



Analysis of a JWST NIRSpec Lab Time Series: Characterizing Systematics, Recovering Exoplanet Transit Spectroscopy, and Constraining a Noise Floor

Zafar Rustamkulov¹ , David K. Sing^{1,2} , Rongrong Liu² , and Ashley Wang² ¹ Department of Earth and Planetary Science, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218, USA; zafar@jhu.edu² Department of Physics and Astronomy, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218, USA

Received 2022 February 16; revised 2022 March 4; accepted 2022 March 7; published 2022 March 21

Abstract

The James Webb Space Telescope’s (JWST) NIRSpec instrument will unveil the nature of exoplanet atmospheres across the wealth of planet types, from temperate terrestrial worlds to ultrahot Jupiters. In particular, the 0.6–5.3 μm PRISM mode is especially well suited for efficient spectroscopic exoplanet observations spanning a number of important spectral features. We analyze a lab-measured NIRSpec PRISM mode Bright Object Time Series observation from the perspective of a JWST user to understand the instrument performance and detector properties. We create two realistic transiting exoplanet time-series observations by performing injection-recovery tests on the lab-measured data to quantify the effects of real instrument jitter, drift, intrapixel sensitivity variations, and $1/f$ noise on measured transmission spectra. By fitting the time-series systematics simultaneously with the injected transit, we can obtain more realistic transit-depth uncertainties that take into account noise sources that are currently not modeled by traditional exposure time calculators. We find that sources of systematic noise related to intrapixel sensitivity variations and point-spread function motions are apparent in the data at the level of a few hundred ppm but can be effectively detrended using a low-order polynomial with detector position. We recover the injected spectral features of GJ 436 b and TRAPPIST-1 d and place a 3σ upper limit on the detector noise floor of 14 ppm. We find that the noise floor is consistent with <10 ppm at the 1.7σ level, which bodes well for future observations of challenging targets with faint atmospheric signatures.

Unified Astronomy Thesaurus concepts: [Exoplanet atmospheres \(487\)](#); [Transmission spectroscopy \(2133\)](#); [Space telescopes \(1547\)](#)

1. Introduction

The James Webb Space Telescope (JWST), now in its L2 orbit, will soon shed new light on transiting exoplanet atmospheres. It will probe at high signal-to-noise ratio planets spanning the phase space of mass, radius, temperature, and age, yielding our most comprehensive view yet. Recent discoveries by the Transiting Exoplanet Survey Satellite (TESS) (Ricker et al. 2015) are beginning to add to an ever-growing catalog of nearby exoplanets that are excellent targets for efficient spectral observations with JWST (e.g., Vanderspek et al. 2019; Winters et al. 2019; Jenkins et al. 2020; Hobson et al. 2021). A variety of spectral modes are available with NIRSpec, each with unique wavelength coverage and spectral resolution. NIRSpec’s versatility made it an extremely popular mode for Cycle 1 GO proposers, with roughly two-thirds of awarded exoplanet time-series programs making use of one or more of its modes.³ Several targets on the Guaranteed Time Observations (GTO) and Early Release Science (ERS) lists also widely make use of NIRSpec.

NIRSpec’s low-resolution ($\lambda/\Delta\lambda \sim 20\text{--}100$) PRISM mode has a uniquely broad spectral range (0.6–5.3 μm), offering a doubling in throughput relative to other JWST mode combinations, which require a minimum of two separate transits to achieve the same wavelength coverage. This property makes PRISM the ideal instrument to carry out efficient, “one-shot” measurements of a

number of carbon-, oxygen-, and even nitrogen-bearing molecules in exoplanet atmosphere in all three of the main observational geometries (i.e., Batalha & Line 2017; Morley et al. 2017; Lustig-Yaeger et al. 2019). Measurements of these information-rich chemical species will allow for exquisite comparative studies between planets and will provide important insights into the mechanisms that sculpt planetary atmospheres and populations.

In this work, we used a 3 hr Bright Object Time Series (BOTS) data set measured from the NIRSpec instrument in the lab to look at the data from the perspective of a JWST observer to quantify how the instrument’s systematics and noise performance affect the final scientific product. The Pandexo (Batalha & Line 2017) and Pandeia (Pontoppidan et al. 2016) packages have served as excellent community tools for estimating exposure times and building synthetic JWST observations. We build off the recent work by Birkmann et al. (2022) and Jakobsen et al. (2022), who characterize many important engineering-side detector diagnostics, and show how these nuances will affect a real exoplanet observation. Similar characterization work with NIRCam has shown that the detector’s noise floor from known systematics sources is 9 ppm and that its systematics are similar in nature and amplitude to the ones outlined in our work (Schlawin et al. 2020, 2021).

As is often seen in time-series data from Spitzer IRAC (e.g., Deming et al. 2015; May & Stevenson 2020), TESS, and Kepler/K2 (Luger et al. 2016, 2018), undersampled point-spread functions (PSFs) often introduce a nonnegligible source of systematic noise due to the PSF’s motion across pixels with nonuniform inter- and intrapixel sensitivity variations. Thankfully, with knowledge of the PSF’s position over time, generalized pixel-level decorrelation methods are able to largely remove these signatures in photometry. With exoplanet time-series

³ <https://www.stsci.edu/jwst/science-planning/user-committees/jwst-users-committee>




Figure 1. Overview of the lamp spectrum. Top: a single 0.45 s integration of the lamp spectrum summed along the cross-dispersion axis. Prominent frost absorption features in NIRSpec’s optical path sculpt the spectrum’s shape. Bottom: a single log-stretched raw NIRSpec PRISM frame. The aspect ratio of the detector is vertically exaggerated by a factor of ~ 2 for clarity.

spectroscopy, analogous jitter decorrelation techniques (e.g., Sing et al. 2019) have shown promise in removing systematic trends and greatly improving light-curve fits. In our work, we leverage a ~ 0.3 pixel, ~ 30 mas drift of the light source on the detector to measure the amplitude of shift-induced systematics in the light curve and the extent to which they can be decorrelated. Further, we fit the observed systematics jointly with the transit light curve at each wavelength channel of our injected transit spectrophotometry and perform a uniform model intercomparison to predict which systematics vectors are important in the detrending of real JWST time-series data. Lastly, we select a wavelength channel that displays weak systematic noise and bin it down through time to place an empirical upper limit on NIRSpec’s noise floor.

2. Overview of the CV3 PRISM Time Series

We analyze an experimental data set collected in 2016 January during the cryo-vacuum testing campaign (CV3) of the JWST Integrated Science Instrument Module (ISIM) at NASA’s Goddard Space Flight Center. The single BOTS exposure was taken with NIRSpec held at a temperature of ~ 40 K (its L2 operating temperature) and consists of 12,000 integrations with three groups each, corresponding to an effective integration time of 0.45 s. The exposure used the SUB512 subarray, with size of 512×32 pixels, which is wide enough to encapsulate the PSF of the light source and the unilluminated region of the detector. The NRSRAPID readout pattern and S1600A1 ($1''.6 \times 1''.6$) slit were used. The data were processed by the NIRSpec ramp-to-slopes pipeline to fit the up-the-ramp count rate, the statistical variance, and the quality flag for each pixel (see Birkmann et al. 2022, Section 7). The CV3 BOTS prism data are publicly available through the ESA website.⁴ Full details of the NIRSpec CV3 observation can be found in Birkmann et al. (2022), which we briefly summarize here.

The NIRSpec instrument was illuminated with a tungsten filament source with a similar PSF width to that expected from JWST (see Figure 1). The PSF displays a spatial drift of ~ 30 mas (~ 0.3 px) over the course of the exposure, which is $\sim 4\times$ larger than the 1σ rms stability requirement for JWST. The light source displays a long-term drift in the flux level at the percent level over several hours, as well as chromatic flickers at the $\sim 0.2\%$ level,

active on the timescale tens of minutes. As is apparent in Figure 2(c), these flickers have a stronger wavelength dependence at the blue end of the detector, suggesting that they are related to sudden temperature changes of the filament. The spectrum further shows deep and shallow absorption features presumably due to condensed ices along the optical path. Some of the absorption features change somewhat in depth throughout the exposure, with a notable concave-down profile at $3.55\text{--}3.7\ \mu\text{m}$ that imparts an erroneous excess in the extracted transit depths in this channel. The first hour of the time series shows a strong, chromatic increase in flux due to the lamp warming up, and the chromaticity of the flickers is stronger here. We therefore use the last $2^{13} = 8192$ integrations, or ~ 2.1 hr, of data to mitigate the strong spectral systematics and to better distill the detector’s performance. The light source, whose spectrum is far removed from a typical stellar spectrum in shape, has an equivalent J -band magnitude of roughly $J \sim 11$, which is approaching PRISM’s saturation limit.

3. Data Extraction and Reduction

In order to prepare our data for injection-recovery testing, we first extract the spectrophotometry from the time series of the raw frames. The overall approach to reduce and calibrate NIRSpec data has been described by Rawle et al. (2016) and Alves de Oliveira et al. (2018) and is also summarized in Birkmann et al. (2022) and Jakobsen et al. (2022). For each integration, we perform the following steps: We first apply the flat-field correction and save the wavelength solution. We next subtract the $1/f$ noise, as measured by the vertical median of the unilluminated regions of the detector at each pixel column. This process is also referred to as “destriping.” The $1/f$ noise component, whose contribution can be seen in Figure 2(a), imparts a near doubling to the white-noise level to the spectrophotometry if not correctly removed. Next, we measure the vertical and horizontal pixel shifts of the trace in each integration relative to the first by performing a 1D cross-correlation between the vertically and horizontally summed counts (Figure 2(b)). We then remove these shifts and align spectra by interpolating each frame with a third-order spline and regriding the image data onto the shifted pixel array. By aligning the spectra, we remove a common time-series systematic effect where the flux of a pixel is observed to change with time as spectral features drift across a pixel.

Next, we fit a Gaussian profile to the PSF at each pixel column to measure the trace position and width. We smooth the trace profile with a median filter, fit it with a sixth-order polynomial, and perform a subpixel extraction of the counts enclosed within $5\times$ FWHM of the trace center. Lastly, we convert from count rate to counts by multiplying by the total integration time of each frame (0.45 s). At this stage, the flat-fielded, $1/f$ -corrected, shift-stabilized, and spatially summed spectrophotometry resembles the true, distilled behavior of the temperamental light source with most of the confounding artifacts removed. The extracted spectrophotometry can be seen in Figure 2(c), where light-source flickers and flux variations imparted by x shifts across sharp spectral features show up strongly in the relative flux at the $\pm 1\%$ level.

Before injecting transit spectra into the time series, we first apply a common-mode correction to each wavelength pixel in order to suppress the effects of the lamp flicker. This common-mode technique to remove systematic errors is an often-used method to remove time-series trends that are the same or similar at each wavelength (e.g., Berta et al. 2012; Deming et al. 2013;

⁴ <https://www.cosmos.esa.int/web/jwst-nirspec/test-data>

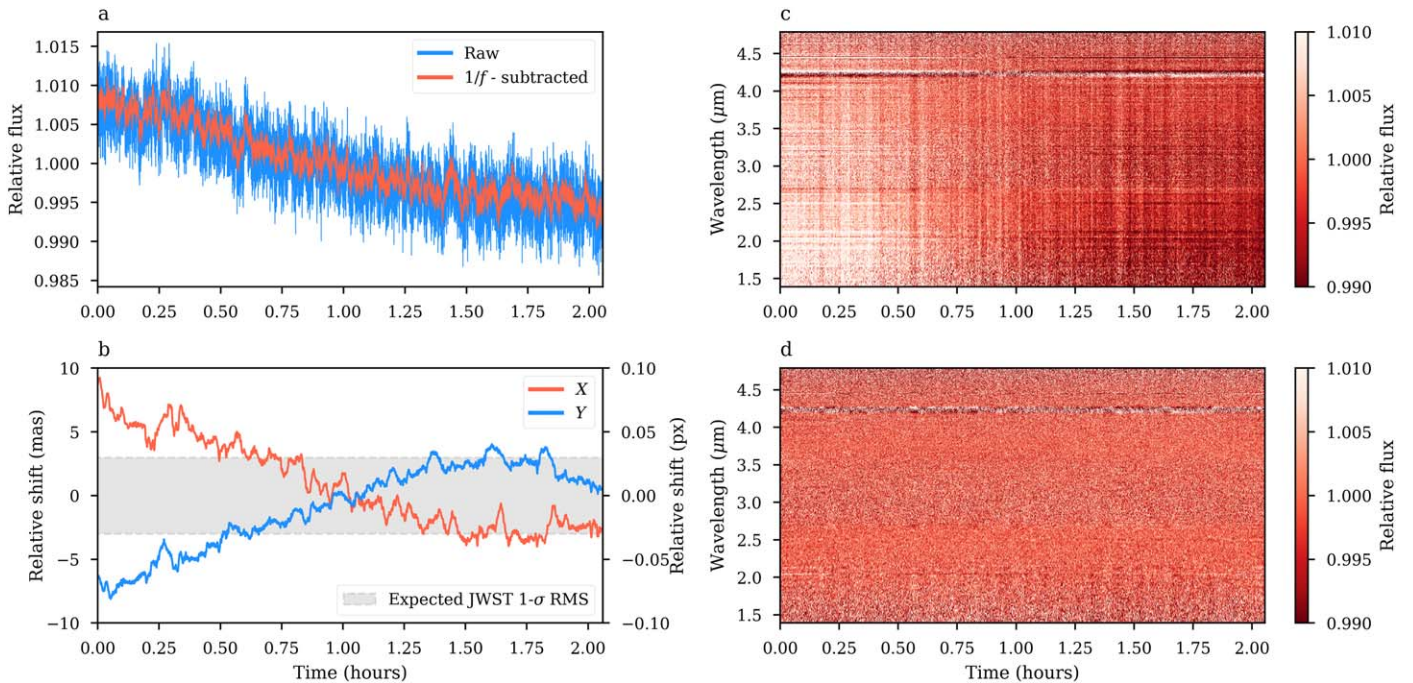


Figure 2. Overview of the NIRSpec PRISM CV3 time-series observation. (a) The mean-normalized white light curve with and without $1/f$ correction. (b) The x and y shifts of the PSF relative to the middle integration. The gray region shows the 1σ rms jitter requirement of JWST over 10^4 s. (c) The raw relative spectrophotometry after shift stabilization, $1/f$ noise subtraction, and extraction. Note the strong light-source flickers and trends. (d) Same as (c), but after a common-mode correction that mostly removes the systematics imparted by the unstable light source.

Sing et al. 2013). While simply dividing the time series of each pixel by the white light curve can remove most of the observed lamp flickers and long-term flux drifts, the amplitude of these variations is found to also have a color dependence, which we remove by applying the common-mode correction on a pixel-by-pixel basis. This is done by performing a linear regression fit of the white light curve (see Figure 2(a)) to the spectral light curves at each wavelength pixel and dividing it out. This correction removes the majority of the light-source flicker while preserving the detector’s noise properties, intrapixel systematics, and most of the light source’s chromaticity. The self-correction effect imposed by the division of the white light curve can only be seen in the postcorrected white light curve, which by definition is a flat line at 1. The common-mode-detrended spectrophotometry is shown in Figure 2(d), where modest residual chromatic light-source systematics are seen. At this stage, the spectrophotometry is largely free from confounding trends and flickers from the light source.

4. Transit Spectrum Injection

Injection-recovery tests are often repeated with several draws from a large data set or noise model, but our study relies on a single real observation and therefore a single noise instance. We perform our tests on two different planet cases, as two separate manifestations of the same noise instance, to show the consistency of our injection-recovery process. We inject the transit spectrophotometry of two planets: the warm Neptune GJ 436 b ($T_{\text{eq}} \sim 700$ K, $J = 6.9$) and the temperate terrestrial planet TRAPPIST-1 d ($T_{\text{eq}} \sim 300$ K, $J = 11.35$). We choose these planets due to their short ≤ 1 hr transit duration and their broad interest to the exoplanet atmosphere community.

GJ 436 b is slated for transmission spectroscopy with NIRSpec G395H (Stevenson et al. 2021), and TRAPPIST-1 d will be observed in transit with the NIRSpec PRISM (Lafreniere 2017). With TRAPPIST-1’s J -band magnitude of

11.4, the count rates from the CV3 light source are similar in our data set to what will be seen on sky with JWST, making it a compelling comparison. GJ 436A, however, is 4 mag brighter (prohibitively bright for PRISM), and the expected precision of its measured transmission spectrum using other modes is significantly higher than what we present. For GJ 436 b, we use a transmission spectrum from the Goyal Grid of generic models (Goyal et al. 2020), with parameters most closely matching those of GJ 436 b. The TRAPPIST-1 d spectrum is taken from the work of Lincowski et al. (2018) and Lustig-Yaeger et al. (2019), where a CO_2 -dominated 10 bar atmosphere is assumed, with high-altitude H_2SO_4 clouds.

For each planet, we bin our model transmission spectra to the wavelength grid of each PRISM pixel and compute four-parameter limb-darkening coefficients between the pixel edges using values from the 3D STAGGER-grid (Magic et al. 2015). These limb-darkening (LD) coefficients are weighted by NIRSpec’s response function and the respective stellar spectrum. Given the narrow wavelength bins used here, we find that the errors in the LD coefficients introduced by using a stellar spectral weighting rather than the lamp’s spectrum are very small. We then compute the transit light-curve time series at each wavelength pixel using *batman* (Kreidberg 2015).

Finally, we inject the spectral transit model while simultaneously dividing out the common-mode correction into the raw data, and reinsert the original time series’ $1/f$ noise and jitter. At this stage, our injected files mimic a JWST observation as closely as possible, with the jitter exaggerated by a factor of 4.

5. Transit Spectrum Recovery

We reduce our transit-injected data files following the procedure described in the first paragraph of Section 3. We bin the transit spectrophotometry into variable-width wavelength channels with

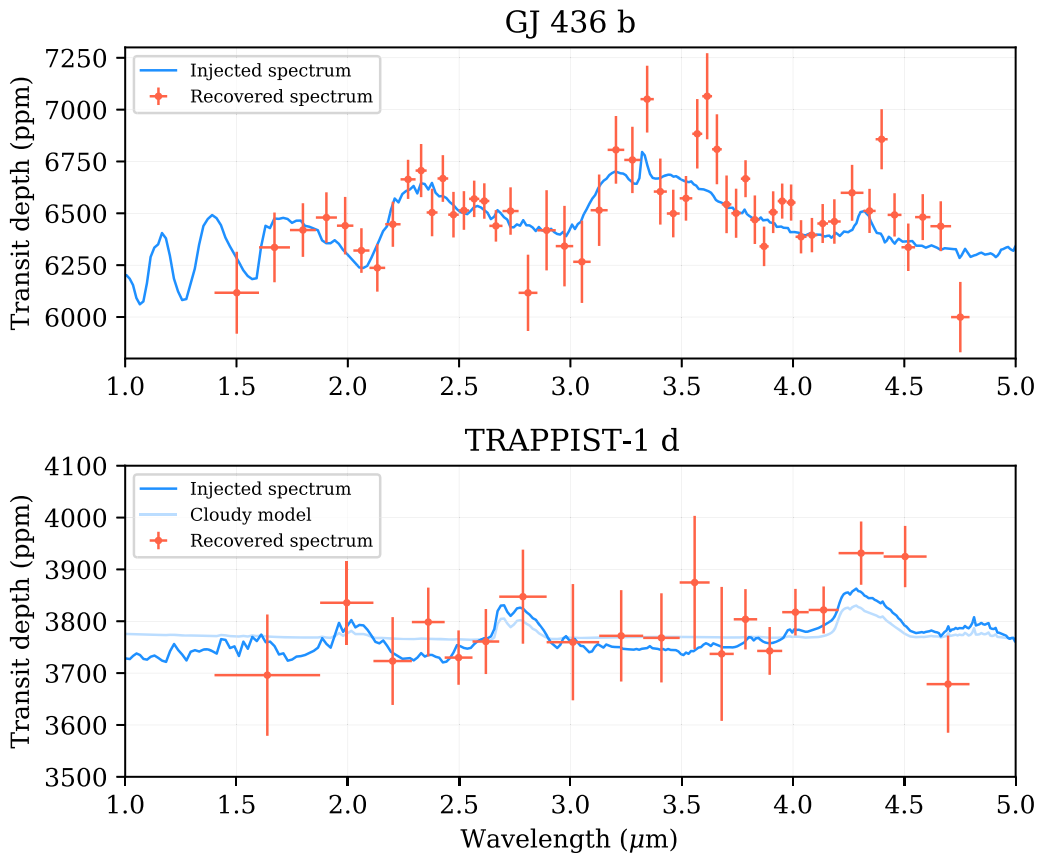


Figure 3. The recovered transmission spectra of GJ 436 b (top) and TRAPPIST-1 d (bottom) recovered from the injected data using our novel pipeline. We include a cloudy TRAPPIST-1 d atmosphere model for an illustrative comparison.

roughly equal counts. We use 51 and 19 wavelength channels for GJ 436 and TRAPPIST-1, respectively (Figure 3).

We model the systematic errors with a parameterized deterministic model, where the photometric trends are found to correlate with a number of external parameters (or optical-state parameters, \mathbf{x}). These parameters describe changes in the instrument or other external factors as a function of time during the observations and are fit with coefficients, p_n , for each optical-state parameter to detrend the light curves. We then fit each spectral light curve jointly with a transit model and systematics model $S(\mathbf{x})$ of the form:

$$S(\mathbf{x}) = p_1 X + p_2 Y + p_3 X^2 + p_4 Y^2 + p_5 XY + p_6 \Delta_{\text{CM}}, \quad (1)$$

where X and Y are the mean-subtracted, standard-deviation-normalized x and y shifts, respectively, and Δ_{CM} is the pre-injection normalized common-mode light curve, normalized in the same way. The full light-curve flux model, $f(t)$, is then assembled with the equation

$$f(t) = f_0 \times T(t, \theta) \times S(\mathbf{x}), \quad (2)$$

where f_0 is the baseline flux and $T(t, \theta)$ is the theoretical transit model that depends on the transit parameters θ . This model captures the behavior of jitter-induced intrapixel variations, as well as leftover light-source variability. The detrending vectors are normalized so that their fit coefficients may be directly compared. We perform joint fits of the systematics model and the planet’s transit depth (while keeping all other transit parameters fixed) using the Python-based Levenburg–Marquardt least-squares algorithm `lmfit` (Newville et al. 2014). Using Python’s

`multiprocessing` package to parallelize the fitting routine, we fit all 51 GJ 436 b channels in 30 s using 16 cores. We initially assume photon-noise uncertainties for the light curve, refit with the residual standard deviation as the uncertainty, and then inflate the final `lmfit` parameter uncertainties on the transit depths by a factor $\beta(\lambda)$. $\beta(\lambda)$ is computed from the fit residuals of each spectral light curve taking into account both white- and red-noise sources. Its median value is 1.35, and it peaks where there are residual systematics due to chromatic lamp flickers and changing cryogenic ice features. $\beta(\lambda)$ is near 1 in wavelength regions where these two effects are small. If $\beta(\lambda)$ were not multiplied into our transit-depth uncertainties, the residual chromatic flicker effects would not be taken into account, giving unrealistic error bars. Further, we avoid removing these residuals using stochastic models such as Gaussian Processes (GPs) because they could remove systematics due to the detector’s noise properties, which we aim to characterize.

We also use the nested sampling package `dynesty` to fit the light curves (Speagle 2020) and find equivalent results in both the fitted depth and the uncertainties. Moreover, we find that there are no complicated degeneracies between the transit depth and the optical-state vectors, indicating that the posterior is well described by a multivariate Gaussian distribution. This means that a Levenburg–Marquardt least-squares minimizer is sufficient to quickly estimate the posterior distribution with accurate uncertainties.

This fit to the transmission spectra of both planets, which we term the “fiducial” fit hereafter, match the injected spectra well, with a reduced chi-squared χ_ν^2 of 1.32 for GJ 436 b and 1.15 for TRAPPIST-1 d (see Section 5). In the latter case, the extracted

Table 1
Systematics Fitting Results Summary

Scenario	Description	N_{params}	$\sigma_{\text{TS resid}}$ (ppm)	χ_{ν}^2	ΔBIC	ΔAIC
1	Fiducial	8	148.70	1.32
2	No X^2	7	148.74	1.25	-6.15	-2.48
3	No XY	7	149.30	1.31	-5.01	-1.35
4	No Y^2	7	151.21	1.28	-4.87	-1.20
5	No X^2, Y^2	6	152.38	1.28	-7.37	-0.08
6	No Y	7	171.28	2.32	39.09	42.75
7	No X	7	184.86	2.25	36.15	39.85
8	Fiducial with X^3, Y^3	10	243.70	3.36	88.69	81.27
9	No X^2, Y^2, XY	5	267.5	4.33	130.17	141.02
10	No X^2, Y^2, XY, X	4	285.93	4.84	154.76	169.10
11	No X^2, Y^2, XY, Y	4	290.39	6.05	211.69	226.03
12	No X^2, Y^2, XY, X, Y	3	293.37	5.66	195.09	212.79
13	Fiducial, no flat field	8	390.98	7.12	248.97	248.97
14	Fiducial, no $1/f$ subtraction	8	N/A	555.69	23837.78	223837.78

Note. Results from a suite of tests for goodness of fit and systematics model quality sorted in decreasing order of residual scatter. Each row represents an independent transmission spectrum extraction from our injected transit spectrophotometry of GJ 436 b, fitted with some linear combination of first-, second-, and third-order X and Y jitter vectors, and the cross-term XY . σ_{resid} is the standard deviation of the residual between the injected and recovered transmission spectrum, and χ_{ν}^2 is the reduced- χ^2 statistic of the fit. The scatter in Scenario 14 is not reported because the noise level was too overwhelming for the fit to prefer a nonzero transit depth.

spectrum qualitatively shows increases in transit depth at the three main absorption features. The χ_{ν}^2 for GJ 436 b’s fiducial fit without taking into account error bar inflation is 3.13. The residuals are largely Gaussian, with the notable exception of the peaky $3.6 \mu\text{m}$ feature, which is imparted by a downward concavity in the ice-deposition profile there. This same feature can be seen in Figure A.2, panel (f), of Birkmann et al. (2022). The average amplitude of the x - and y -shift vectors are found to be 455 and 445 ppm, respectively, across both spectra. This is in excellent agreement with the jitter systematic noise modeled by Birkmann et al. (2022). Rescaling these amplitudes linearly by a factor proportional to the expected jitter of JWST, we estimate the on-sky jitter-induced photometric signatures to be at the ~ 100 ppm level.

6. Systematics Model Intercomparison

We explore the relevance of each term in the systematics model, as well as the flat-field and $1/f$ subtraction, by performing a model intercomparison on the GJ 436 b data. We refit the transit spectrophotometry using the methods described in Section 5 but with key elements added or removed to see how the goodness of fit of the final transmission spectrum changes. We recommend this technique as an informative method for systematics model selection and to explore the response of the final spectrum to different procedures in the data reduction. At each scenario, we compute the standard deviation of the fit residual, the reduced chi-squared statistic χ_{ν}^2 , and the change in Bayesian and Akaike inference criteria (ΔBIC and ΔAIC , respectively) relative to the fiducial model. We note that in each instance, we keep the uncertainties on the transit depths fixed to those of the fiducial fit in order to avoid erroneously inflating the error bars due to poor model selection and to preserve the uniformity of the comparison. The results of this intercomparison are in Table 1.

Decorrelation against the linear jitter vectors X , Y , and XY all substantially improve goodness of fit over a model without them. The second-order terms X^2 and Y^2 are marginally, if at all, important, as evidenced by the negative changes to their Bayesian and Akaike information criteria relative to the fiducial

model (ΔBIC and ΔAIC , respectively). This indicates that the intrapixel sensitivity profile of the typical NIRSspec pixel is well approximated by a low-order polynomial. Adding the third-order terms X^3 and Y^3 to the systematics model results in a much poorer goodness of fit, likely due to the even nature of the detector’s intrapixel sensitivity function.

Flat fielding has a major impact on the quality of the resulting fit, nearly halving the resulting residual scatter when the flat field is applied. From this large effect, we determine that significant efforts should be made to obtain several, accurate flat fields for the subarrays used in JWST BOTS data. Subpixel flat fields could help further reduce the impact of jitter on the photometry. Perhaps most important, though, is the $1/f$ correction—this added noise component overwhelms the transit depths, resulting in nondetections of the transit feature in all spectroscopic channels when not removed.

A number of systematic sources that are not explored by our work, or potentially fully captured by this data set, are likely to be present in on-sky JWST data. Birkmann et al. (2022) discuss many potential sources of systematic sources in more depth, including charge-trapping-induced persistence in the detector, which results in exponential “ramps” in measured flux as is regularly seen in HST/WFC3 data (e.g., Deming et al. 2013; Wakeford et al. 2013; Kreidberg et al. 2014; Zhou et al. 2017). This effect, which cannot be measured in the CV3 data due to the instability of the lamp, is expected to be greatly reduced for the relatively new JWST detectors (Rauscher et al. 2014; Birkmann et al. 2022). High-gain antenna moves, increased cosmic rays, and thermal drift of the telescope may also introduce systematics and are not captured in the CV3 data. Birkmann et al. (2022) estimates that slit losses and intrapixel sensitivity variations will both be on the order of several hundred ppm with BOTS data, pending JWST’s real on-sky PSF profile and pointing stability. The CV3 data displays these trends at the expected levels given the X - and Y -position dependence of the spectrophotometry, which are well decorrelated by our jitter detrending technique.

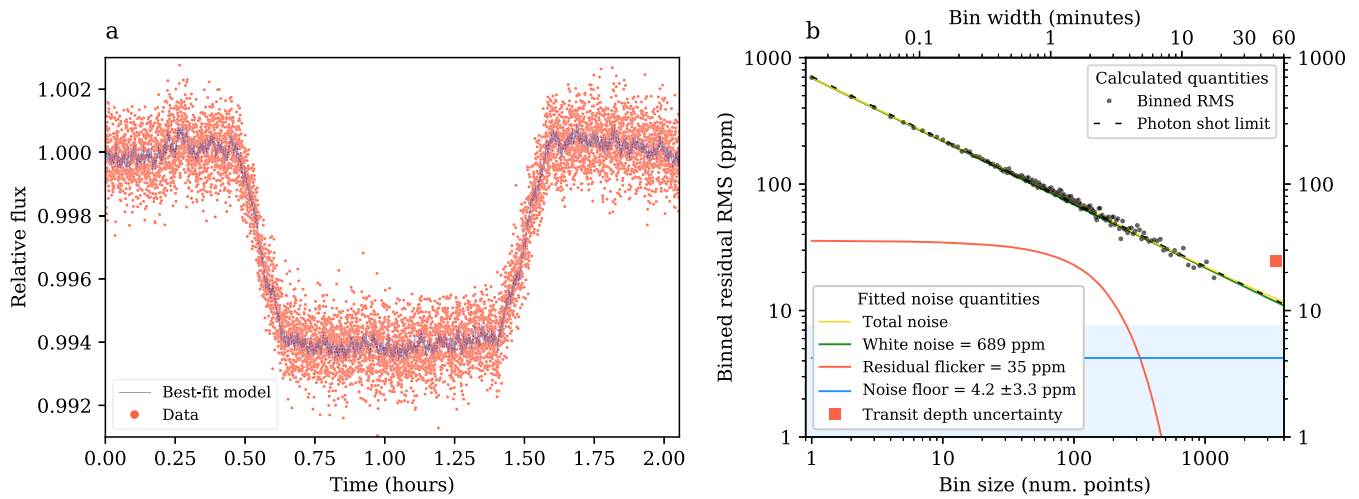


Figure 4. (a) The injected transit light-curve photometry of GJ 436 b in the relatively systematics-free $3.77\text{--}4.17\ \mu\text{m}$ channel (tomato points) and the best-fitting joint model (gray line). (b) Residual scatter as a function of bin size for the $3.77\text{--}4.17\ \mu\text{m}$ channel (black points) with the best-fitting noise components (solid lines) and the photon shot noise limit (dashed line).

7. Searching for the Noise Floor

To search for the detector’s noise floor, we first select a wide wavelength region that is largely free of systematics and light-source spectral features. A bin sufficiently narrow to minimize the enclosed chromatic systematics, and wide enough to encapsulate maximal counts, is ideal. The relatively clean $3.77\text{--}4.17\ \mu\text{m}$ region is chosen for these reasons. We then fit the light curve of GJ 436 b at this channel as described previously and study its residual by binning it through time (see Figure 4). The residual rms bins down closely with $1/\sqrt{N}$ (where N is the number of points per bin). We fit this rms curve with a quadrature sum of the white-noise component, a noise floor, and a logistic decay law to capture the leftover chromatic flickers. We find that the fitted white noise at this bin closely matches the expected photon shot limit and that the residual flicker noise contributes 35 ppm at the ~ 2 minute timescale, which is consistent with the flicker timescale seen in the white light curve. The rms curve bins to ~ 15 ppm. Most importantly, we recover a 3σ upper bound of 14 ppm for the noise floor, with the fit preferring a < 10 ppm floor at 1.7σ . Given the joint probability distribution with our systematics model, the transit-depth uncertainty at this channel is $1.35\times$ larger than what is expected due to photon noise alone, or 74% of the photon-noise limit. We expect on-sky JWST, with its expected weaker systematics, to reach closer to the photon limit, making our value somewhat pessimistic.

8. Summary and Conclusion

In this work, we analyzed NIRSpec lab time-series data, injecting and recovering transmission spectra including real detector noise, making it the closest analog available before science data are collected. The data set, complete with the detector’s actual noise performance and intrapixel sensitivity variations, informed our optimistic expectations as future JWST users. We demonstrate that subtraction of the $1/f$ noise is a crucial step in the data reduction and caution that this step may be complicated in the narrower SUB512S subarray, which is only 16 pixels tall. Further, we show that ignoring the flat field increases the scatter of the final spectrum by a factor of 2.6 (Table 1). Given JWST’s 7 mas (0.07 pixel) on-sky stability requirement we expect the amplitude of jitter-induced systematics to be at the ~ 100 ppm

level. We demonstrate that such jitter systematics can be readily detrended using a low-order polynomial with detector position, binning the data down close to the photon limit. While we do not confidently detect a noise floor, we show that it is < 14 ppm and likely below 10 ppm. Using our fast reduction and analysis pipeline, we successfully recover our injected transmission spectra and show that Levenburg–Marquardt least-squares fitters are sufficient to quickly estimate the posterior distribution. We also find that, when simulating observations, transit-depth uncertainties should be scaled conservatively by a factor of 1.35 relative to the photon-noise limit expectation. All of these facts considered, we believe NIRSpec PRISM is well suited to characterize challenging exoplanet atmosphere targets with ~ 10 ppm spectral amplitudes. We expect that the precision of such spectra will be limited by the photon limit and confounding astrophysical signatures such as stellar activity rather than instrumental systematics.

DKS acknowledges support from the STScI INS division. RL and AW were supported through the JHU BDP Summer Program. We thank Giovanna Giardino, Stephan Birkmann, and Pierre Ferruit for their useful discussions and the NIRSpec team for their efforts in collecting and sharing the CV3 data set.

Facility: JWST (NIRSpec).

Software: Matplotlib (Hunter 2007), Numpy (van der Walt et al. 2011), Scipy (Virtanen et al. 2020), astropy (Astropy Collaboration et al. 2013), dynesty (Speagle 2020), lmfit (Newville et al. 2014).

ORCID iDs

Zafar Rustamkulov <https://orcid.org/0000-0003-4408-0463>
 David K. Sing <https://orcid.org/0000-0001-6050-7645>
 Rongrong Liu <https://orcid.org/0000-0003-0685-3525>
 Ashley Wang <https://orcid.org/0000-0001-6542-999X>

References

- Alves de Oliveira, C., Birkmann, S. M., Böker, T., et al. 2018, *Proc. SPIE*, 10704, 107040Q
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33
- Batalha, N. E., & Line, M. R. 2017, *AJ*, 153, 151
- Berta, Z. K., Charbonneau, D., Désert, J.-M., et al. 2012, *ApJ*, 747, 35

- Birkmann, S. M., Ferruit, P., Giardino, G., et al. 2022, arXiv:2202.03309
- Deming, D., Knutson, H., Kammer, J., et al. 2015, *ApJ*, **805**, 132
- Deming, D., Wilkins, A., McCullough, P., et al. 2013, *ApJ*, **774**, 95
- Goyal, J. M., Mayne, N., Drummond, B., et al. 2020, *MNRAS*, **498**, 4680
- Hobson, M. J., Brahm, R., Jordán, A., et al. 2021, *AJ*, **161**, 235
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jakobsen, P., Ferruit, P., de Oliveira, C. A., et al. 2022, arXiv:2202.03305
- Jenkins, J. S., Díaz, M. R., Kurtovic, N. T., et al. 2020, *NatAs*, **1**, 1148
- Kreidberg, L. 2015, *PASP*, **127**, 1161
- Kreidberg, L., Bean, J. L., Désert, J.-M., et al. 2014, *Natur*, **505**, 69
- Lafreniere, David 2017, JWST Proposal, Cycle 1, ID #1201
- Lincowski, A. P., Meadows, V. S., Crisp, D., et al. 2018, *ApJ*, **867**, 76
- Luger, R., Agol, E., Kruse, E., et al. 2016, *AJ*, **152**, 100
- Luger, R., Kruse, E., Foreman-Mackey, D., Agol, E., & Saunders, N. 2018, *AJ*, **156**, 99
- Lustig-Yaeger, J., Meadows, V. S., & Lincowski, A. P. 2019, *AJ*, **158**, 27
- Magic, Z., Chiavassa, A., Collet, R., & Asplund, M. 2015, *A&A*, **573**, A90
- May, E. M., & Stevenson, K. B. 2020, *AJ*, **160**, 140
- Morley, C. V., Kreidberg, L., Rustamkulov, Z., Robinson, T., & Fortney, J. J. 2017, *ApJ*, **850**, 121
- Newville, M., Stensitzki, T., Allen, D. B., & Ingargiola, A. 2014, LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python (0.8.0), Zenodo, doi:10.5281/zenodo.11813
- Pontoppidan, K. M., Pickering, T. E., Laidler, V. G., et al. 2016, *Proc. SPIE*, **9910**, 991016
- Rauscher, B. J., Boehm, N., Cagiano, S., et al. 2014, *PASP*, **126**, 739
- Rawle, T. D., Alves de Oliveira, C., Birkmann, S. M., et al. 2016, *Proc. SPIE*, **9904**, 990446
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *JATIS*, **1**, 014003
- Schlawin, E., Leisenring, J., McElwain, M. W., et al. 2021, *AJ*, **161**, 115
- Schlawin, E., Leisenring, J., Misselt, K., et al. 2020, *AJ*, **160**, 231
- Sing, D. K., Lavvas, P., Ballester, G. E., et al. 2019, *AJ*, **158**, 91
- Sing, D. K., Lecavelier des Etangs, A., Fortney, J. J., et al. 2013, *MNRAS*, **436**, 2956
- Speagle, J. S. 2020, *MNRAS*, **493**, 3132
- Stevenson, Kevin, Lustig-Yaeger, Jacob, Munazza, Alam, et al. 2021, JWST Proposal, Cycle 1, ID #1981
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Vanderspek, R., Huang, C. X., Vanderburg, A., et al. 2019, *ApJL*, **871**, L24
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nat. Methods*, **17**, 261
- Wakeford, H. R., Sing, D. K., Deming, D., et al. 2013, *MNRAS*, **435**, 3481
- Winters, J. G., Medina, A. A., Irwin, J. M., et al. 2019, *AJ*, **158**, 152
- Zhou, G., Bakos, G. Á., Hartman, J. D., et al. 2017, *AJ*, **153**, 211