorption, where the spectral flux approaches zero. We therefore fit the product of the spectrum and a Chebyshev polynomial model of the profile directly to the observed data. Mathematically, if i represents the order of the Chebyshev polynomial T_i and j is the column number, we look for the x that minimizes $Ax = b$, where A and b are

$$A_{i,j} = S_j T_i \left(\frac{j - N/2}{N} \right) / \sigma_j \quad (5)$$

$$b_j = I_j / \sigma_j, \quad (6)$$

$N = 72$ is the size of the window, and $\frac{j - N/2}{N}$ normalizes the column numbers to range from -1 to $+1$. After obtaining the least-squares solution to x , the profile can be computed as

$$P_j = \sum_{i=0}^5 x_i T_i \left(\frac{j - N/2}{N} \right), \quad (7)$$

where 5 is the the maximum order of Chebyshev polynomials we fit.

We note that this is mathematically equivalent to fitting monomials. If we had replaced $T_i(y)$ with y^i in all the equations above, we would arrive at an identical P_j . The advantage of Chebyshev polynomials comes from the numerical stability and robustness of the linear algebra solver. y^i becomes extremely small when i is large for all but the boundary values -1 and 1 , causing numerical problems, whereas the Chebyshev polynomials are well behaved for every order. The Chebyshev polynomials are also mutually orthogonal; the monomials are not. In practice, it appears that neither of these matter for our low-order profile fit, but since Chebyshev fitting is theoretically superior, we adopt it as our preferred form of polynomial fitting throughout our code.

We make several additional modifications to the Horne (1986) algorithm. The background cannot be estimated from the difference image alone, so we instead use the estimate from Equation (2). Instead of identifying bad pixels solely from the image using sigma clipping, we start with the bad pixel mask from the calibration stage (Section 3.2) and identify additional bad pixels using the image. Finally, we use an extremely high threshold $\sigma_{\text{clip}} = 12$ to reject bad pixels during the optimal extraction iterations. We require this high threshold because of a phenomenon that we observe with our high-S/N data: the

⁶ <https://www2.keck.hawaii.edu/inst/nirspec/redspeg.html>




Figure 2. Top: residuals of optimal extraction, scaled to the standard deviation of each pixel and clipped from -5σ to 5σ . If optimal extraction worked perfectly, this image would consist of Gaussian-distributed noise with a mean of 0 and a standard deviation of 1. The four-leaf-clover pattern in the middle of the residuals image corresponds to a deep stellar line at $1.0830\ \mu\text{m}$. The position of the helium $1.08333\ \mu\text{m}$ line is marked in red. Although a stellar helium line is present, it is much shallower than the $1.0830\ \mu\text{m}$ line, and there does not appear to be any strong residual structure in this region. Bottom: optimally extracted spectrum with accompanying variance. The spikes in variance are due to masks on bad pixels.

four-leaf-clover pattern in the residuals image, shown in Figure 2. This pattern occurs wherever the spectrum is changing rapidly—namely, in the vicinity of deep stellar or telluric lines. The pattern traces the cross derivative $\frac{dF}{dx dy}$ and is more pronounced where the trace is highly tilted with respect to the horizontal. This clover pattern indicates that optimal extraction does not give optimal results in the vicinity of deep lines. Regions without deep lines—including the vicinity of the $1.08333\ \mu\text{m}$ helium line—are minimally affected.

We initially thought that the clover pattern was due to the misalignment between the image axes and the wavelength/spatial axes, or due to the nonorthogonality between the wavelength and spatial axes. We tried many different optimal estimation schemes that do not make the rectilinear assumptions of Horne (1986), including Marsh (1989) and an unpublished algorithm,⁷ in an attempt to get rid of the clovers. However, we eventually realized that the clover pattern is caused by the inseparability of the 2D point-spread function (PSF). The optimal extraction algorithm assumes that $\text{PSF}_{\lambda,x} = S(\lambda)P(x)$ for some spectrum $S(\lambda)$ and profile $P(x)$. A 2D Gaussian with its axes aligned with the wavelength and spatial axes is separable, but a tilted 2D Gaussian is not. Bolton & Schlegel (2010) simulate spectra under the assumption of PSF nonseparability, apply optimal extraction to the simulated spectra, and obtain clover patterns strikingly similar to ours (see their Figure 1, right panel). The nonseparability of the PSF means that only algorithms like spectro-perfectionism (Bolton & Schlegel 2010) that adopt a 2D PSF can avoid the clovers. Spectro-perfectionism is not necessary for our work here because the helium line is not in a region where the spectrum varies rapidly.

3.4. Wavelength Solution

After extracting the 1D spectra, we determine the wavelength solution for each spectrum. We first create a template containing stellar and telluric lines at known wavelengths, with the stellar lines shifted to account for Earth’s velocity relative to the star on that night. We adopted a $T_{\text{eff}} = 5200\ \text{K}$, $\log(g) = 4.5$, and $[\text{M}/\text{H}] = 0.0$ PHOENIX model spectrum (Husser et al. 2013) for the star. We used a telluric transmission spectrum from the Gemini website,⁸ which assumes a precipitable water vapor of 1.6 mm and an air mass of 1.5. This spectrum was calculated using ATRAN (Lord 1992). We

multiply the stellar spectrum by the telluric transmission and then down-sample to instrumental resolution ($R = 37,500$) to get our final template.

We parameterize the wavelength solution as a third-order polynomial function of the normalized column number, $j' = \frac{j - N/2}{N}$:

$$\lambda(j') = \sum_{i=0}^3 c_i T_i(j'), \quad (8)$$

where T_i are the Chebyshev polynomials. Similarly, we parameterize the continuum as a fifth-order polynomial of j' and multiply it by the template. We then use `scipy`’s differential evolution minimizer to minimize χ^2 , which we calculate as the difference between the observations and the continuum-adjusted template interpolated using the proposed wavelength solution coefficients. A visual inspection of the resulting fit indicates that the line positions match to better than a pixel.

4. Analysis

After extracting the 1D spectra, we place the data from each night on a uniform wavelength grid and remove signals not related to the planet. This includes instrumental effects like detector fringing, telluric absorption lines, and absorption lines from the star itself. We then shift each spectrum into the planetary rest frame and quantify the amount of excess absorption in the helium triplet during transit. We describe each step of this process in detail below.

4.1. Better Ephemeris

Due to the high radial acceleration of the planet during transit, an accurate ephemeris is necessary to shift spectra into the planet frame. We calculate an updated ephemeris by combining the epoch derived by Demory et al. (2016b) using four Spitzer $4.5\ \mu\text{m}$ transits, the epoch derived by Sulis et al. (2019) using 143 MOST transits, and 35 individual transit timings from TESS. We obtain the TESS transit timings using a procedure similar to that described in Dai et al. (2019). In short, we downloaded the TESS photometry from the Mikulski Archive for Space Telescopes (MAST). Using the archival ephemeris, we isolated data within a wide window of three times the archival transit duration. We fitted each individual transit with the BATMAN (Kreidberg 2015) package and a quadratic function of time to account for local stellar variability. We adopted quadratic limb darkening and imposed Gaussian priors with widths of 0.3

⁷ https://code.obs.carnegiescience.edu/Algorithms/ghlb/at_download/file

⁸ <https://www.gemini.edu/observing/telescopes-and-sites/sites>

centered around the theoretical values from EXOFAST⁹ (Eastman et al. 2013). We then fitted all TESS transits globally after removing the stellar variability component. With this global model as a template, we revisit each individual transit, allowing only the mid-transit time and local quadratic function to vary. This process is iterated a few times until convergence. Finally, the individual TESS transit epochs are fitted together with the archival transit times. The updated ephemeris is $P = 0.73654604 \pm 1.6 \times 10^{-7}$ days and $T_0 = 2,458,723.38328 \pm 0.00014$ (BJD_{TDB}) with negligible covariance $C = -10^{-14}$ days. The covariance C is defined such that the prediction error at epoch E from this ephemeris is

$$\sigma_T^2 = \sigma_{T_0}^2 + E^2 \sigma_P^2 + 2CE. \quad (9)$$

We chose the initial epoch T_0 to make C as close to 0 as possible, allowing the last term to be ignored. The prediction error evaluates to 13 s at the time of our 2019 December observations.

4.2. Making a Spectral Grid

The first step is to linearly interpolate all spectra on a given night onto a common wavelength grid. We choose a wavelength grid spanning 1.0826–1.0840 μm , which encompasses the locations of the three lines in the helium triplet. We select a resolution of 110,000 for our grid, approximately matching the native pixel resolution. The process of interpolation introduces covariances between adjacent wavelength bins, artificially smoothing out the interpolated spectrum. We keep track of the covariance matrix during our analysis and eventually use it to calculate likelihoods.

We do an initial check for helium absorption during the transit by dividing all spectra from both nights into two categories, in-transit and out-of-transit. We take the mean of the in-transit spectra to create a master in-transit spectrum and the mean of the out-of-transit spectra to create a master out-of-transit spectrum. We then calculate the excess absorption in the stellar frame as $F_{\text{in}}/F_{\text{out}} - 1$. The two master spectra and the excess absorption are plotted in Figure 3. There does not appear to be any detectable increase in absorption during the transit at the position of the helium line.

4.3. Residuals Image

Since we are interested in fractional changes in the spectrum, we take the natural log of the spectral grid and subtract the mean of every row and column, producing what we call the residuals image. Every pixel in the residuals image approximately represents the fractional flux change at that epoch and wavelength from the mean spectrum. Taking the natural log also has the advantage of linearizing the effect of changing air mass. The observed flux is roughly $F(\lambda) = F_{\text{vac}}(\lambda)e^{-\alpha z}$, where z is the air mass, so $\ln F(\lambda) = \ln F_{\text{vac}}(\lambda) - \alpha z$.

We show the residuals image through various steps of the pipeline in Figure 4 (top). The most striking feature in the median-subtracted image are the vertical bars caused by fringing, which are spaced 20 pixels apart. A fast Fourier transform (FFT) of each spectrum reveals a prominent peak at a frequency of 0.052 pixel^{-1} . To remove the fringing, we use `scipy` to apply a second-order infinite impulse response notch digital filter with a frequency of 0.052 pixel^{-1} and a quality factor of 15. We apply




Figure 3. Top: in-transit vs. out-of-transit spectra in the stellar rest frame, calculated using data from both nights. Bottom: excess absorption ($F_{\text{in}}/F_{\text{out}} - 1$), plotted in the stellar frame; fringing has been removed with a band-stop filter, but additional corrections for telluric and instrumental effects described in Section 4.3 have not been applied. The red vertical lines mark the locations of the three helium lines. The dashed vertical black lines represent strong stellar lines, while the dashed-dotted vertical line is a telluric water line. The most prominent features in the excess absorption plot are the continuum variation and the sudden spikes/dips at the position of the strong lines. The latter is due to the poor performance of optimal extraction in the vicinity of strong lines. Helium absorption would manifest as a spike in the vicinity of the red vertical lines, which is not seen.

the filter twice, which is sufficient to suppress the peak in the FFT without bringing the spectral power substantially below that of neighboring frequencies. The subsequent panel in Figure 4 shows the residuals image after fringing correction, demonstrating that the notch filter has effectively removed the fringing.

The most prominent features in the fringing-corrected residuals image (Figure 4) are the strongly variable columns. These are deep telluric and stellar lines that vary for a variety of reasons, including the suboptimal performance of optimal extraction in deep lines (discussed in Section 3.3), the inaccuracy of interpolation across deep lines, and the time-variable telluric absorption. The line immediately to the right of center in the residuals image, for example, is a water absorption line at 1.08351 μm . Its dependence on air mass is clear: as the night progresses, 55 Cnc first rises and then sets. The line is dark at the beginning, brightens until the air mass reaches its

⁹ astroutils.astronomy.ohio-state.edu/exofast/limbdark.shtml




Figure 4. Pipeline steps: the original spectra; after removing fringing; after removing one SYSREM component; after masking variable lines and subtracting off continuum variations; after subtracting off continuum variations by fitting a polynomial to each spectrum. In the last panel, the vertical white lines mark a stellar Si I line at $1.0830057 \mu\text{m}$ (left) and a strong water line at $1.08351 \mu\text{m}$ (right).

minimum, and dims again. The other strong line is an Si I line at $1.0830057 \mu\text{m}$.

We correct for variability in telluric absorption using SYSREM (Mazeh et al. 2007). SYSREM is a generalization

of principal component analysis (PCA) that takes into account the errors on the data. Like PCA, it identifies eigenvectors and eigenvalues that, when linearly combined, best explain the residuals image. When we apply SYSREM to our residuals image, the first component has eigenvalues that closely track the air mass, showing that the algorithm is successfully identifying telluric variability.

After subtracting the first principal component identified by SYSREM, we compute the standard deviation of each column and mask the most variable columns. This removes the prominent telluric line and all four of the most prominent stellar lines. One iteration of SYSREM is not sufficient to remove the continuum variation seen in Figure 4. We therefore remove it by fitting a third-order polynomial with respect to column number for each row and subtracting the polynomial. The residuals image, averaged across both nights, is shown in the bottom panel of Figure 4. In Figure 5, we additionally show a plot of excess absorption in the stellar frame. There is no sign of helium absorption in either plot.

4.4. Shifting into the Planetary Rest Frame

55 Cnc has an extremely high orbital speed of 230 km s^{-1} , causing it to accelerate by 130 km s^{-1} over the duration of the transit. At every epoch, we use the barycentric Julian date to compute the radial velocity of the planet relative to the star, shift the residual spectrum accordingly, and resample onto the common wavelength grid using linear interpolation. The covariances introduced by linear interpolation are properly computed and propagated. We then take the mean of all the shifted in-transit residual spectra to arrive at the final excess absorption spectrum. The average spectrum across both nights is taken to be the fiducial excess absorption spectrum.

The excess absorption from both nights, as well as the averaged excess absorption, is shown in Figure 6. The combined spectrum has a standard deviation of 146 ppm, smaller than the planet’s white-light transit depth of 350 ppm, although some amount of correlated noise is present. No sign of helium absorption can be seen.

The bottom panel of Figure 6 shows what the excess absorption spectrum would look like with a 1300 ppm helium absorption signature. The absorption profile was taken from the 2.5D hydrodynamic model with the weakest absorption (solid green curve in Figure 15). This figure shows that the reduction process subtracts out at most 20% of the signal and that even the 2.5D hydrodynamic model with the weakest absorption is ruled out by our data.

4.5. Constraints on Absorption

To quantify the constraint on excess absorption that our data provide, we used the nested sampling code *dynesty* (Speagle 2020). We modeled the data as a scaled and vertically shifted version of the 2.5D hydrodynamic model with the weakest absorption (solid green curve in Figure 15, corresponding to an H-dominated atmosphere with 10^{-10} envelope fraction and 1 yr dispersal timescale). The observed excess absorption spectrum (Figure 6) is truncated to 10831–10835 Å to avoid the need to model any unsubtracted low-frequency variations in the data. Instead, the low-frequency variations are accounted for by vertically shifting the model to match the data.




Figure 5. Top: combined residuals image for both nights, showing excess absorption (ppm) in the stellar rest frame after all pipeline steps except shifting into the planetary rest frame. The top and bottom horizontal black lines mark the beginning and end of transit, respectively. The vertical red lines mark the helium triplet positions. No helium absorption is evident. Bottom: excess absorption during transit ($1 - F_{in}/F_{out}$) computed from the final residuals image for both nights.




Figure 6. Top: excess absorption during the transit in the planet frame as a function of wavelength, averaged over both nights. Helium absorption from the planet would result in a positive signal at the position of the red vertical lines. Due to the nature of ground-based high-resolution spectra, the white-light transit cannot be detected. Bottom: same as above, except an absorption signal was injected into the spectra before any reduction was undertaken. The injected signal has a maximum excess absorption of 1300 ppm, corresponding to the 2.5D hydrodynamic model with the weakest absorption (dashed blue curve in Figure 15).

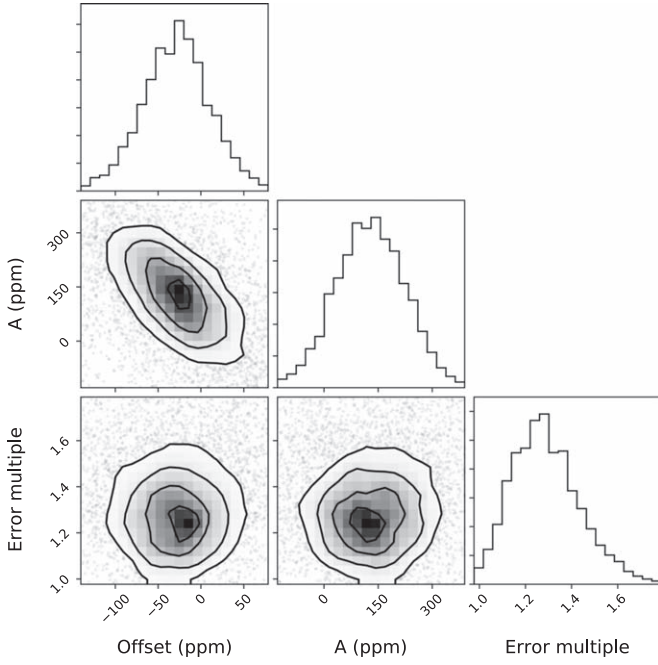


Figure 7. 2D posteriors from our fit to the data, constraining the maximum excess absorption (A). A is 1.3σ from 0.

We compute a covariance matrix for the excess absorption spectrum using the errors from the optimal extraction algorithm propagated forward through the pipeline. This matrix accounts for the covariances caused by the two linear interpolations—first onto a common wavelength grid, and second into the planetary rest frame. To account for unmodeled systematics, we multiply the covariance matrix by a free parameter e^2 , where $e > 1$. The log likelihood is given by

$$\ln L = -\frac{1}{2}r^TK^{-1}r - \frac{1}{2}|K| - \frac{N}{2}\ln 2\pi, \quad (10)$$

where K is the covariance matrix and r are the residuals.

Figure 7 shows the results of our nested sampling run. The free parameters are the vertical offset of the model, the maximum excess absorption (A), and the error multiple e whose square multiplies the covariance matrix. Our data are consistent with zero excess absorption, and the best-fit error inflation parameter e indicates that our errors are likely underestimated by approximately 25%. The preferred vertical offset is consistent with zero, as expected from visual inspection. Our fit places a 90% upper limit on A of 250 ppm, corresponding to an equivalent width of 0.27 mÅ. To see whether A is significantly nonzero, we performed a nested sampling run with the amplitude removed as a free parameter. The resulting log Bayesian evidence $\log(Z)$ was 1.3 higher in this run, indicating that the data prefer a model with no absorption.

The constraint on A can also be expressed in scale heights, where the scale height $H_{\text{eq}} = \frac{k_B T_{\text{eq}}}{\mu g}$ is computed assuming an equilibrium temperature of 2000 K and a hydrogen-dominated atmosphere. This metric was first proposed by Nortmann et al. (2018). We obtain an upper limit on the increase in apparent planet radius at 1083 nm: $\delta_{Rp}/H_{\text{eq}} < 11$ with $H_{\text{eq}} = 350$ km, which we put into context in Figure 8 by plotting the EUV flux and $\delta_{Rp}/H_{\text{eq}}$ of all planets with helium detections or nondetections. The EUV flux experienced by 55 Cnc e is obtained from Salz et al. (2016), who calculate $F_{\text{EUV}} = 7.4$




Figure 8. 55 Cnc e in context, with the size of each point inversely proportional to the planetary radius. The y-axis shows $\delta_{Rp}/H_{\text{eq}}$, the ratio of the increase in apparent radius at 1083 nm to the scale height at equilibrium temperature. The only super-Earth with a helium measurement or nondetection, 55 Cnc e is highly irradiated and has a very tight observational constraint compared to its peers. Data for other planets were taken from Kasper et al. (2020). We did not include the helium nondetections of AU Mic b (Hirano et al. 2020), WASP-52b (Vissapragada et al. 2020), K2-100b (Gaidos et al. 2020a), or K2-25b (Gaidos et al. 2020b), which are less sensitive (the upper limits on $\delta_{Rp}/H_{\text{eq}}$ are all greater than 100).

W m^{-2} by applying the scaling relation of Linsky et al. (2014) to the Ly α luminosity. Following Kasper et al. (2020), we conservatively adopt $3\times$ error bars on the EUV flux on either end, as different EUV reconstruction techniques give results that are discrepant by an order of magnitude in some cases (i.e., Salz et al. 2015b; Drake et al. 2020). Figure 8 shows that, relative to other planets with helium measurements, 55 Cnc e is a highly EUV-irradiated planet with a tight upper limit that is far below the $\delta_{Rp}/H_{\text{eq}}$ of successful detections for other planets.

5. Modeling

We next turn to helium outflow models to explore what our nondetection of helium absorption might mean for 55 Cnc e's atmospheric composition and corresponding mass-loss rate. We start by performing order-of-magnitude estimates of the mass-loss rate from analytical formulae. Next, we use three independent models to interpret the observational result. First, an isothermal Parker wind model (Oklopčić & Hirata 2018) puts constraints on the 2D parameter space of temperature and mass-loss rate but does not constrain either parameter independently. Our second model, The PLUTO-CLOUDY Interface (TPCI), can model the outflow in a 1D fashion given the stellar XUV spectrum and predict the mass-loss rate, temperature profile, and absorption spectrum. Our third and most sophisticated model is a 2.5D model (Wang & Dai 2018), which combines ray-tracing radiative transfer, real-time nonequilibrium thermochemistry, and hydrodynamics to model the outflow, assuming that it is symmetric about the star-planet axis. This model can also predict the mass-loss rate, temperature profile, and absorption spectrum, given an assumed atmospheric composition and envelope fraction.

We compare the TPCI and 2.5D predictions directly to observations. Since the Parker wind model requires both the mass-loss rate and the exosphere temperature as input parameters, we take a typical exosphere temperature from the TPCI model as a reasonable estimate and use it within the

Table 2
Summary of Model Predictions

Model	Dimensions	\dot{M} (g s ⁻¹)	Peak Absorption (ppm)
Energy-limited	0	$>2 \times 10^9$	N/A
Semiempirical	0	$8 \times 10^9 - 1 \times 10^{11}$	N/A
Parker wind	1	N/A	900–1800 ^a
TPCI	1	$0.75 - 1.1 \times 10^{10}$	800–1000
2.5D	2.5	$0.70 - 1.5 \times 10^{10}$	1400–2400

Note.

^a Assumes the mass-loss rate and exosphere temperature predicted by TPCI.

framework of the Parker wind model to constrain the mass-loss rate.

For convenience, we compile the takeaway results of the models we considered in Table 2. In many of these models, we consider a range of possible parameters and atmospheric compositions. The range of values we give in the table reflects the range of parameters and compositions.

5.1. Analytic Estimates of Escape Rate

Before running complex hydrodynamic simulations, it is useful to perform rough analytical estimates of the escape rate to illustrate the typical magnitudes involved and their dependence on stellar and planetary quantities. The simplest way to estimate the escape rate is to assume that it is limited by the X-ray and extreme-UV stellar radiation hitting the planet. The energy per unit mass required to escape the planet’s gravity well is $dE/dm = -GM_p/R_{XUV}$, while the rate at which XUV radiation impinges on the planet is $dE/dt = \frac{L_{XUV}}{4\pi a^2} \pi R_p^2$. Assuming a fraction η of the energy goes toward driving mass loss, we obtain

$$\frac{dm}{dt} = \frac{\eta R_p^3 L_{XUV}}{4 GM_p a^2}. \quad (11)$$

R_p is the radius of the XUV photosphere, but we take it to be the optical transit radius in order to obtain a conservatively low mass-loss rate. We obtain $L_{XUV} = 10^{27.70}$ erg s⁻¹ from Salz et al. (2016), who in turn obtain it by applying the scaling law in Linsky et al. (2014) to the star’s Ly α luminosity. We assume a low efficiency $\eta = 0.15$, in line with Monte Carlo heating models (Shematovich et al. 2014; Ionov & Shematovich 2015), and obtain 2×10^9 g s⁻¹. Because the values of η and L_{XUV} are both highly uncertain, this should be regarded as an order-of-magnitude estimate only.

A slightly more sophisticated approach to estimating the mass-loss rate is to adopt the semiempirical expressions of Wang & Dai (2018), who derived their expression from 2.5D numerical simulations of mass loss. Their equations are

$$\dot{M} = 4.5 \times 10^{-10} M_{\oplus} \text{ yr}^{-1} \cdot \left(\frac{R_{\text{EUV}}}{5R_{\oplus}} \right)^2 \times \max(F', F'^{0.6}) \min(1, M'^{-0.5}) \quad (12)$$

$$M' = \frac{M_c}{10M_{\oplus}} \quad (13)$$

$$F' = \frac{F_{\text{EUV}}}{5 \times 10^{14} \text{ erg s}^{-1} \text{ \AA}^{-1}}. \quad (14)$$

For 55 Cnc, the X-ray luminosity is only 10% of the total XUV luminosity, so the terms EUV and XUV are more or less

interchangeable. We use three methods to determine R_{EUV} , the apparent radius of the planet in the EUV. First, we use the radius implied by the optical transit depth: $1.9 R_{\oplus}$. This radius leads to a mass-loss rate of 8×10^9 g s⁻¹. Second, we use the Hill radius, $R_{\text{Hill}} = a \left(\frac{M_p}{3M_c} \right)^{1/3} = 7.5 R_{\oplus}$. This leads to a mass-loss rate of 1×10^{11} g s⁻¹. These two mass-loss rates are the lower and upper limits of what is reasonable under the semiempirical framework of Wang & Dai (2018).

The third method for estimating R_{EUV} is to follow Wang & Dai (2018) in assuming that the EUV photosphere is at $\rho = 10^{-13}$ g cm⁻³. To find the radius that corresponds to this density, we assume that the atmosphere is isothermal, and we need to find one point with known r and ρ . Wang & Dai (2018) (following Owen & Jackson 2012) pick the radiative–convective boundary (RCB) and use a parameterized opacity to calculate the position of the RCB. Unfortunately, this calculation involves many unknown quantities, such as the Kelvin–Helmholtz timescale and the envelope fraction. It may also be particularly ill-suited to thin atmospheres, where the radiative–convective boundary may be very close to the surface and at a much higher temperature than T_{eq} . We therefore assume that $P = 100$ mbar corresponds to the radius inferred from the optical transit depth, $1.9 R_{\oplus}$. The value of 100 mbar is the approximate photosphere of transit observations in the optical for solar-metallicity planets with thick atmospheres, making this guess more accurate than trying to estimate the RCB radius and pressure.

Assuming $r(P = 100 \text{ mbar}) = 1.9 R_{\oplus}$, we obtain $\rho_{\text{phot}} = \frac{P\mu}{k_B T_{\text{eq}}} = 1.4 \times 10^{-6}$ g cm⁻³ and $\beta \equiv \frac{GM_c \mu}{R_p k_B T_{\text{eq}}} = 37$, and we can calculate R_{EUV} :

$$R_{\text{EUV}} = \frac{R_{\text{phot}}}{1 + \beta^{-1} \ln(\rho_{\text{EUV}}/\rho_{\text{phot}})} \quad (15)$$

$$= 3.4 R_{\oplus}. \quad (16)$$

This implies $\dot{M} = 3 \times 10^{10}$ g s⁻¹.

5.2. 1D Parker Wind Model

Now that we have a rough sense for the expected mass-loss rate for a hydrogen- and helium-rich atmosphere, we can ask whether mass-loss rates in this range are definitively ruled out by our helium nondetection. We use the Parker wind model of Oklopčić & Hirata (2018) to translate our helium nondetection into a joint constraint on the mass-loss rate and the exosphere temperature. This study treats the outflow as an isothermal, radially symmetric Parker wind from a hydrogen-dominated atmosphere. The model calculates the population levels of singlet and triplet states by balancing recombination, photo-ionization, collisional (de-)excitation, and radiative decay as a function of altitude. Using the helium number density and the population level of the triplet state as a function of altitude, it then calculates the transit depth as a function of wavelength.

We simulate absorption spectra for a 2D grid of mass-loss rates ranging from 10^7 to 10^{11} g s⁻¹ and exospheric temperatures ranging from 2500 to 9000 K, in order to better quantify the limits our data place on the planet’s present-day mass-loss rate. For each combination of mass-loss rate and exosphere temperature, we compute an excess absorption spectrum and calculate the corresponding log likelihood of the observational data given the model. We then compute $\Delta \ln(L)$:




Figure 9. Confidence with which each combination of mass-loss rate and exosphere temperature is ruled out, according to the 1D Parker wind model of Oklopčić & Hirata (2018). The red cross indicates the mass-loss rate and approximate exosphere temperature predicted by the nonisothermal Salz et al. (2016) model.

the difference between the $\ln(L)$ of the model and that of a zero-absorption model.

We show the resulting $\Delta \ln(L)(\dot{M}, T_0)$ in Figure 9. In these models higher-temperature exospheres have faster outflow velocities and correspondingly weaker helium absorption for a given mass-loss rate. At the planetary surface, for example, a 5000 K exosphere with a mass-loss rate of $1.4 \times 10^{10} \text{ g s}^{-1}$ requires a 0.16 km s^{-1} wind; a 10,000 K exosphere with the same mass-loss rate requires a 3 km s^{-1} wind. A faster wind implies lower density at the same mass-loss rate, decreasing helium absorption. This is both because there are fewer helium atoms per cubic volume and because the fraction of helium atoms in the triplet state is lower. The latter, in turn, is because the triplet state is primarily populated by recombination (see Figure 10), and a lower density means a lower recombination rate. For these reasons, the mass-loss constraint is less stringent at higher temperatures.

The Parker wind model does not provide a way to estimate the exosphere temperature, which makes it impossible to constrain the mass-loss rate without the help of another model. The TPCI model predicts a peak exosphere temperature of 5000–6000 K (Salz et al. 2016; see next subsection for details). At $T_0 = 5000 \text{ K}$, the mass-loss rate is constrained to $\dot{M} < 10^{9.1} \text{ g s}^{-1}$ ($\Delta \ln(L) < 7$), or $\dot{M} < 10^{8.8} \text{ g s}^{-1}$ ($\Delta \ln(L) < 3$). At $T_0 = 6000 \text{ K}$, the mass-loss rate is constrained to $\dot{M} < 10^{9.5} \text{ g s}^{-1}$ ($\Delta \ln(L) < 7$), or to $\dot{M} < 10^{9.2} \text{ g s}^{-1}$ ($\Delta \ln(L) < 3$), where $\Delta \ln(L)$ thresholds of 7 and 3 correspond to likelihood ratios of 1100 and 20, respectively. In summary, the Parker wind model gives an upper limit on the mass-loss rate of $\sim 10^9 \text{ g s}^{-1}$ for exospheric temperatures predicted by TPCI.

Lastly, we use this model to gain intuition on the physical processes that determine the triplet helium fraction and therefore the helium absorption strength. Figure 10 plots, as a function of radius, the production and destruction rates of triplet helium due to the processes considered by Oklopčić & Hirata (2018): recombination, radiative decay, ionization, collisional excitation/de-excitation with electrons, and collisional de-excitation with neutral hydrogen atoms. Recombination is the most important production mechanism, while destruction is due to a combination of collisional de-excitation

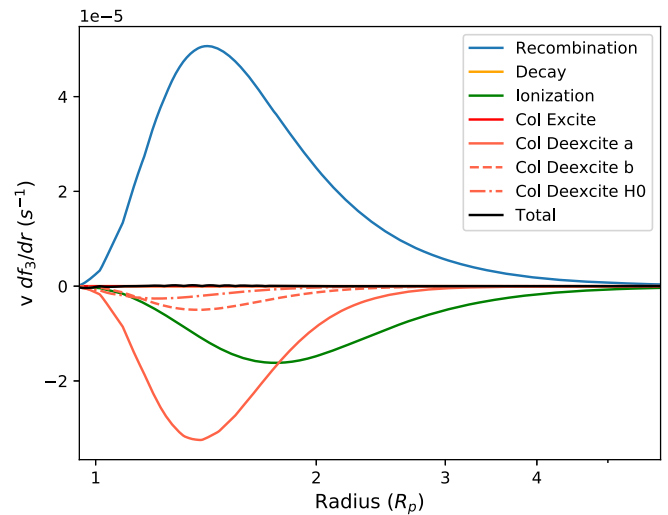


Figure 10. Production and destruction rates of triplet helium due to various processes, from a Parker wind model of 55 Cnc e with a mass-loss rate of $1.38 \times 10^{10} \text{ g s}^{-1}$ and exosphere temperature of 5000 K.

with electrons and ionization: the former is dominant close to the planet, while the latter is dominant far from the planet, as one would expect. Collisional excitation and radiative decay are negligible. The production and destruction rates very nearly cancel out, indicating that the triplet helium fraction is mostly in equilibrium and that advection is not significant.

5.3. 1D PLUTO-CLOUDY Hydrodynamic Model

In this section we use a 1D, spherically symmetric radiative-hydrodynamical simulation to predict the exospheric temperature structure, mass-loss rate, and metastable helium absorption signal for 55 Cnc e under several different scenarios. Salz et al. (2016) previously used this model to simulate hydrogen-rich exospheres for several planets, including 55 Cnc e, but did not calculate the predicted absorption signal from metastable helium. The Salz et al. (2016) model predicts an exospheric temperature of 3000–6000 K and an outflow velocity at large distances of 10–15 km s^{-1} . It predicts a mass-loss rate of $1.4 \times 10^{10} \text{ g s}^{-1}$, corresponding to 0.9% of the planetary mass per Gyr. As the paper shows, the predicted Ly α signal is consistent with the nondetection by Ehrenreich et al. (2012). This mass-loss rate would imply that 55 Cnc e started life as a sub-Neptune with a hydrogen/helium envelope of $\sim 10\%$ by mass. We use TPCI to reconstruct the Salz et al. (2016) model and calculate the corresponding helium absorption signal. This also allows us to explore other compositions, including a helium-dominated atmosphere.

This model couples PLUTO, a hydrodynamics code that can work in one, two, or three dimensions (Mignone et al. 2007), and CLOUDY (Ferland et al. 2013), a 1D plasma simulation and spectral synthesis code. PLUTO and CLOUDY are linked in TPCI (Salz et al. 2015a). In TPCI, CLOUDY calculates the equilibrium chemistry and ionization state of the medium given a radiation field, computes the heating and cooling rates of the new state, and feeds this information to PLUTO. PLUTO heats or cools the medium appropriately, evolves the medium hydrodynamically, and provides the new state to CLOUDY, after which the cycle restarts. Both PLUTO and CLOUDY are sophisticated, publicly available, and general-purpose codes that have been applied to a variety of astrophysical problems, ranging from exoplanet mass loss to the magnetic fields of




Figure 11. Profiles of various physical quantities for the hydrogen- and helium-dominated TPCI models in addition to the 2.5D model. The low-amplitude oscillations in the helium-dominated profiles are numerical artifacts resulting from the nonzero advection length. In green are the results from Salz et al. (2016), who compute all quantities plotted here except for the helium triplet density.

neutron stars (PLUTO; e.g., Sur et al. 2020) to high-redshift gamma-ray bursts (CLOUDY; e.g., Shaw & Ferland 2020). Salz et al. (2016) adopt a 1D spherically symmetric model and do not include the planetary magnetic field. CLOUDY includes the 30 lightest elements, from hydrogen to zinc, and accounts for many physical processes, including radiative and collisional ionization/recombination, inner shell ionization, and charge exchange. However, Salz et al. (2016) assume a purely atomic hydrogen and helium atmosphere.

As input to the PLUTO-CLOUDY model, the authors use the X-ray luminosity measured by XMM-Newton in 2009 April, namely, $4.6 \times 10^{26} \text{ erg s}^{-1}$ between 5 and 100 Å (Sanz-Forcada et al. 2011). This is 2.4 times the flux measured by Chandra, which was obtained simultaneously with the HST Ly α observations in 2012 March/April (Ehrenreich et al. 2012). This is likely because 2012 March/April was right at the minimum of the 10.5 ± 0.5 yr stellar activity cycle (Bourrier et al. 2018), while 2009 April was 3 yr from minimum and 2.3 yr from maximum. Because our 2019 December observations were almost exactly one stellar cycle after 2009 April, we expect the star’s X-ray properties to be similar to those observed by Sanz-Forcada et al. (2011) and adopted by Salz et al. (2018), and more favorable for mass-loss measurements than the conditions encountered by Ehrenreich et al. (2012).

Because we do not know exactly what settings or what version of TPCI they used, we are unable to replicate their results exactly, but we match their temperature profile to better than 500 K and their mass-loss rate to within 25% accuracy (see Figure 11). This is far smaller than the factor of several uncertainty they report as the inherent model error (their Table 2), resulting from factors such as the neglected 3D structure ($4\times$ uncertainty), uncertain irradiation strength ($3\times$), and neglect of magnetic fields ($2\times$). Following Salz et al., we neglect molecules and elements other than hydrogen and helium in our simulation.

We find that the advection length—roughly speaking, the resolution of CLOUDY’s advection calculations—is a crucial parameter for these models. Smaller values lead to more accurate results but take longer to converge. Large values lead to spurious spatial oscillations in the temperature and ionization state. We adopt an advection length of 0.15 planetary radii, which we find is small enough that the oscillations in temperature are of order 0.2% for the helium-dominated run and 0.06% for the hydrogen-dominated run. We run the simulations for 150 days of model time (roughly 1000 sound-crossing times) while neglecting advection. We then turn on advection, which slows down the run by a factor of ~ 80 , and let the simulation run for another 150 days of model time. We monitor the evolution of the temperature, density, velocity, and

mass-loss rate (calculated as $4\pi r^2 \rho(r)v(r)$) profiles to verify that the simulation has in fact converged, with temperatures fluctuating by less than 50 K and the mass-loss rate by less than 0.1%.

After reproducing the hydrogen-dominated model from Salz et al. (2016), we run a second TPCI simulation for a helium-dominated atmosphere with 99% He and 1% H by number to explore a scenario where slow mass loss over many gigayears fractionated the atmosphere.

Since CLOUDY computes the level populations of every species, we configure it to report the number density of helium atoms in the triplet state at every radial coordinate. We use this number density profile, in addition to the temperature and velocity profiles, to compute the excess absorption spectrum.

Figure 11 compares the profiles for temperature, velocity, hydrogen density, helium triplet density, and neutral fraction for the hydrogen- and helium-dominated TPCI models. Using these profiles, we computed the mass-loss rate as $\dot{M} = 4\pi r^2 \rho v / 4$. The division by 4 follows Salz et al. (2016) and is meant to account for the 3D nature of the outflow, as we simulate only the substellar point. We obtain a mass-loss rate of $1.1 \times 10^{10} \text{ g s}^{-1}$ for the hydrogen-dominated scenario and $7.5 \times 10^9 \text{ g s}^{-1}$ for the helium-dominated scenario.

In Figure 12, we compare the predicted excess absorption from our TPCI models to the observed excess absorption spectrum. Taking into account the interpolation-induced covariance between the data points, we find that a zero-absorption model is preferred over the TPCI hydrogen-dominated model by $\Delta \ln(L) = 126$ and preferred over the helium-dominated model by $\Delta \ln(L) = 100$. Although these estimates do not account for the error due to systematics or variable tellurics, it can be seen visually that the predicted absorption for both models is ruled out.

Critical to all photoevaporation models is the stellar EUV luminosity, which cannot be directly observed and must be reconstructed from X-ray and/or Ly α flux. However, different EUV reconstruction techniques give results that are discrepant by an order of magnitude in the worst-case scenario (i.e., Salz et al. 2015b; Drake et al. 2020), although inactive main-sequence stars like 55 Cnc e can be modeled more accurately. To determine the effect that uncertainties in the EUV luminosity have on the helium absorption signal, we ran TPCI simulations for the hydrogen-dominated scenario at twice and half the nominal stellar flux. The profiles of various physical quantities with respect to radius are shown in Figure 13 for the low-flux, nominal, and high-flux scenarios, while the helium absorption signals are compared in Figure 12.

The high-flux scenario has a mass-loss rate 3 times higher than the low-flux scenario, implying a slightly sublinear scaling of mass-loss rate with respect to flux. Doubling the stellar flux substantially increases the exosphere temperature and outflow velocity while substantially decreasing the neutral fraction of both hydrogen and helium. The hydrogen number density is moderately higher at higher flux, but the triplet helium number density is remarkably insensitive to flux. This insensitivity means that the excess absorption is barely affected by uncertainties in the stellar flux—it is only 20% lower in the high-flux scenario than in the low-flux scenario.

5.4. 2.5D Hydrodynamic Simulations

Our final and most sophisticated model utilizes the approach outlined in Wang & Dai (2018), which combines ray-tracing radiative transfer, real-time nonequilibrium thermochemistry,

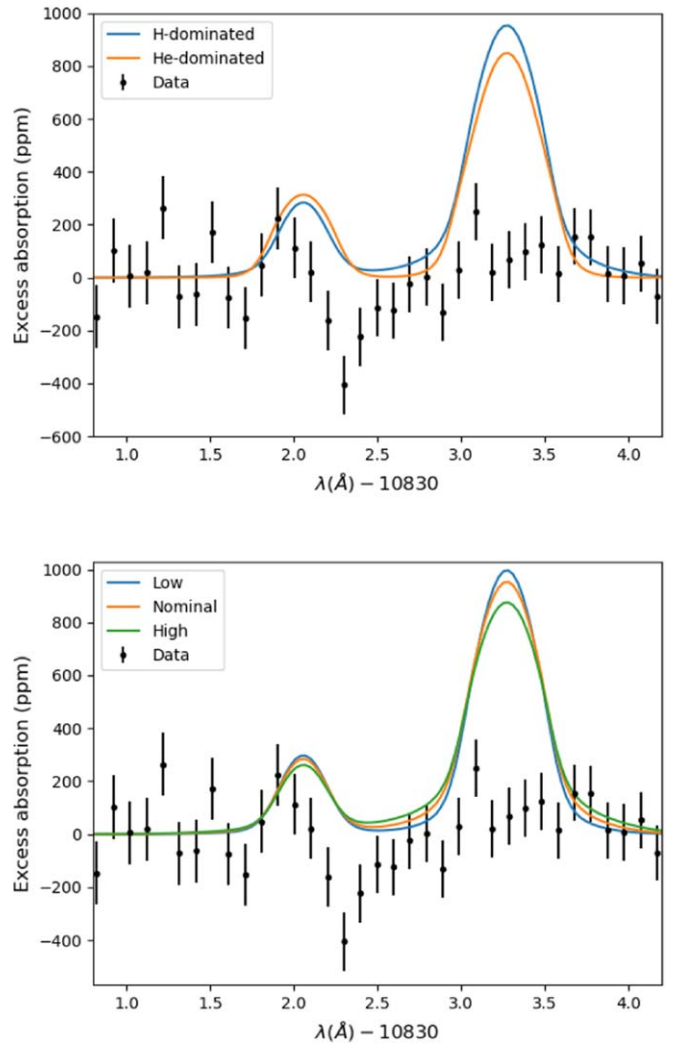


Figure 12. Top: TPCI hydrogen-rich (90% H, 10% He) and helium-rich (1% H, 99% He) models are both ruled out by our observational data. Bottom: increasing or decreasing the stellar XUV flux by a factor of 2 does not significantly change the absorption signal in the TPCI hydrogen-rich model.

and hydrodynamics based on the higher-order Godunov method code Athena++ (Stone et al. 2020). Photoevaporation is inherently not a spherically symmetric phenomenon, since it is driven by stellar XUV flux and the star is only in one direction. Compared to 1D LTE models, this axisymmetric (2.5D) model better captures the anisotropy of the outflow pattern, while the non-LTE thermochemistry self-consistently predicts the mass-loss rate and the line profiles. Like in Oklopčić & Hirata (2018) and Wang & Dai (2020), metastable helium is added as a chemical species, and key reactions that form and destroy this species are included in the thermochemical network. The model incorporates a total of 26 species and 135 reactions, including various relevant heating and cooling processes (e.g., photoionization, photodissociation of atomic hydrogen, and Ly α cooling).

Our simulations are done in a spherical coordinate system centered on the planet, with the polar axis pointing from the center of the planet to the host star. The simulation domain is in the radial and polar directions $[r, \theta]$, while symmetry is assumed in the ϕ direction. Photons for the ray-tracing calculation are divided into five energy bins: (1) $h\nu = 2 \text{ eV}$ for infrared, optical, and near-ultraviolet (NUV) photons; (2)




Figure 13. Profiles of various physical quantities for hydrogen-dominated TPCI models of different stellar XUV flux: nominal, $2\times$ nominal (high), and half nominal (low).

$h\nu = 7$ eV for “soft” far-ultraviolet (FUV) photons; (3) $h\nu = 12$ eV for the Lyman–Werner band FUV photons that can photodissociate molecular hydrogen but cannot ionize them; (4) $h\nu = 20$ eV for “soft” extreme-ultraviolet (soft EUV) photons that can ionize hydrogen but *not* helium; (5) $h\nu = 40$ eV for hard EUV (and soft X-ray; denoted by “hard EUV” hereafter for simplicity) photons that ionize hydrogen *and* helium. Photon fluxes in each energy bin are determined according to the planet’s orbital separation and the corresponding luminosities for a typical K star, according to the reviews in Oklopčić (2019). The EUV flux for this model star is $10,926 \text{ erg s}^{-1} \text{ \AA}^{-1}$, 30% higher than what we adopted in the 1D TPCI simulations—a negligible difference given the inherent uncertainty in EUV flux. In addition to the opacities caused by photochemical reactions, we also added an effective opacity term to all bands, particularly in the optical band $h\nu = 2$ eV, where our opacity calculation did not include the Thomson cross section $\sigma/H \simeq 6.7 \times 10^{-25} \text{ cm}^2$.

We note that the models we present here do not account for the effect of the stellar wind on the predicted mass-loss rate. In previously published 1D models (i.e., Murray-Clay et al. 2009), the stellar wind acts as a simple scalar suppression force. In Wang & Dai (2020), they discuss the effect of the stellar wind on planetary outflows using the same hydrodynamic model presented here. They find that in higher dimensions the stellar wind only suppresses the dayside mass loss; the outflow is

simply redirected toward the night side of the planet, forming a comet-like tail. In these models, adding a solar-like stellar wind (we note that 55 Cnc e is an old, Sun-like star) only changes the 3D mass-loss rate by a few percent, corresponding to a change in the equivalent width of He absorption of less than 10% (see Table 2 in Wang & Dai 2020). In short, stellar wind may alter the line profile of He absorption, but it is unable to quench the outflow in three dimensions and is therefore unlikely to be the reason behind the nondetection of He absorption for 55 Cnc e.

A typical planetary atmosphere consists of a convective interior and a quasi-isothermal exterior (e.g., Rafikov 2006), but our numerical tests found that even a thin convective isentropic layer will cause the whole atmosphere to become overdense and unbounded for such a close-in low-mass planet (the boil-off regime; see discussions in Owen & Wu 2016 and Wang & Dai 2018). We therefore set up the model atmospheres with a quasi-isothermal layer directly above the rocky core. In contrast, the TPCI model did not set a solid surface or impose a truncation of the gas reservoir; instead, it sets pressure and density boundary conditions at the inner radius. We summarize the key quantities that define the 2.5D fiducial model in Table 3. In addition to the fiducial model, which has a density of $\rho(R_p) = 10^{-7} \text{ g cm}^{-3}$ at the rocky surface, we also considered models with $\rho(R_p) = 10^{-5}$ and $10^{-9} \text{ g cm}^{-3}$, respectively. As with the TPCI models, we also considered a scenario with an He-dominated atmosphere (1% hydrogen and




Figure 14. Density, temperature, and triplet helium density from the fiducial 2.5D model. The star is toward the left, and the simulation volume is represented by revolving the semicircles about their straight edge. The white lines are streamlines, while the dashed black lines represent the sonic surface.

Table 3

Setups of the Fiducial 2.5D Numerical Models of the Evaporating Atmosphere

Item	Value
Simulation domain	
Radial range	$1.89 \leq (r/R_{\oplus}) \leq 20$
Latitudinal range	$0 \leq \theta \leq \pi$
Resolution ($N_{\log r} \times N_{\theta}$)	144×96
Atmospheric properties ^a	
$\rho(R_p)$	$10^{-7} \text{ g cm}^{-3}$
$T(R_p) \simeq T_{\text{eq}}$	1990 K
Radiation flux (photons $\text{cm}^{-2} \text{ s}^{-1}$)	
2 eV (IR/optical)	1.1×10^{21}
7 eV (soft FUV)	4.3×10^{16}
12 eV (Lyman–Werner FUV)	1.5×10^{13}
20 eV (soft EUV)	1.2×10^{13}
40 eV (hard EUV)	1.6×10^{14}
Initial abundances [n_X/n_{atom}] ^b	
H ₂	0.455
He	0.091
H ₂ O	1.8×10^{-4}
CO	1.4×10^{-4}
S	2.8×10^{-5}
Si	1.7×10^{-6}

Notes.

^a $R_p \simeq 1.89 R_{\oplus}$ is the size of the rocky planet core.

^b Dust grains are not included, since $T_{\text{eq}} \simeq 1990 \text{ K}$ is well above the dust sublimation temperature.

99% helium by atom number) spanning the same three surface pressures.

Figure 14 shows the 2D profiles of density, temperature, and triplet helium density, along with streamlines and the sonic surface. As can be seen, the antistellar side of the planet has drastically different physical conditions from the star-facing side. However, due to its compactness, much of the colder, less dense material on the antistellar side would not block any more light during transit than the planet itself. Except for this region, the rest of the simulation domain is largely spherically symmetric, especially in the number density of triplet helium, which directly determines the helium absorption signature.

Figure 11 shows the radial profiles of temperature, velocity, hydrogen number density, triplet helium number density, neutral H fraction, and neutral He fraction for the fiducial model (hydrogen dominated, $\rho(R_p) = 10^{-7} \text{ g cm}^{-3}$), showing a fairly typical supersonic photoevaporative outflow that carries




Figure 15. Metastable helium transmission spectra of model atmospheres for the simulations described in Section 5.4 and Tables 3, 4, presenting the time-averaged (from the end of ingress through the start of egress) excess absorption. The horizontal dashed line indicates zero excess absorption for reference. Different reference densities (measured at the planet surface R_p) are marked by different line styles, with $\log(\rho_0) = -5, -7, -9$ corresponding to roughly 500, 5, and 0.05 mbar, respectively. Orange lines represent helium-dominated models, while blue lines represent hydrogen-dominated models.

metastable helium atoms. The assumed EUV flux of 55 Cnc e produces a relative abundance of metastable helium in 55 Cnc e's atmosphere of $\sim 10^{-7}$, nearly identical to the abundance predicted by TPCI. Despite this low abundance, the transmission spectra in Figure 15 still have clearly recognizable excess absorption with amplitudes of a few thousand parts per million. The absorption is still greater than 1000 ppm even for the thinnest atmosphere ($\rho(R_p) = 10^{-9} \text{ g cm}^{-3}$). The mass-loss timescales of these models, summarized in Table 4, range between years and thousands of years. The mass-loss rates are within 50% of those predicted by TPCI for both hydrogen- and helium-dominated atmospheres. These mass-loss rates are time-averaged after the simulation reaches a quasi-steady state after dozens of dynamical timescales elapsed since the start of the simulations. The nondetection of helium absorption is therefore in good agreement with the short dispersal timescales for these atmospheres, implying that 55 Cnc e would not have been able to retain a primordial atmosphere for very long after the dispersal of the gas disk.

Table 4
2.5D Simulations: Parameters and Results

Model No.	$\rho(R_p)$ (g cm^{-3})	M_{env}/M_p	$P(R_p)$ (mbar)	\dot{M} (10^{10} g s^{-1})	$\tau_{\text{disp}}^{\text{a}}$ (yr)	$\langle W_{\lambda} \rangle^{\text{b}}$ (10^{-3} \AA)	$\Delta \ln(L)^{\text{c}}$
H/He Envelope							
1	10^{-5}	0.98×10^{-7}	500	1.5	5000	2.4	748
2 ^d	10^{-7}	0.98×10^{-9}	5	0.84	86	1.4	243
3	10^{-9}	0.98×10^{-11}	0.05	0.85	0.9	1.4	252
He-dominated							
4	10^{-5}	0.80×10^{-7}	400	1.1	4000	3.3	712
5	10^{-7}	0.80×10^{-9}	4	0.94	65	2.0	538
6	10^{-9}	0.80×10^{-11}	0.04	0.70	0.9	1.6	342

Notes.

^a Dispersal timescale, defined by the current atmospheric mass divided by the steady-state mass-loss rate.

^b Dimensional equivalent width of excess absorption, time-averaged from the end of ingress through the start of egress.

^c Difference in log likelihood between this model and a model with zero absorption. A bigger number means a worse fit.

^d Fiducial model (Section 5.4).

6. Discussion

Now that we have explored the limits that can be placed on 55 Cnc e’s atmosphere from our helium observations alone, we consider how these results relate to other published observations of 55 Cnc e’s atmosphere. We first focus on Ehrenreich et al. (2012), which presents an Ly α transit of 55 Cnc e from HST/STIS. Ly α and the helium 1083 nm line are both good probes of the outflowing atmosphere, but the former is sensitive only to the high-speed portions of the outflow, while the latter can only probe the low-speed portions of the outflow where metastable helium exists. Ehrenreich et al. (2012) measure the transit depth between 1215.36 and 1215.67 \AA , obtaining a value of $0.3\% \pm 2.4\%$. From this, they constrain the mass-loss rate to be below $3 \times 10^8 \text{ g s}^{-1}$ (3σ). They interpret their data using the Bourrier & Lecavelier des Etangs (2013) model, which is a 3D particle-based simulation that takes into account radiation pressure, photoionization, the stellar wind, and the effects of self-shielding for both stellar photons and protons. The inclusion of outward forces from radiation pressure and stellar wind gives the outflow a trailing cometary tail: a consequence of the Coriolis acceleration $-2\Omega \times \mathbf{v}$. This model has five tunable parameters: the mass-loss rate, the stellar EUV flux, and three stellar wind parameters: bulk velocity, temperature, and density. (By contrast, our 2.5D model is a hydrodynamics code that predicts the mass-loss rate given the EUV flux but does not take the stellar wind into account and cannot model a tail.) As with all Ly α observations of even the nearest stars, the line core is completely absorbed by the interstellar medium, and only the far wings are visible. This means that the detectability of an absorption signal depends as much on the highly uncertain kinematic structure of the outflow as on the quantity of outflowing gas.

Even though we quote a slightly higher upper limit on the mass-loss rate ($\sim 10^9 \text{ g s}^{-1}$ from the isothermal Parker wind model) than Ehrenreich et al. (2012; $3 \times 10^8 \text{ g s}^{-1}$), our models are in fact consistent with their Ly α nondetection even at mass-loss rates that are definitively ruled out by our helium nondetection. Figure 12 of Salz et al. (2016) compares the Ly α data to their TPCI simulation of 55 Cnc e under the assumption of a hydrogen-dominated atmosphere, finding that the two are consistent. As we have seen, however, the helium absorption signal predicted by TPCI for the hydrogen-dominated atmosphere case is highly inconsistent with our

nondetection. In addition, a 1D model is not ideal for modeling Ly α because the strength of the line absorption, combined with interstellar absorption that eliminates the line core, makes the kinematic structure of the outflow crucially important in predicting the observable signal. For this reason, we used the same code we used to run our 2.5D models to run a helium-dominated 3D hydrodynamic model (resembling model 5 in Table 4, corresponding to an intermediate-mass helium atmosphere). We found that just 1% hydrogen in the outflow is enough to give a 12% excess absorption depth in the Ly α line core, but the line core is not observable owing to interstellar absorption.

To determine what is observable, we start with the stellar intrinsic Ly α profile, corrected for interstellar absorption, that Ehrenreich et al. (2012) provide in Figure 6. We model the observed out-of-transit spectrum by convolving the absorbed profile with the line-spread function of STIS, as provided by STScI.¹⁰ We model the observed in-transit spectrum by multiplying the post-ISM-absorption profile by the predicted Ly α absorption profile from the exosphere and convolving the product with the line-spread profile. The resulting excess absorption spectrum for our helium-dominated 2.5D model is shown in Figure 16. We did not show the hydrogen-dominated case as the high-opacity region extends well beyond the domain of our simulation. It is clear that the predicted Ly α absorption is fully consistent with the observations. As a caveat, we note that absorption is often seen in the high-velocity wings of Ly α observations, which could be due to charge exchange with the stellar wind (Holmström et al. 2008; Tremblin & Chiang 2013). Since our model does not include stellar wind, it may severely underestimate the absorption at observable wavelengths. Were the stellar wind to be included, it is possible that the Ly α nondetection would become a stronger constraint on mass loss than the helium nondetection.

As noted in the previous section, the observations by Ehrenreich et al. (2012) coincided with the stellar activity minimum. As far as we know, nobody observed 55 Cnc e in either Ly α or He I 1083 nm during the stellar maximum in mid-2017. Our observations took place in between said maximum and the next minimum in late 2022. The star’s X-ray luminosity is 2.4 times higher in this phase than it was in 2012, making it

¹⁰ <https://www.stsci.edu/hst/instrumentation/stis/performance/spectral-resolution>




Figure 16. $\text{Ly}\alpha$ observations of Ehrenreich et al. (2012), compared to predictions of our helium-dominated 2.5D model (model 5 in Table 4). Within the gray shaded region, the intrinsic stellar flux is fully absorbed by the ISM, but the observed flux is nonzero because of the broad instrumental line-spread profile.

conceivable that any potential outflow from the planet would have increased in strength relative to the epoch of the $\text{Ly}\alpha$ observations during the stellar minimum. Both the TPCI models and the 2.5D simulations assume the higher X-ray flux relevant for the epoch of our helium observations, and as a result they may overpredict the magnitude of the $\text{Ly}\alpha$ absorption at the epoch of the HST observations. Since the predicted $\text{Ly}\alpha$ absorption from these higher X-ray flux models is already undetectable, there is no need to recalculate the models with a lower X-ray flux level.

Prior to our observations, it was also conceivable that 55 Cnc e might have had a helium-dominated atmosphere. However, our observations show no evidence of helium absorption, and our 2.5D simulations show that the mass-loss rates are much too fast for the planet to keep either a hydrogen- or helium-dominated atmosphere. Fast mass loss also makes it difficult to create a helium-dominated atmosphere in the first place, because escape rates much faster than the diffusion-limited rate do not significantly fractionate the elements (Hu et al. 2015).

Our observations appear to contradict those of Tsiaras et al. (2016), who analyzed two transits of 55 Cnc e observed with HST/WFC3 and found an upward-sloping spectral feature consistent with HCN absorption in a lightweight atmosphere. They perform a free retrieval on the transmission spectrum with TauREx and find a mean molecular weight μ of 2–6 (their Figure 9). The atmospheric scale height is inversely proportional to μ ; thus, the amplitude of features in a transmission spectrum is also inversely proportional to μ . Values of $\mu = 2$ –6 are consistent with a hydrogen-dominated atmosphere ($\mu = 2.3$) or a helium-dominated atmosphere ($\mu \approx 4$) but exclude atmospheres dominated by heavier molecules (i.e., water with $\mu = 18$, N_2 with $\mu = 14$, O_2 with $\mu = 16$). If the atmosphere is not hydrogen or helium dominated, absorption features from any molecule, including HCN, would be undetectable. If $\mu = 16$, for example, the scale height would be 47 km (assuming equilibrium temperature) and the change in transit depth corresponding to one scale height would be 3 ppm. This lies far below the detection threshold of the observations reported in Tsiaras et al. (2016). We conclude that

our nondetection of escaping helium, together with the nondetection of escaping hydrogen by Ehrenreich et al. (2012) and the strong irradiation of the planet over its long life, makes it unlikely that the features detected by Tsiaras et al. in the HST/WFC3 spectrum are planetary in nature.

Two more pieces of evidence complicate the picture: the significant phase offset in the Spitzer 4.5 μm phase curve (Demory et al. 2016a), and the year-to-year variability in the secondary eclipse depth in this same band (Demory et al. 2016b). To explain the eclipse variability, one could invoke a magma world with no atmosphere, or an extremely tenuous mineral atmosphere in equilibrium with the molten surface (Ito et al. 2015), in which case ejecta from volcanic eruptions could periodically shroud the surface. The mass and radius of 55 Cnc e are consistent with a world without an atmosphere (Bourrier et al. 2018). However, the phase offset suggests a thick, high molecular weight atmosphere. This point was discussed in Angelo & Hu (2017), who suggested that an N_2 -dominated atmosphere would be consistent with the Spitzer data. The possible composition of a nitrogen-dominated atmosphere and the observability of spectral features are explored in Miguel (2019) and Zilinskas et al. (2020). However, such a thick atmosphere would be unlikely to change substantially on year-long timescales. Our nondetection of helium is agnostic to all potential high mean molecular weight atmospheres, as helium is not expected to be a significant component of secondary (e.g., nonprimordial) atmospheres. While photodissociation at the top of a water-rich atmosphere might create a detectable $\text{Ly}\alpha$ signal, high-resolution spectral observations rule out water-rich atmospheres (volume mixing ratio $> 0.1\%$) with a mean molecular weight of $\leq 15 \text{ g mol}^{-1}$ at a 3σ confidence level (Jindal et al. 2020).

7. Conclusion

We observed two transits of 55 Cnc e using Keck/NIRSPEC to look for metastable helium absorption in the 1083 nm line. We found no absorption greater than 250 ppm (90% upper limit) and used three independent models to interpret this result. First, an isothermal Parker wind model (Oklopčić & Hirata 2018) puts constraints on the temperature and mass-loss rate (Figure 9), with the mass-loss rate constrained to less than $\sim 10^9 \text{ g s}^{-1}$ for exosphere temperatures of 5000–6000 K. This exosphere temperature is obtained from our second model, TPCI, which can model the outflow in a 1D fashion given the stellar XUV spectrum. TPCI predicts a mass-loss rate of $1.1 \times 10^{10} \text{ g s}^{-1}$ for a hydrogen-dominated atmosphere and $7.5 \times 10^9 \text{ g s}^{-1}$ for a helium-dominated atmosphere, both of which result in absorption several times stronger than what is (not) observed (Figure 12). Our third and most sophisticated model is a 2.5D model (Wang & Dai 2018) that combines ray-tracing radiative transfer, real-time nonequilibrium thermochemistry, and hydrodynamics to model the outflow, assuming that it is symmetric about the star–planet axis. Even for extremely thin atmospheres with dispersal timescales of millennia or less, the model still predicts high mass-loss rates of $\sim 10^{10} \text{ g s}^{-1}$ for both hydrogen- and helium-dominated atmospheres (see Table 4), which result in 1500–2500 ppm excess absorption—many times higher than what is observed (Figure 15). Although the significant model uncertainties must be kept in mind, our observations provide strong evidence against the existence of a low mean molecular weight primordial atmosphere on 55 Cnc e.

If 55 Cnc e instead possesses a high mean molecular weight secondary atmosphere, detection via transit spectroscopy will be extremely challenging. Ultimately, we believe that emission spectroscopy with next-generation telescopes—JWST, TMT, ELT, and GMT—represents the best path forward. This planet’s high dayside temperature (~ 2700 K) makes it a particularly favorable target for emission spectroscopy, and the magnitude of spectral features seen in emission is independent of the mean molecular weight of the atmosphere. While the brightness of 55 Cnc makes it feasible to search for relatively small signals, systematics that are negligible for low-S/N targets become important when the photon noise is small. In this paper, we saw that cross talk, fringing, and the inseparability of the 2D PSF each required special handling; observations with both JWST and next-generation ground-based telescopes will likely encounter similar technical challenges. Despite these challenges, atmosphere modeling studies suggest that it is indeed possible to detect high mean molecular weight atmospheres for 55 Cnc e using next-generation telescopes (e.g., Zilinskas et al. 2020). To date, atmospheric absorption features have only been detected for planets with relatively massive, hydrogen-rich atmospheres. Detecting a high mean molecular weight atmosphere around a high-density planet like 55 Cnc e would provide invaluable insights into the nature and origin of the broader population of short-period super-Earths.

M.Z. would like to acknowledge Joe Ninan, who was indispensable in deciphering the mysteries of NIRSPEC systematics. He also acknowledges Jacob Bean for sending most of the data plotted in Figure 8.

Software: `numpy` (van der Walt et al. 2011), `scipy` (Virtanen et al. 2020), `matplotlib` (Hunter 2007), `dynesty` (Speagle 2020), `corner` (Foreman-Mackey 2016).

ORCID iDs

Michael Zhang  <https://orcid.org/0000-0002-0659-1783>
 Heather A. Knutson  <https://orcid.org/0000-0002-5375-4725>
 Lile Wang  <https://orcid.org/0000-0002-6540-7042>
 Fei Dai  <https://orcid.org/0000-0002-8958-0683>
 Antonija Oklopčić  <https://orcid.org/0000-0002-9584-6476>
 Renyu Hu  <https://orcid.org/0000-0003-2215-8485>

References

Angelo, I., & Hu, R. 2017, *AJ*, 154, 232
 Bolton, A. S., & Schlegel, D. J. 2010, *PASP*, 122, 248
 Bourrier, V., Dumusque, X., Dorn, C., et al. 2018, *A&A*, 619, A1
 Bourrier, V., & Lecavelier des Etangs, A. 2013, *A&A*, 557, A124
 Chachan, Y., & Stevenson, D. J. 2018, *ApJ*, 854, 21
 Dai, F., Masuda, K., Winn, J. N., & Zeng, L. 2019, *ApJ*, 883, 79
 Demory, B.-O., Gillon, M., de Wit, J., et al. 2016a, *Natur*, 532, 207
 Demory, B.-O., Gillon, M., Madhusudhan, N., & Queloz, D. 2016b, *MNRAS*, 455, 2018
 Drake, J. J., Kashyap, V. L., Wargelin, B. J., & Wolk, S. J. 2020, *ApJ*, 893, 137
 Eastman, J., Gaudi, B. S., & Agol, E. 2013, *PASP*, 125, 83
 Ehrenreich, D., Bourrier, V., Bonfils, X., et al. 2012, *A&A*, 547, A18
 Ferland, G. J., Porter, R. L., van Hoof, P. A. M., et al. 2013, *RMxAA*, 49, 137
 Foreman-Mackey, D. 2016, *JOSS*, 1, 24

Fulton, B. J., Petigura, E. A., Howard, A. W., et al. 2017, *AJ*, 154, 109
 Gaidos, E., Hirano, T., Mann, A. W., et al. 2020a, *MNRAS*, 495, 650
 Gaidos, E., Hirano, T., Wilson, D. J., et al. 2020b, *MNRAS*, 498, L119
 George, E. M., Tulloch, S. M., & Ives, D. J. 2020, *JATIS*, 6, 011003
 Ginzburg, S., Schlichting, H. E., & Sari, R. 2018, *MNRAS*, 476, 759
 Hirano, T., Krishnamurthy, V., Gaidos, E., et al. 2020, *ApJL*, 899, L13
 Holmström, M., Ekenbäck, A., Selsis, F., et al. 2008, *Natur*, 451, 970
 Home, K. 1986, *PASP*, 98, 609
 Hu, R., Seager, S., & Yung, Y. L. 2015, *ApJ*, 807, 8
 Hunter, J. D. 2007, *CSE*, 9, 90
 Husser, T.-O., Wende-von Berg, S., Dreizler, S., et al. 2013, *A&A*, 553, A6
 Ionov, D. E., & Shematovich, V. I. 2015, *SoSyR*, 49, 339
 Ito, Y., Ikoma, M., Kawahara, H., et al. 2015, *ApJ*, 801, 144
 Jindal, A., de Mooij, E. J. W., Jayawardhana, R., et al. 2020, *AJ*, 160, 101
 Kasper, D., Bean, J. L., Oklopčić, A., et al. 2020, *AJ*, 160, 258
 Kim, S., Prato, L., & McLean, I. 2015, REDSPEC: NIRSPEC Data Reduction, Astrophysics Source Code Library, ascl:1507.017
 Kreidberg, L. 2015, *PASP*, 127, 1161
 Lee, E. J., & Chiang, E. 2017, *ApJ*, 842, 40
 Linsky, J. L., Fontenla, J., & France, K. 2014, *ApJ*, 780, 61
 Lopez, E. D., & Fortney, J. J. 2013, *ApJ*, 776, 2
 Lord, S. D. 1992, A New Software Tool for Computing Earth’s Atmospheric Transmission of near- and Far-infrared Radiation, NASA Tech. Memo. 103957
 Malsky, I., & Rogers, L. A. 2020, *ApJ*, 896, 48
 Marsh, T. R. 1989, *PASP*, 101, 1032
 Martin, E. C., Fitzgerald, M. P., McLean, I. S., et al. 2018, *Proc. SPIE*, 10702, 107020A
 Mazeh, T., Tamuz, O., & Zucker, S. 2007, in ASP Conf. Ser. 366, The System Detrending Algorithm: Implementation and Testing, ed. C. Afonso, D. Welldrake, & T. Henning (San Francisco, CA: ASP), 119
 Mignone, A., Bodo, G., Massaglia, S., et al. 2007, *ApJS*, 170, 228
 Miguel, Y. 2019, *MNRAS*, 482, 2893
 Murray-Clay, R. A., Chiang, E. I., & Murray, N. 2009, *ApJ*, 693, 23
 Ninan, J. P., Stefansson, G., Mahadevan, S., et al. 2020, *ApJ*, 894, 97
 Nortmann, L., Pallé, E., Salz, M., et al. 2018, *Sci*, 362, 1388
 Oklopčić, A. 2019, *ApJ*, 881, 133
 Oklopčić, A., & Hirata, C. M. 2018, *ApJL*, 855, L11
 Owen, J. E. 2019, *AREPS*, 47, 67
 Owen, J. E., & Jackson, A. P. 2012, *MNRAS*, 425, 2931
 Owen, J. E., & Wu, Y. 2016, *ApJ*, 817, 107
 Palte, E., Nortmann, L., Casasayas-Barris, N., et al. 2020, *A&A*, 638, A61
 Rafikov, R. R. 2006, *ApJ*, 648, 666
 Ridden-Harper, A. R., Snellen, I. A. G., Keller, C. U., et al. 2016, *A&A*, 593, A129
 Salz, M., Banerjee, R., Mignone, A., et al. 2015a, *A&A*, 576, A21
 Salz, M., Czesla, S., Schneider, P. C., & Schmitt, J. H. M. M. 2016, *A&A*, 586, A75
 Salz, M., Czesla, S., Schneider, P. C., et al. 2018, *A&A*, 620, A97
 Salz, M., Schneider, P. C., Czesla, S., & Schmitt, J. H. M. M. 2015b, *A&A*, 576, A42
 Sanz-Forcada, J., Micela, G., Ribas, I., et al. 2011, *A&A*, 532, A6
 Shaw, G., & Ferland, G. J. 2020, *MNRAS*, 493, 5153
 Shematovich, V. I., Ionov, D. E., & Lammer, H. 2014, *A&A*, 571, A94
 Spake, J. J., Sing, D. K., Evans, T. M., et al. 2018, *Natur*, 557, 68
 Speagle, J. S. 2020, *MNRAS*, 493, 3132
 Stone, J. M., Tomida, K., White, C. J., & Felker, K. G. 2020, *ApJS*, 249, 4
 Sulis, S., Dragomir, D., Lendl, M., et al. 2019, *A&A*, 631, A129
 Sur, A., Haskell, B., & Kuhn, E. 2020, *MNRAS*, 495, 1360
 Tremblin, P., & Chiang, E. 2013, *MNRAS*, 428, 2565
 Tsiaras, A., Rocchetto, M., Waldmann, I. P., et al. 2016, *ApJ*, 820, 99
 Valencia, D., Ikoma, M., Guillot, T., & Nettelmann, N. 2010, *A&A*, 516, A20
 van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, 13, 22
 Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, 17, 261
 Vissapragada, S., Knutson, H. A., Jovanovic, N., et al. 2020, *AJ*, 159, 278
 Wang, L., & Dai, F. 2018, *ApJ*, 860, 175
 Wang, L., & Dai, F. 2020, arXiv:2101.00045
 Zilinskas, M., Miguel, Y., Mollière, P., & Tsai, S.-M. 2020, *MNRAS*, 494, 1490