CrossMark

No Evidence for Lunar Transit in New Analysis of *Hubble Space Telescope* Observations of the Kepler-1625 System

Laura Kreidberg[1,2] , Rodrigo Luger[3] , and Megan Bedell[3]
[1] Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA
[2] Harvard Society of Fellows, 78 Mount Auburn Street, Cambridge, MA 02138, USA
[3] Flatiron Institute, Simons Foundation, 162 Fifth Avenue, New York, NY 10010, USA

Abstract

Observations of the Kepler-1625 system with *Kepler* and the *Hubble Space Telescope* have suggested the presence of a candidate exomoon, Kepler-1625b I, a Neptune-radius satellite orbiting a long-period Jovian planet. Here we present a new analysis of the *Hubble* observations, using an independent data reduction pipeline. We find that the transit light curve is well fit with a planet-only model, with a best-fit χ^2_ν equal to 1.01. The addition of a moon does not significantly improve the fit quality. We compare our results directly with the original light curve from Teachey & Kipping, and find that we obtain a better fit to the data using a model with fewer free parameters (no moon). We discuss possible sources for the discrepancy in our results, and conclude that the lunar transit signal found by Teachey & Kipping was likely an artifact of the data reduction. This finding highlights the need to develop independent pipelines to confirm results that push the limits of measurement precision.

Key words: planetary systems

Supporting material: data behind figure

1. Introduction

Moons are abundant in the solar system, and provide clues to the formation history, evolution, and even habitability of the planets that they orbit. The great scientific potential of moons has prompted an extensive search for lunar companions in exoplanetary systems (exomoons), and creative development of new search techniques (e.g., Kipping 2009a, 2009b; Simon et al. 2010; Kipping et al. 2013; Peters & Turner 2013; Heller et al. 2014; Noyola et al. 2014; Agol et al. 2015; Hippke 2015; Sengupta & Marley 2016; Vanderburg et al. 2018).

Recently, a potential exomoon candidate was identified in the Kepler-1625 system (Teachey et al. 2018). The host planet, Kepler-1625b, has a radius consistent with that of Jupiter and an orbital period of 287 days. The first evidence for the exomoon candidate, Kepler-1625b I, was based on observations from *Kepler*. The light curve showed drops in stellar flux that were attributed to a transiting exomoon; however, later analysis called this result into question, showing that the moon transit features were highly sensitive to the *Kepler* reduction pipeline and the algorithm used to detrend the data (Rodenbeck et al. 2018; Teachey & Kipping 2018).

Subsequent follow-up observations with *Hubble Space Telescope* (*HST*) revived the possibility of an exomoon in the system based on two factors: a small drop in the system flux after the planet's transit egress, and a transit timing variation (Teachey & Kipping 2018, hereafter TK18). The best-fit moon had a large radius (comparable to that of Neptune), and if real, is unlike any moon in the solar system.

As the primary evidence for the exomoon now rests on the *HST* transit light curve, in this work we perform an independent reduction and fit to the *HST* data and compare it to the results from TK18.

2. Observations and Data Reduction

The Kepler-1625 system was observed with 26 continuous *HST* orbits on 2017 October 28–29 (Program GO 15149: PI: A. Teachey). The observations used the Wide Field Camera 3 (WFC3) G141 grism in staring mode, which fixed the spectrum in a constant position on the detector. At the beginning of the visit, there was a single exposure taken with the F130N filter, which is used to determine the position of the spectral trace. The following exposures used the G141 grism. See TK18 for additional description of the observation design.

We reduced the *HST* data using custom software developed in Kreidberg et al. (2014). This software has yielded consistent results with multiple independent pipelines (e.g., Knutson et al. 2014; Spake et al. 2018). We ran our pipeline on the flt data product provided by the Space Telescope Science Institute (STScI). In keeping with previous WFC3 analysis, we discarded the first orbit of data, where the instrument systematics have larger amplitude. We also discarded exposures taken during the South Atlantic Anomaly passage (exposures 107, 116, 125, and 126).

To begin the data reduction, we fit the centroid of the direct image with a two-dimensional Gaussian. The centroid position determines the position of the spectral trace, which we calculated using the coefficients provided in the configuration file from STScI: G141.F130N.V4.32.conf.[4] To process the spectra, we flatfielded the raw data using the spectroscopic flatfield coefficients provided by STScI in WFC3.IR.G141.flat.2.fits, following the instructions in Section 6 of the aXe User Manual.[5] We then created an extraction box centered on the spectral trace. We varied the height and width of the box in 1 pixel increments to find the window that minimized the root

---

[4] Available at http://www.stsci.edu/hst/wfc3/analysis/grism_obs/calibrations/wfc3_g141.html.
[5] http://ane-info.stsci.edu/

mean square (rms) deviation from the best fit to the transit light curve. The best was $450 < X < 574$, and $522 < Y < 536$, where X and Y are physical pixels in the spectral and spatial direction, respectively.

We reduced the grism exposures with the optimal extraction routine of Horne ([1986](#)), which minimizes background noise in the extracted spectrum by preferentially weighting pixels that are dominated by the target spectrum. The inputs for optimal extraction are the background-subtracted data array, and estimate of the error per pixel, an initial guess for the spectrum and its uncertainty, and a mask array for bad pixels. For the initial guess of the spectrum, we did a simple box extraction (sum over all rows in the extraction window), and assumed that the variance was equal to the box-extracted spectrum. We subtracted the background from the data array as described in [2.1](#). For the error array, we used a quadrature sum of the photon noise (the square root of the pixel counts), the read noise (12 photoelectrons for `flt` files; WFC3 Data Handbook,[6]) and the error due to background subtraction. The initial pixel mask marked all pixels as good.

To optimally extract the spectrum, we first created a smoothed image by median-filtering each row of the data with a 9-pixel-wide window. We then normalized the smoothed image by dividing each column by its sum, and multiplied it by the best-guess spectrum. We compared the smoothed image to the real data and masked outliers in the data that are greater than a threshold $\sigma_{cut} = 7.5$. We then recomputed the best-guess spectrum with the new mask and the optimal weights from Horne ([1986](#)). The process was iterated until no outliers greater than the threshold remained. To create the broadband transit light curve, we summed each optimally extracted spectrum over all wavelengths.

The broadband light curve is shown in Figure [1](#), in comparison to the light curve from [TK18](#). We note that there are differences between the two data sets, particularly a kink near the moon-like transit feature identified by [TK18](#).

2.1. Background Subtraction

The star Kepler-1625 is faint (H mag $= 14.0$) relative to most other exoplanet host stars observed with WFC3, which makes accurate background subtraction especially important for this target. Moreover, the host star is in a crowded field, so the pixels used to estimate the background must be chosen carefully to avoid contamination from other stars. To estimate the background counts, we masked pixels with total counts larger than 800 electrons (2.7 electrons s^{-1}) and took the median count in the unmasked pixels. The per pixel uncertainty due to background subtraction is 1.4826 times the median absolute deviation.

2.2. Pointing Drift Measurement

The position of the spectrum on the detector shifts slightly over time ($\sim$0.1 pixel day^{-1}) due to the spacecraft's pointing drift. This drift can change the flux measured for the target star: if the spectrum moves onto less sensitive pixels, fewer photoelectrons are recorded. To enable a correction for this effect, we measured the position of the spectrum over time. Figure [2](#) shows the best-fit shifts.
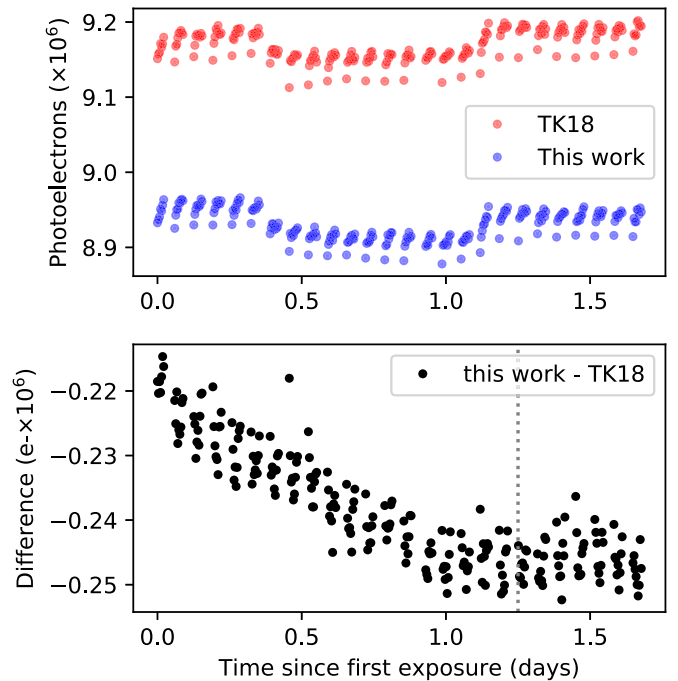
Figure 1. Top panel: extracted light curve from this work (blue) compared to [TK18](#) (red). Bottom panel: the difference in photoelectron counts between the data sets. The moon ingress identified by [TK18](#) is marked by the dotted gray line. The data used to create this figure are available.
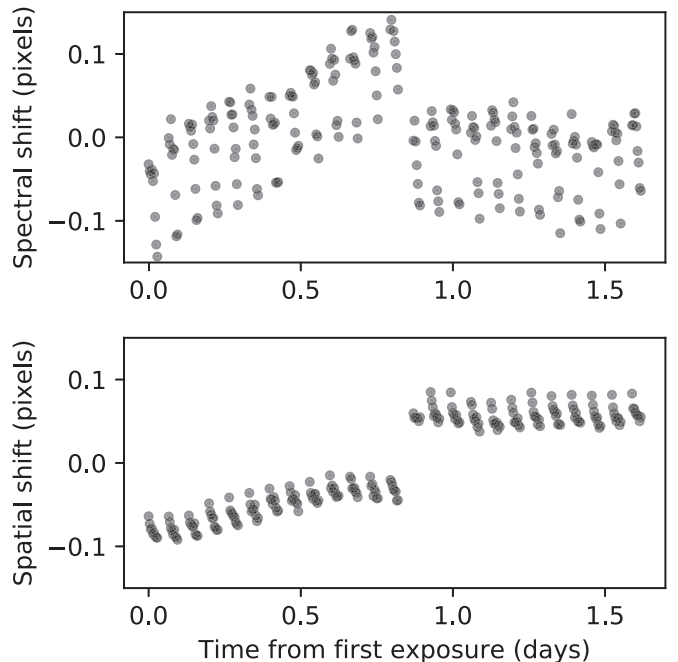
Figure 2. Shift (in pixels) relative to the mean position of the spectrum in the spectral direction (top panel) and spatial direction (bottom panel). The largest shift occurs after orbit 14 due to a guide star reacquisition.

To measure shifts in the spatial direction, we summed each `flt` image over all columns (which we dub the "column sum"). We used the first exposure in the visit as a template, and for each subsequent exposure, we used least-squares minimization to calculate the shift in pixels that minimized the difference between its column sum and the template. The shifts are a

fraction of a pixel, so we used the NumPy `interp` routine to do linear interpolation on a sub-pixel scale. The WFC3 point-spread function is undersampled, so we convolved each column sum with a 4-pixel-wide Gaussian before the interpolation (following Deming et al. 2013).

To measure the spectral shifts, we repeated this procedure with two differences: (1) we used the optimally extracted spectrum rather than the column sum; and (2) in addition to calculating the best-fit shift, we also calculated a best-fit normalization factor (a scalar multiple for the whole spectrum) to ensure that our results are not biased by the varying brightness of the host star during the planet's transit.

3. Analysis

The extracted transit light transit light curve contains both astrophysical signals and instrument systematic noise, which we model simultaneously.

3.1. Astrophysics Model

For the astrophysics, we used the `planet-planet` package (Luger et al. 2017), a photodynamical code that calculates light curves for multiple occulting bodies orbiting a star. Within `planet-planet`, the orbits are computed with the N-body integrator `REBOUND` (Rein & Liu 2012), which calculates the three-dimensional motion of the star, planet, and moon over time under the influence of gravity.

In our analysis, we considered two scenarios: a no-moon model and a moon model. The free parameters for the no-moon model were: the stellar radius, the planet radius, the planet's time of central transit, the planet inclination, and the planet mass. For the moon model, we added a third body with six free parameters: moon radius, moon time of central transit, moon orbital period, moon inclination, moon mass, and longitude of the ascending node relative to that of the planet. We fixed the eccentricity of all bodies to zero. We also fixed the orbital period of the planet–moon barycenter to 287.378949 days (the best fit from TK18). We elected not to vary the period of the planet–moon barycenter because it is poorly constrained from a single transit observation. A longer period would cause a longer transit duration for the planet, but this could equivalently result from a smaller impact parameter, a larger star, or a massive moon that significantly perturbs the planetary orbit.

We used the following priors: the stellar radius was normally distributed, $R_* \sim N(1.81, 0.17)R_\odot$ (see the next section). The planet radius was uniform from 8 to $14R_\oplus$. The planet transit time was uniform over the timespan of the observations, and inclination was uniform from 0° to 90°. The planet mass was log-normally distributed, $M_p \sim 10^{2.5\pm0.5}M_\oplus$, based on the expectation for Jupiter-radius objects from Ning et al. (2018). For the moon model, we allowed the transit time to vary uniformly over the entire visit. The moon period was uniform between 1.6 and 260 days. These limits span the duration of the *HST* observations (so there is one possible moon occultation event), to the orbit at 0.5 the Hill radius, based on the stability limit for prograde moon orbits (Domingos et al. 2006). The Hill radius calculation assumed that the planet and stellar masses are $1M_{Jup}$ and $1.37M_\odot$. The moon mass varied uniformly from 0 to $30M_\oplus$. The longitude of the ascending node was also uniform from 0° to 360°. The moon inclination was uniform from 0° to 90°, and was defined relative to the line of sight. We

assigned zero prior probability to scenarios where the moon did not transit or experienced a grazing transit. We made this choice to put an upper limit on the radius of a fully transiting moon; for grazing transits or non-transits the moon radius could be arbitrarily large.

3.1.1. Stellar Parameters

For both the moon and no-moon scenarios, we used a quadratic stellar limb-darkening law and fixed the coefficients to the prediction for a 5700 K, solar metallicity `PHOENIX` model from Espinoza & Jordán (2015); u_1, u_2 = [0.216, 0.183].

We estimated the host star parameters using the *Gaia* second data release (DR2) parallax (Gaia Collaboration et al. 2016, 2018) along with *UBV* photometry from Everett et al. (2012) and *JHK* photometry from the Two Micron All-Sky Survey (Skrutskie et al. 2006). We employed the isochrone python package (Morton 2015) with the Dartmouth isochrone grid (Dotter et al. 2008) to obtain posterior constraints on the stellar parameters. The resulting parameters indicate that Kepler-1625 has stellar mass $1.37^{+0.13}_{-0.16}\,M_\odot$, radius $1.81^{+0.18}_{-0.16}\,R_\odot$, and age $2.8^{+1.6}_{-1.2}$ Gyr. In our analysis, we fixed the stellar mass to the best-fit value ($1.37M_\odot$), and used a Gaussian prior on the radius, $R_* \sim N(1.81, 0.17)$.

3.2. Instrument Systematics Model

There are two systematic trends in the data. One is the orbit-long ramp, attributed to charge traps in the detector filling up over the orbit (Zhou et al. 2017). The other is a visit-long trend over multiple orbits, which could be due to shifts in the target star position onto more/less sensitive pixels.

For our primary fit, we used a linear combination of the spectrum's X and Y position (shown in Figure 2), multiplied by the non-parametric orbital ramp model from TK18, which assigns each of the nine exposures per orbit a normalization constant, c_1, ..., c_9. In sum, for exposure number i, the systematics model S is

$$S_i = c_j \times (1 + aX_i + bY_i) \qquad (1)$$

where a and b, and c_j are free parameters, and $j = i \bmod 9 + 1$ is the exposure number relative to the first exposure in the orbit.

For comparison, we also tested the "exponential" systematics model from TK18, which combines an exponential visit-long trend, an offset after orbit 14 when the guide stars were reacquired, and the non-parametric orbital ramp model.

3.3. Light Curve Fits

We fit the broadband transit light curve using the models described above. We determined the best-fit model parameters with least-squares minimization. We also ran a Markov chain Monte Carlo (MCMC) fit to determine the posterior probability of the parameters. For the MCMC, we held the ramp parameters c_1, ..., c_9 fixed at their best-fit values. The MCMC used the `emcee` package (Foreman-Mackey et al. 2013) with 50 walkers and ran for 5000 steps. We discarded the first 20% of the MCMC chain as burn-in. As a quick test for convergence, we divided the remainder of the chain in half and confirmed that the results from the first half were consistent with the second half.
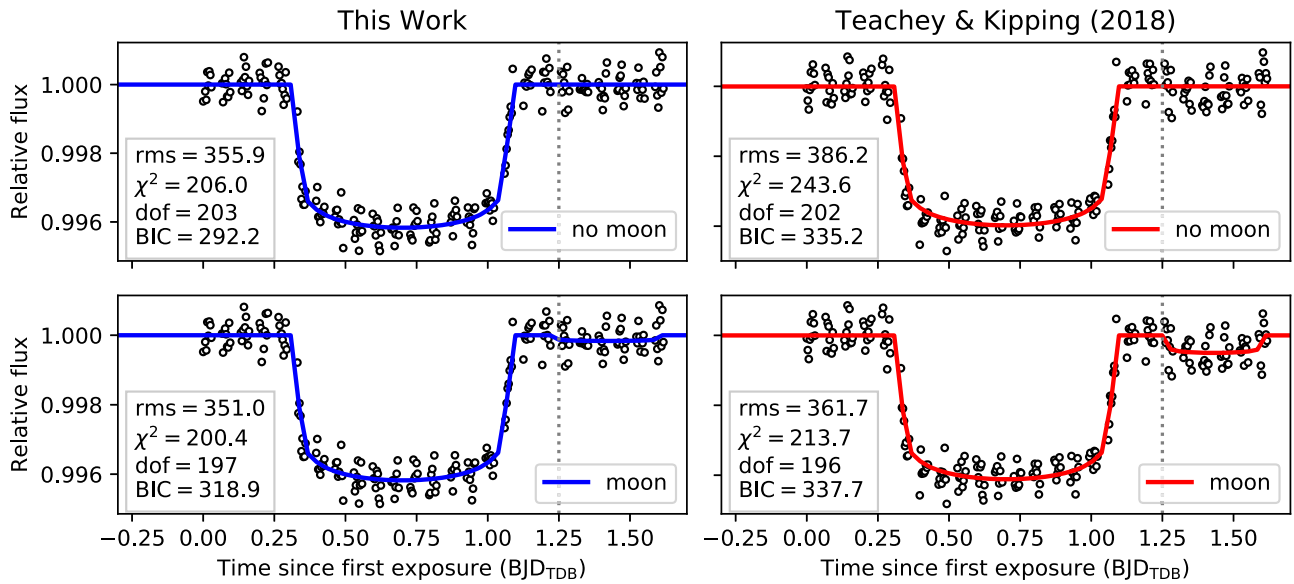
Figure 3. Best-fit models compared to normalized transit light curves from this work (left, blue) and from TK18 (right, red). The top panel shows the best-fit no-moon model, and the bottom panel shows the best-fit moon model. The lower left of each panel indicates the fit rms (in ppm), the χ_2, the degrees of freedom, and the Bayesian information criterion (BIC). Each light curve is divided by its best-fit systematics model (*XY* decorrelation for this work; exponential and offset for TK18). The dotted gray line marks the possible moon ingress identified by TK18.

4. Results

We obtained an excellent fit to the light curve with the no-moon model, as illustrated in Figure 3. The residuals to the no-moon model fit have rms equal to 356 parts per million (ppm), which is within 3% of the predicted photon shot noise (367 ppm), and yields a $\chi_\nu^2 = 1.01$. The binned rms decreases with the square root of the number of points per bin, as expected for photon noise-limited statistics (see rms versus bin size in Figure 4).

We obtained a slightly lower rms with the moon model (351 ppm); however, this is not a large enough improvement in fit quality to merit the addition of six additional free parameters. The moon model has a small increase in χ_ν^2 to 1.02. According to the Bayesian information criterion (BIC), which penalizes unnecessary model complexity, the moon model is disfavored with ΔBIC = 26.7. This constitutes strong evidence against the inclusion of a moon (Kass & Raftery 1995). In addition, as shown in Figure 5, the posterior distribution for the moon transit time spans the entire duration of the observations (2σ confidence). The upper limit on the moon radius is $3.6\,R_\oplus$ at 3σ confidence.

We found that our results were unchanged when we used the exponential systematics model from TK18 (described in Section 3.2). For this case, the moon model is also strongly disfavored (ΔBIC = 32.2). The fit rms is within 5 ppm of the *XY* decorrelation model.

The posterior probabilities for the planet's mass and radius are consistent with either a gas giant planet or brown dwarf. The radius is $11.1^{+0.5}_{-0.2}\,R_\oplus$ and mass is $10^{2.2\pm0.5}M_\oplus$.

4.1. Comparison with Teachey & Kipping (2018)

TK18 found evidence for the transit of a Neptune-size moon in their analysis of the *HST* data, in contrast to the findings presented here. To compare our results with theirs, we fit the TK18 light curve directly. We fit the astrophysical signal with both the no-moon and moon models, and used the
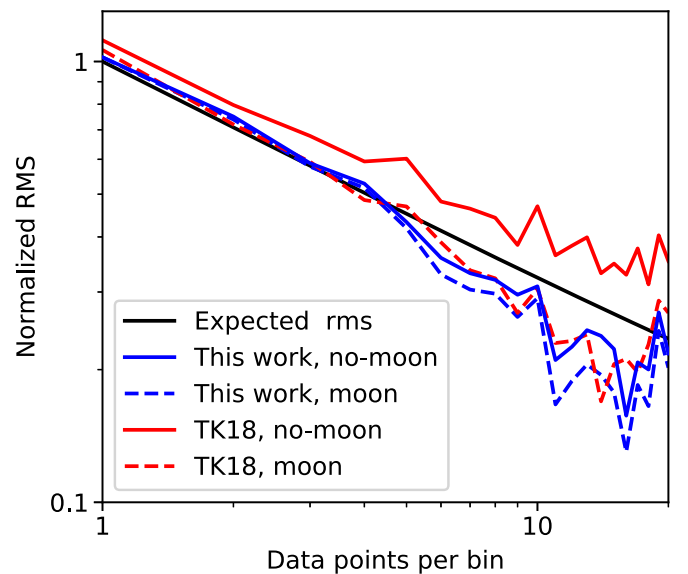

Figure 4. Light curve rms vs. bin size for the best fit no-moon model (solid lines) and moon model (dashed lines), for data from this work (red) and TK18 (blue). The fits to data from this work agree well with the expected photon-limited, $\sqrt{N}$ decrease in rms with bin size (black line). We also reach the photon limit for the TK18 data, but only for the moon model. The rms for the no-moon fit (red dashed line) ranges from 1.1 to 1.5× the photon limit for 1–20 points per bin.

exponential systematics model. Figure 3 shows the best-fit models.

Similar to the findings of TK18, the moon model improved the fit quality by a $\Delta\chi^2 = 29.9$. Notably, however, the moon model fit to the TK18 data does not perform better than the no-moon model fit to our new data (rms of 362 versus 356 ppm), even with the additional seven free parameters.

The moon model also yields qualitatively different posterior distributions for the two data sets. As shown in Figure 5, for the TK18 data the moon radius and transit time are peaked at
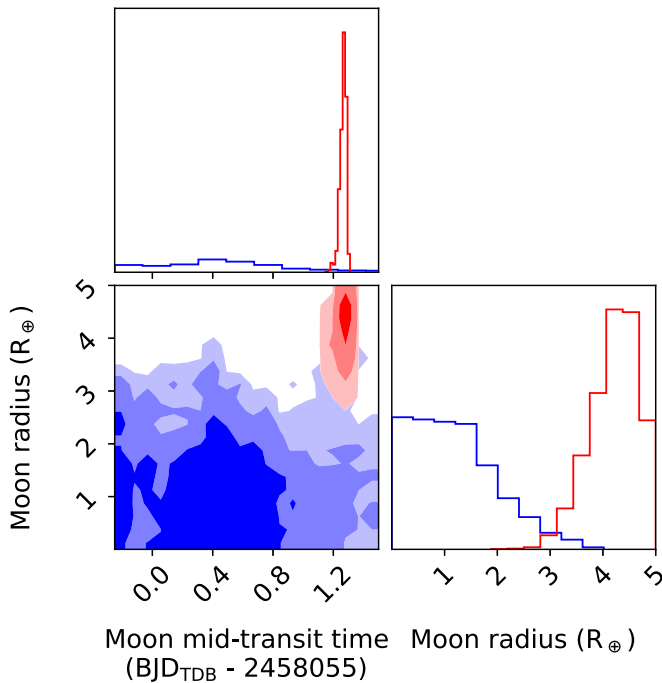
Figure 5. Posterior distributions for the moon radius and time of transit based on an MCMC fit to data from this work (blue) and from TK18 (red). The shading corresponds to 1, 2, and 3σ contours (from darkest to lightest). These values are marginalized over all other model parameters.

$r_{\rm moon} = 4.3 \pm 0.5\,R_{\oplus}$ and $t_{\rm moon} = 2458056.29^{+0.06}_{-0.04}\,{\rm BJD_{TDB}}$. By contrast, the fit to the new data presented here yields an upper limit on the moon radius of $3.6\,R_{\oplus}$ at 3σ confidence, and the transit time is unconstrained.

Although the two data sets yield different constraints on the moon properties, the planet's mid-transit time agrees to better than 1σ for the two fits. The transit time is earlier than expected based on the *Kepler* data (3σ confidence; TK18), suggesting that there are transit timing variations in the system. Such a variation could arise from the presence of a moon, as suggested by TK18; however, the variation could also be caused by another planet in the system.

5. Discussion and Conclusions

A natural question arising from our analysis is the source of the discrepancy between TK18 and the new results presented here. We find that with our new data, there is strong evidence against the moon ($\Delta {\rm BIC} > 25$), even when we use a comparable model to TK18 (a six-parameter moon model and an exponential fit to the visit-long trend).

If the model is not the source of the difference, it must be the extracted transit light curves themselves. Figure 1 shows a direct comparison of the light curves. The count rate we measure is 2.46%–2.74% lower than the TK18 light curve, and there is a small bump in the difference between the two data sets near the location of the moon transit identified in the TK18 data (see the bottom panel). This bump may be the source of the moon feature reported in TK18.

We explored several modifications to our pipeline to attempt to reproduce the TK18 data reduction. These included rotating the image by $0°.5$, using the same aperture as TK18 to extract the spectrum, and scaling the master sky flat for the background subtraction (rather than just subtracting the median). None of these modifications had a significant effect on our results.

There are a few other steps in the TK18 data reduction that would require substantial modification of our pipeline to recreate, but seem unlikely to be responsible for the difference. One of these is outlier masking. TK18 identify outliers with a Gaussian process fit to the pixel-level light curves, compared to our optimal extraction approach. Despite the difference, both methods flag $\sim 0.01\%$ of pixels as bad. We also do not use the STScI software aXeprep to embed the raw 256×256 image in a larger array; however, this process primarily affects the edge of the image, many pixels distant from the extraction box, so it is unclear how this step would bias the light curve. We conclude that no single choice in the data reduction provides an easy explanation for the difference in our light curves.

During the referee process for this work, we learned of another manuscript that also reanalyzed the *HST* transit observation (Heller et al. 2019). The best fit favored a moon model similar to that found by TK18; however, an MCMC analysis did not converge on this model, leading the authors to conclude that the highest likelihood solution may be an outlier.

Taken together, these findings illustrate the challenge of pushing measurement precision to the 100 ppm level, and highlight the importance of developing multiple independent pipelines to confirm cutting-edge results.

ORCID iDs

Laura Kreidberg ⬤ https://orcid.org/0000-0003-0514-1147
Rodrigo Luger ⬤ https://orcid.org/0000-0002-0296-3826
Megan Bedell ⬤ https://orcid.org/0000-0001-9907-7742

References

Agol, E., Jansen, T., Lacy, B., Robinson, T. D., & Meadows, V. 2015, ApJ, 812, 5
Deming, D., Wilkins, A., McCullough, P., et al. 2013, ApJ, 774, 95
Domingos, R. C., Winter, O. C., & Yokoyama, T. 2006, MNRAS, 373, 1227
Dotter, A., Chaboyer, B., Jevremović, D., et al. 2008, ApJS, 178, 89
Espinoza, N., & Jordán, A. 2015, MNRAS, 450, 1879
Everett, M. E., Howell, S. B., & Kinemuchi, K. 2012, PASP, 124, 316
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306
Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, A&A, 616, A1
Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, A&A, 595, A1
Heller, R., Rodenbeck, K., & Bruno, G. 2019, A&A, 624, A95
Heller, R., Williams, D., Kipping, D., et al. 2014, AsBio, 14, 798
Hippke, M. 2015, ApJ, 806, 51
Horne, K. 1986, PASP, 98, 609
Kass, R. E., & Raftery, A. E. 1995, J. Am. Stat. Assoc., 90, 773
Kipping, D. M. 2009a, MNRAS, 392, 181
Kipping, D. M. 2009b, MNRAS, 396, 1797
Kipping, D. M., Hartman, J., Buchhave, L. A., et al. 2013, ApJ, 770, 101
Knutson, H. A., Dragomir, D., Kreidberg, L., et al. 2014, ApJ, 794, 155

Kreidberg, L., Bean, J. L., Désert, J.-M., et al. 2014, Natur, 505, 69

Luger, R., Lustig-Yaeger, J., & Agol, E. 2017, ApJ, 851, 94

Morton, T. D. 2015, Isochrones: Stellar Model Grid Package, Astrophysics Source Code Library, ascl:1503.010

Ning, B., Wolfgang, A., & Ghosh, S. 2018, ApJ, 869, 5

Noyola, J. P., Satyal, S., & Musielak, Z. E. 2014, ApJ, 791, 25

Peters, M. A., & Turner, E. L. 2013, ApJ, 769, 98

Rein, H., & Liu, S.-F. 2012, A&A, 537, A128

Rodenbeck, K., Heller, R., Hippke, M., & Gizon, L. 2018, A&A, 617, A49

Sengupta, S., & Marley, M. S. 2016, ApJ, 824, 76

Simon, A. E., Szabó, G. M., Szatmáry, K., & Kiss, L. L. 2010, MNRAS, 406, 2038

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163

Spake, J. J., Sing, D. K., Evans, T. M., et al. 2018, Natur, 557, 68

Teachey, A., & Kipping, D. M. 2018, SciA, 4, eaav1784

Teachey, A., Kipping, D. M., & Schmitt, A. R. 2018, AJ, 155, 36

Vanderburg, A., Rappaport, S. A., & Mayo, A. W. 2018, AJ, 156, 184

Zhou, Y., Apai, D., Lew, B. W. P., & Schneider, G. 2017, AJ, 153, 243