



An Extreme-precision Radial-velocity Pipeline: First Radial Velocities from EXPRES

Ryan R. Petersburg^{1,2} , J. M. Joel Ong (王加冕)² , Lily L. Zhao² , Ryan T. Blackman² , John M. Brewer^{2,3} ,
Lars A. Buchhave⁴ , Samuel H. C. Cabot² , Allen B. Davis² , Colby A. Jurgenson^{2,5} , Christopher Leet² ,
Tyler M. McCracken^{2,6} , David Sawyer^{2,7} , Mikhail Sharov⁸ , René Tronsgaard⁴ , Andrew E. Szymkowiak^{1,2} , and
Debra A. Fischer² 

¹ Department of Physics, Yale University, 217 Prospect Street, New Haven, CT 06511, USA

² Department of Astronomy, Yale University, 52 Hillhouse Avenue, New Haven, CT 06511, USA; debra.fischer@yale.edu

³ Department of Physics & Astronomy, San Francisco State University, 1600 Holloway Avenue, San Francisco, CA 94132, USA

⁴ DTU Space, National Space Institute, Technical University of Denmark, Elektrovej 328, DK-2800 Kgs. Lyngby, Denmark

⁵ Department of Astronomy, The Ohio State University, 4055 McPherson Laboratory, 140 West 18th Avenue, Columbus, OH 43210, USA

⁶ Ball Aerospace and Technologies Corporation, 1600 Commerce Street, Boulder, CO 80301, USA

⁷ Lowell Observatory, 1400 West Mars Hill Road, Flagstaff, AZ 86001, USA

⁸ Department of Astronomy, Boston University, 725 Commonwealth Avenue, Boston, MA 02215, USA

Received 2019 September 3; revised 2020 February 10; accepted 2020 March 7; published 2020 April 2

Abstract

The EXtreme-PREcision Spectrograph (EXPRES) is an environmentally stabilized, fiber-fed, $R = 137,500$, optical spectrograph. It was recently commissioned at the 4.3 m Lowell Discovery Telescope near Flagstaff, Arizona. The spectrograph was designed with a target radial-velocity (RV) precision of 30 cm s^{-1} . In addition to instrumental innovations, the EXPRES pipeline, presented here, is the first on-sky, optical, fiber-fed spectrograph to employ many novel techniques—including an “extended flat” fiber used for wavelength-dependent quantum efficiency characterization of the CCD, a flat-relative optimal extraction algorithm, chromatic barycentric corrections, chromatic calibration offsets, and an ultra-precise laser frequency comb for wavelength calibration. We describe the reduction, calibration, and RV analysis pipeline used for EXPRES and present an example of our current sub-meter-per-second RV measurement precision, which reaches a formal, single-measurement error of 0.3 m s^{-1} for an observation with a per-pixel signal-to-noise ratio of 250. These velocities yield an orbital solution on the known exoplanet host 51 Peg that matches literature values with a residual rms of 0.895 m s^{-1} .

Unified Astronomy Thesaurus concepts: [High-resolution spectroscopy \(2096\)](#); [Radial velocity \(1332\)](#); [Astronomy data analysis \(1858\)](#); [Observational astronomy \(1145\)](#)

Supporting material: machine-readable table

1. Introduction

Results from the NASA Kepler mission show that small planets with radii between 1 and $4 R_{\oplus}$ are found orbiting 20%–50% of main-sequence stars (Winn & Fabrycky 2015). While such transit surveys, such as Kepler, K2, and TESS, have revealed a wealth of planets, few of these planets have had their masses measured; mass estimates that do exist are typically radial-velocity (RV), dynamical masses. More precise RV measurements are required to determine mass estimates for these planets, particularly small rocky ones, than are possible with pre-existing RV spectrographs. Were they available, these mass measurements would shed light on planetary structure, bulk density, and the mass–radius relation for sub-Neptune-mass planets.

To meet these needs, a new generation of RV spectrographs is now emerging: the Echelle SPectrograph for Rocky Exoplanets Search and Stable Spectroscopic Observations (ESPRESSO: Pepe et al. 2013) and the EXtreme-PREcision Spectrograph (EXPRES: Jurgenson et al. 2016) are now on-sky taking data, while the NN-explOre Exoplanet Investigations with Doppler spectroscopy spectrograph (Schwab et al. 2016) is in the commissioning phase. These new extreme-precision radial-velocity (EPRV) spectrographs are driving toward the sub- 10 cm s^{-1} RV precision needed to detect a true Earth twin around a Sun-like star.

Extremely precise spectrographs, in turn, demand extremely precise data reduction pipelines. The fidelity of the data and

error estimates returned by these data reduction pipelines is paramount if groundbreaking discoveries returned from such instruments are to be credible. To this end, a complete science reduction, extraction, and analysis pipeline was newly developed and tailored for EXPRES data. In what follows, we begin with a brief description of the instrument and the calibration strategy. We then describe the analysis performed by our pipeline on EXPRES data, and present the first radial-velocity measurements from the instrument.

2. Instrument Description

EXPRES is a fiber-fed, white-pupil EPRV spectrograph with a design resolution of $R = 150,000$ and a wavelength range of 3800–7800 Å. In practice, Blackman et al. (2020) show that the median resolution is better characterized as $R \sim 137,500$ with a maximum resolution reaching $R \sim 150,000$ in some regions of the detector. EXPRES is environmentally stabilized in a vacuum enclosure and is situated at Lowell Observatory’s 4.3 m Lowell Discovery Telescope (LDT) near Flagstaff, Arizona. The multi-instrument port configuration of the LDT allows for high-cadence, flexible scheduling of stars (up to 280 partial nights per year).

Owing to various changes to the hardware configuration during commissioning, we have divided radial velocities from EXPRES into different calibration epochs, with an independent RV offset fitted for each epoch. Each epoch demarcates changes introduced to the configuration of the echellogram on

Table 1
Instrumental Epochs

Epoch	Start	End	Changes before Epoch Start
0	...	2018 Apr 15	Commissioning
1	2018 Apr 15	2018 Jun 15	Increased CCD pre-settle time
2	2018 Jun 15	2018 Nov 8	Fiber change; CCD rotation
3	2018 Nov 8	2019 Feb 7	Original science fiber replaced; Calibration unit rebuilt
4	2019 Feb 7	2019 Aug 4	LFC beat frequency mitigated; Fiber agitator repaired
5	2019 Aug 4	...	Realignment of FEM; Replacement of LFC PCF

the CCD, the shape of the point-spread function (PSF), or the stability of our calibration sources. A summary of the changes that delineate the beginning of each epoch can be found in Table 1.

The inputs to EXPRES are a $33 \times 132 \mu\text{m}$ science fiber as well as a $60 \times 180 \mu\text{m}$ extended fiber. The extended fiber is wider than the science fiber in the cross-dispersion direction, allowing for higher signal-to-noise ratio (S/N) flat-fielding for all pixels illuminated by light from the science fiber, particularly at the cross-dispersion edges. Both fibers are also supplemented by their own square simultaneous fibers, though these are not used in normal operation.

The flat-field light source is provided by a custom solution: a collection of 25 LEDs integrated on a single compact chip. The wavelength range of the LEDs cover the entire bandwidth of EXPRES, and the relative power of each LED is tuned to approximately match the inverse response of the spectrograph. Light from these LEDs, averaging 12.5 W, is coupled into an integrating sphere and subsequently injected into both the extended and science fibers.

Wavelength calibration is carried out with a Menlo Systems laser frequency comb (LFC; similar to those in Steinmetz et al. 2008; Probst et al. 2016). Three LFC exposures are taken through the science fiber every 15–30 minutes throughout a night. Two ThAr exposures are taken each night—once at the beginning and once at the end. The LFC is set to standby between each set of calibrations to reduce wear on nonlinear optical elements, and is only turned on approximately 1 minute before the next exposure is needed, to suppress turn-on transients.

Barycentric corrections are derived from the EXPRES exposure meter, a $R \approx 100$ spectrograph with an EMCCD detector and a bandpass that covers the spectral range of the LFC. During each science observation, the exposure meter takes a continuous series of 1 s exposures. Further technical details can be found in Blackman et al. (2019).

3. Analysis of EXPRES Data

The EXPRES pipeline is written in Python and makes heavy use of the SciPy stack (Virtanen et al. 2020). We show a schematic representation of it in Figure 1. In brief, the following steps are taken:

1. At the start of each calibration epoch, several hundred extended flat images are taken. These are used to construct a master extended flat-field image, which is

divided out from all exposures in the corresponding calibration epoch.

2. Each night, 30 dark and 30 science flat images are taken. They are used to reduce and extract the science frames taken the same night.
3. Echellogram orders are traced using the reduced science flats, a scattered light model is removed, and a flat-relative optimal extraction is performed.
4. Wavelength solutions are interpolated for all science frames using bracketed LFC exposures, seeded by a nightly Thorium Argon source, as a calibration reference.
5. Telluric lines are identified empirically with SELENITE (Leet et al. 2019) and labeled for later analysis.
6. For exposures marked for RV analysis, we obtain radial velocities with two methods: cross-correlation against a line mask and a forward model.

We describe each of these steps in some detail in the following sections.

3.1. Reduction

The EXPRES detector is an STA1600LN CCD backside-illuminated image sensor with a $10,560 \times 10,560$ array containing $9 \mu\text{m} \times 9 \mu\text{m}$ pixels. The CCD is divided into 16 equal 5280×1320 pixel sections (two rows of eight), each with their own independent output amplifier and, thus, corresponding gain (see Table 2). Further information about the EXPRES CCD can be found in Blackman et al. (2020).

In the context of the EXPRES pipeline, “reduction” refers to the conversion of these 16 independent regions read in analog-to-digital units (ADU) to a full two-dimensional frame in units of photoelectrons. The reduction steps are as follows, dependent on the type of image being reduced (science, dark, science flat, or extended flat):

1. Subtract a bias frame constructed from the overscan regions (all).
2. Multiply each amplifier region by the corresponding gain coefficient (all).
3. Median combine calibration frames (dark, science flat, extended flat).
4. Subtract the reduced dark image (science, science flat, extended flat).
5. Divide by the reduced master extended flat (science, science flat).
6. Approximate the noise model using photon (Poisson) and read noise (science, science flat).
7. Trace the echelle orders (science flat).
8. Approximate the scattered light using a two-dimensional b-spline model (science, science flat).

In the following subsections, we go into detail about each of the above steps. Information about the reduction of exposure meter data can be found in Blackman et al. (2019).

3.1.1. Overscan

Each of the EXPRES CCD amplifier regions have overscans along both the serial (horizontal) and parallel (vertical) registers. The serial overscan is 5300×180 pixels along the right side of each amplifier region, and the parallel overscan is 20×1320 pixels along the center line of the CCD (at the bottom of the upper regions and the top of the lower regions). Because these overscans contain virtual pixels read out by the

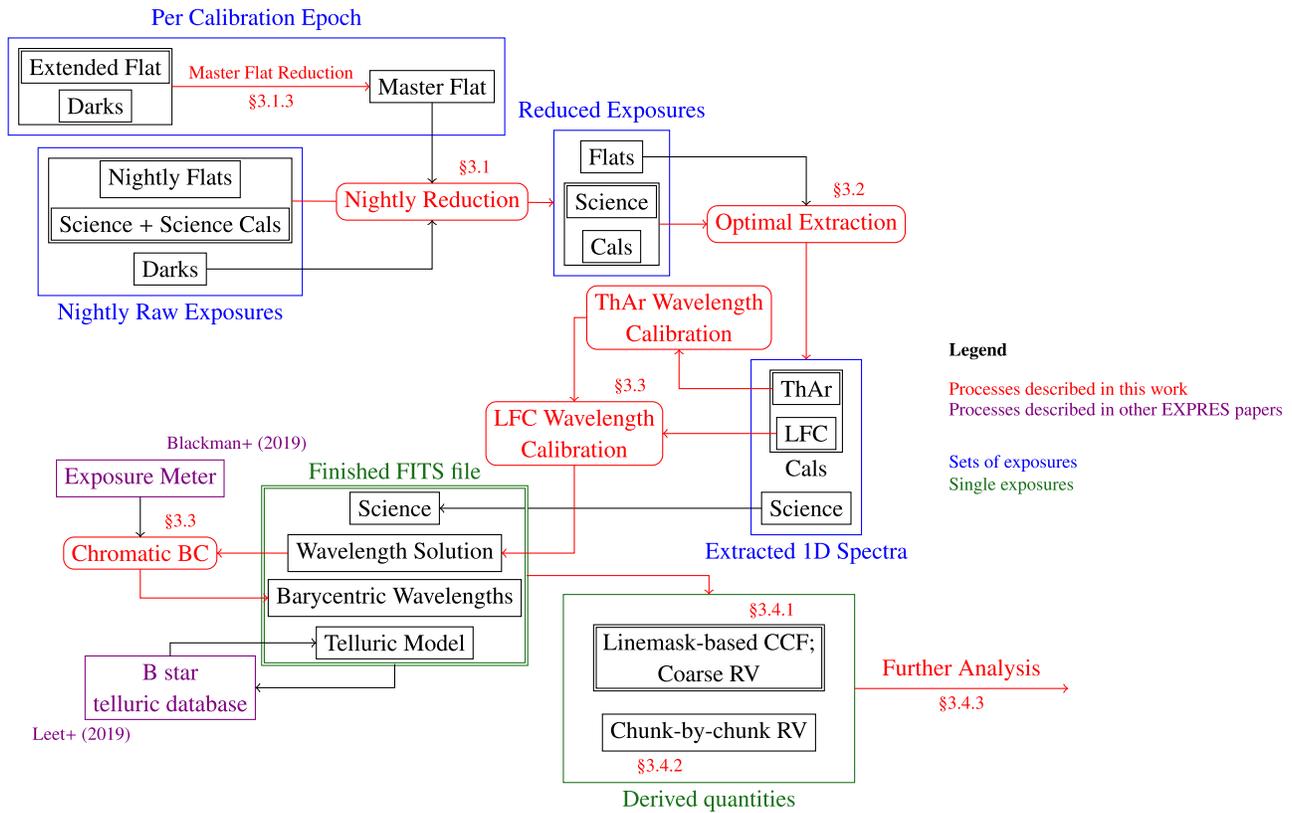


Figure 1. Data flow for the EXPRES pipeline. Small black boxes represent different kinds of exposures (if in blue boxes) or different data associated with an exposure (if in green boxes). Double-boxes indicate finished data products, ready for use in subsequent science analysis.

Table 2
EXPRES CCD Gain

Amplifier	Empirical	STA	Amplifier	Empirical	STA
0	2.57980	2.6645	8	2.69787	2.6945
1	2.55171	2.5352	9	2.65100	2.6422
2	2.53844	2.5218	10	2.64354	2.6367
3	2.52444	2.4065	11	2.60344	2.5502
4	2.51480	2.6024	12	2.60497	2.5571
5	2.49382	2.5686	13	2.59183	2.5630
6	2.52745	2.4816	14	2.62881	2.5691
7	2.53478	2.4960	15	2.72938	2.6111

same electronics as the real amplifier region, the EXPRES reduction pipeline uses them to approximate the bias of the CCD. The process is as follows:

1. Calculate the mean of the serial overscan region along its horizontal axis.
2. Smooth this mean using a cubic b-spline with knots every ~ 100 pixels.
3. Correct the rows in the parallel overscan region by subtracting the overlapping smoothed serial overscan region.
4. Calculate the mean of the parallel overscan region along its vertical axis.
5. Smooth this mean using a cubic b-spline with knots every ~ 100 pixels.
6. Construct a bias for each pixel in the amplifier region by summing the corresponding row from the serial overscan and column from the parallel overscan.

The mean and subsequent spline fit for the overscan regions of a single amplifier region are shown in Figure 2. This process is executed for all exposures, including darks.

3.1.2. Gain

The independent amplifier gains for the EXPRES CCD (Table 2) were determined empirically by matching the edges of neighboring amplifier regions based on stacked bias-subtracted extended flat images. We median combine 20 columns along each edge of two adjacent amplifier regions and determine the factor that minimizes the difference between them. We repeat this process along each row of amplifier regions, yielding a set of relative corrections for each row. Then, using the 20 median-combined rows along the center line of the detector, we determine the single factor that relates the top row of corrected amplifier regions to the bottom row.

Since these corrections are merely relative, we assume that the mean gain of the 16 amplifier regions given by the manufacturer is approximately correct. Thus, we match the mean of the empirically determined gain corrections to that given by the manufacturer. We note that the gains given by the manufacturer were not sufficient at matching the boundary conditions of the EXPRES CCD; some gain corrections were tweaked by more than 3%.

We were not able to calculate the gains in the typical manner—relating the variance of each pixel to its mean—because our flat-fielding LED source has an intrinsic variability that cannot be modeled by photon noise alone. See Blackman et al. (2020) for further details.

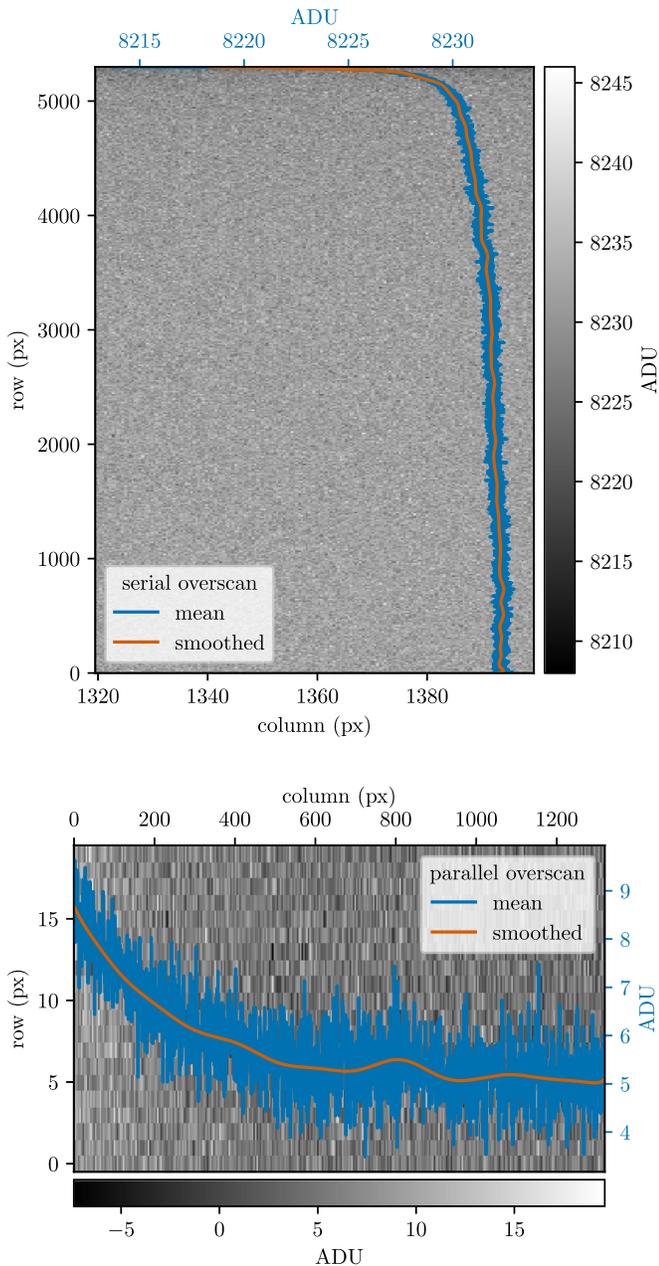


Figure 2. Bias approximation using the serial and parallel overscans. The mean and subsequent smoothed mean for each overscan are shown in blue and orange, respectively. Note that the counts shown in the parallel overscan were first subtracted by the first 20 rows of the serial overscan smoothed mean, generating values close to zero.

3.1.3. Master Extended Flat

In order to measure the quantum efficiency (QE) variations with high signal across the relevant areas of the CCD, EXPRES employs an “extended flat” fiber. This rectangular fiber has slightly larger dimensions of $60 \times 180 \mu\text{m}$ —as compared to the “science” rectangular fiber with dimensions $33 \times 132 \mu\text{m}$ —and is aligned with the center of the science fiber using a motor-driven mirror. Periodically throughout the operation of EXPRES, and at least once per epoch (Table 1), the flat-fielding LED—typically used for order tracing and optimal extraction slit modeling—is injected into the extended flat fiber, and a series of more than 100 images are taken of the resultant spectrum.

Using a median combination of these extended flat images, we construct a master extended flat-field image by dividing out a smooth fit to its echellogram. For each column of each order, we fit a parametric slit function: the squared convolution of a rectangle function with a Gaussian function, which has the analytic form

$$P_{x,y}(A, d, \sigma) = A \left[\Phi\left(\frac{y + d/2}{\sqrt{2}\sigma}\right) - \Phi\left(\frac{y - d/2}{\sqrt{2}\sigma}\right) \right]^2, \quad (1)$$

where Φ is the error function, d is the width of the rectangle, σ is the standard deviation of the Gaussian, and A is the amplitude.

This choice of functional PSF form was motivated by the physical nature of EXPRES and Fourier optics approximations. The input near field of the spectrograph is a rectangle function generated by the rectangular fiber. Taking the Fourier transform of the near field yields an approximation for the far field, which is subsequently morphed as it travels along the optical path of the instrument. We approximate this morph as a Gaussian function. Therefore, the resultant near field that is captured by the detector is the inverse Fourier transform of the modified far field. Invoking the convolution theorem, this detected near field is simply the convolution of the input rectangular function with a Gaussian function, which we subsequently square to approximate the total energy of the light as it hits the detector, yielding Equation (1).

Finally, we smooth the best-fit values for A , d , and σ along each order using a cubic spline with 10 equally spaced knots. The parameters from this smooth fit yield the profile with which we can divide the original extended flat image to generate the master extended flat-field template.

We only use pixels along the approximately flat portion of the extended slit in this constructed smooth profile. This decision was made for two reasons:

1. There is less signal along the top and bottom edges of the slit; thus there is inherently more scatter in the flat corrections for these pixels.
2. Even with the analytic function that we use, the steep edges of the slit function are difficult to fit, which leads to systematic problems in the master extended flat.

Since the slit function of the extended flat was designed to be only 50% larger than the science fiber slit function, the range of pixels covered by the flat portion of the extended slit (as shown in Figure 3) does not correct the entire cross-dispersion profile of the science order. Thus, we also change the motor position of the extended flat injection mirror, which moves the extended flat slit function along the cross-dispersion direction of the echellogram. We thereby expand the number of pixels included in the master extended flat by generating a master extended flat image for each mirror position and then mean-combining these images. This process is completed for each epoch.

3.1.4. Noise Model

The two largest contributions to the noise model for any given pixel on the EXPRES CCD are photon noise and read noise, where these two quantities are measured and summed in quadrature for each pixel. Photon noise is assumed to be Poisson, such that the standard deviation is equal to the square root of the photoelectron counts. Read noise, on the other hand, is calculated empirically for each amplifier. First, the standard

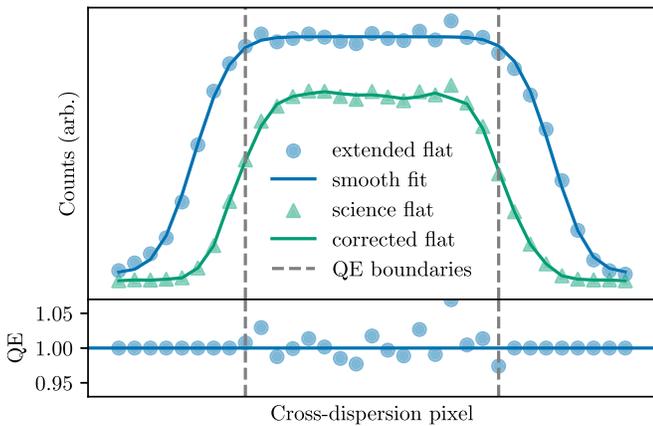


Figure 3. Example cross-dispersion, cross section of an extended flat as compared to a science flat. The best fit of Equation (1) is given for the extended flat. The resultant quantum efficiency (QE) corrections are shown in the lower plot, and the bounds of the corrected region are demarcated. Note that these are corrections *relative* to the mean QE. The QE corrected science flat is also shown in the upper plot.

deviation of the nightly stack of dark frames is determined for every bias-subtracted gain-corrected pixel. Then, the median of these standard deviations is assigned as the read noise for each amplifier region. We assume that the read noise is consistent throughout each night of observation.

For the median-combined science flat frame, there is an additional noise term: intrinsic variability in the flat-fielding LED. As described in Blackman et al. (2020), the LED brightness varies by about 0.5% over time. The uncertainty from this variability is proportional to the counts (as opposed to photon noise, which is proportional to the square root of the counts). Therefore, it is also summed in quadrature to produce the total noise for the science flats. As with all median- or mean-combined frames, the noise model for the science flat is divided by the square root of the number of frames.

3.2. Spectral Extraction

“Extraction,” in the context of the EXPRES pipeline, refers to the process of converting reduced two-dimensional CCD data into a series of one-dimensional, normalized spectra—one for each order of the echelleogram. This process involves tracing the echelle orders, removing scattered light, executing the optimal extraction, and continuum normalizing the resultant spectra. Details for these steps are found in the following sections.

3.2.1. Order Tracing

The orders of the echellogram are traced using the reduced science flat frame. First, the orders are detected using a peak-finding algorithm along the mean-combined center three columns. Then, for each order, moving from this center-line outward one column at a time, triplets of neighboring columns are mean-combined and the centroid of the resultant array is calculated. The right and left ends of each order are determined by setting a $S/N > 30$ threshold and stopping the trace once 50 subsequent columns do not reach this threshold. Finally, these centroids are smoothed along each echelle order using a 6th degree polynomial. A single set of traces is calculated for each night of observations.

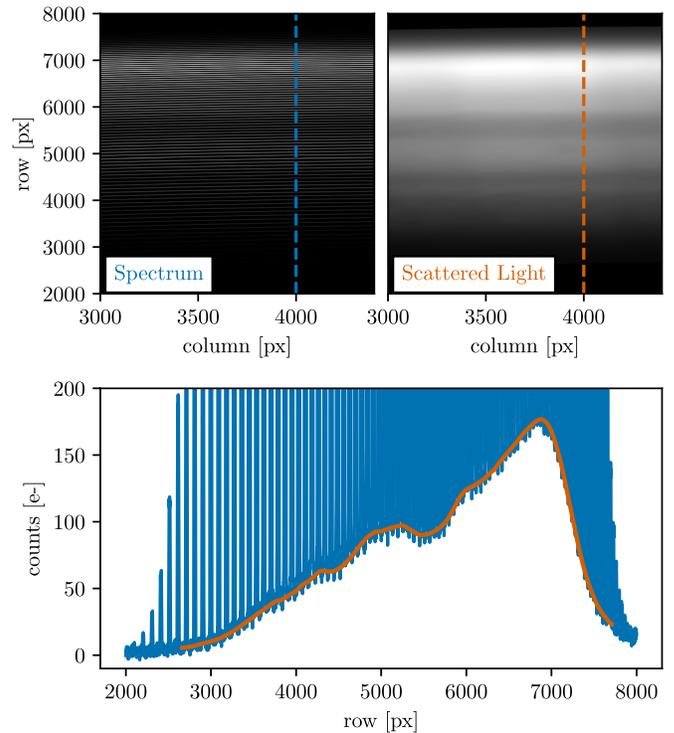


Figure 4. In the upper plot, two-dimensional scattered light approximation for reduced and median-combined science flat frames taken on 2019 October 24 are shown. The scales for the two images are not matched in order to emphasize the scattered light. The cross section of a single column (4000) for both the spectrum and calculated scattered light is shown in the lower plot. Note that the range of traced orders sets the upper and lower limits of the calculated scattered light.

3.2.2. Scattered Light

Due to imperfections in the EXPRES optics, some scattered light from the instrument hits the CCD. This diffuse, scattered light is assumed to be smoothly varying across the detector and is estimated from the counts in the regions between the orders of the echellogram. First, a variance-weighted mean and associated uncertainty is calculated for each column of each inter-order region (including those immediately above and below the traced echellogram) using the seven pixels set halfway between adjacent traced orders. These inter-order background approximations are then smoothed using a cubic b-spline with knots set every ~ 100 columns. Cosmic rays that could potentially skew this smoothed fit are iteratively rejected using a 5σ outlier cut.

The two-dimensional scattered light image is generated through a quadratic interpolation along each column of the smoothed inter-order backgrounds. The calculated scattered light is subtracted from the associated reduced image before any further extraction. This process is completed for the science flats, stellar frames, and all wavelength calibration frames. An example of a subregion of the resultant scattered light approximation is shown in Figure 4.

3.2.3. Optimal Extraction

The optimal extraction algorithm that we implement in the EXPRES pipeline is an updated version of the algorithms developed by Horne (1986), Piskunov & Valenti (2002), and Zechmeister et al. (2014). Following along the traces (from Section 3.2.1) of each echelle order, we construct a least-

squares estimator for each 33 pixel tall column x :

$$\chi_x^2 = \sum_y (D_{x,y} - P_{x,y}s_x)^2 w_{x,y}, \quad (2)$$

where $D_{x,y}$ are the photoelectron counts for each pixel (x,y) in the reduced data (see Figure 5), $P_{x,y}$ is the model of the slit function corresponding to those same pixels, s_x is the extracted spectral intensity, and $w_{x,y}$ are the weights for each pixel. These weights are inversely proportional to the variance of $D_{x,y}$ ($\sigma_{x,y}^2$, see Section 3.1.4) and include a binary cosmic ray mask, $M_{x,y}$ (i.e., $w_{x,y} = M_{x,y}/\sigma_{x,y}^2$; see Zechmeister et al. 2014). The minimization of Equation (2) has an analytic solution:

$$s_x = \frac{\sum_y w_{x,y} D_{x,y} P_{x,y}}{\sum_y w_{x,y} P_{x,y}^2}. \quad (3)$$

The propagated uncertainty of the extraction

$$\sigma_{s_x} = \sqrt{\frac{1}{\sum_y w_{x,y} P_{x,y}^2}} \quad (4)$$

is rescaled by $\chi_{\text{red},x}$ as in Zechmeister et al. (2014). However, we smooth $\chi_{\text{red},x}$ across each order using a third-order polynomial before applying it to σ_{s_x} to avoid low number statistic variance.

Cosmic rays are rejected by iteratively adding to the cosmic ray mask, $M_{x,y}$, based on a tiered outlier rejection algorithm. After calculating s_x for an entire order, the residual $(D_{x,y} - P_{x,y}s_x)w_{x,y}$ for each pixel is calculated. For each column in the order with a pixel that exceeds 8σ , the pixel with the largest residual is rejected and s_x is re-calculated for that column. This process repeats until all 8σ outliers are rejected. Next, this same rejection process repeats for 2σ outliers that also neighbor previously rejected pixels, thus rejecting the dim tails of otherwise bright cosmic rays. Importantly, this second step is executed using two-dimensional pixel information, meaning that a bright cosmic ray on a pixel in a given column can have its tail rejected in an adjacent column.

The EXPRES pipeline has two distinct methods of approximating the model $P_{x,y}$ in Equation (3):

1. Fit each column of the science flat with a parametric slit function (Equation (1)), smooth the parameters along each order using a b-spline, and then normalize the function.
2. Use the science flat, without normalization, as in Zechmeister et al. (2014).

This yields two modes of operation when extracting data, wherein (1) keeps the echelle blaze function of each order intact while (2) intrinsically removes the blaze.

We note that using method (2) requires inclusion of the variance prior in the flat $\sigma_{P_{x,y}}^2$, as well as that for the data $\sigma_{D_{x,y}}^2$, when determining the weights $w_{x,y}$ of the least-squares estimator. Therefore, the variance prior for each pixel should instead be constructed as

$$\sigma_{x,y}^2 = \sigma_{D_{x,y}}^2 + s_x^2 \sigma_{P_{x,y}}^2, \quad (5)$$

and an additional nonlinear cost term must be added to Equation (2) to prevent $s_x \rightarrow \infty$. In our case, we choose $\sum_y \ln \sigma_{x,y}^2$.

Thus, the minimization of Equation (2) must be solved numerically:

$$s_x^{(n+1)} = \frac{\sum_y w_{x,y}^{(n)} D_{x,y} P_{x,y}}{\sum_y w_{x,y}^{(n)} P_{x,y}^2 + R_x^{(n)}}, \quad (6)$$

where the weights $w_{x,y}^{(n)}$ for each iteration are calculated using the previous iteration's solution $s_x^{(n)}$ and $s_x^{(0)}$ is calculated assuming all $\sigma_{P_{x,y}}^2 = 0$. The relaxation factor $R_x^{(n)}$ is also chosen to be

$$R_x^{(n)} = \sum_y \sigma_{P_{x,y}}^2 w_{x,y}^{(n)} - \sum_y \sigma_{P_{x,y}}^2 (w_{x,y}^{(n)})^2 (D_{x,y} - s_x^{(n)} P_{x,y})^2 \quad (7)$$

from the minimization of the modified Equation (2). This numerical process is repeated until a relative tolerance of 10^{-10} is met. Thus, method (2) is operationally slower than method (1), but intrinsically accounting for un-corrected QE variations and automatically removing the blaze without relying on a blaze model outweigh this minor increase in computational cost.

Removal of the blaze through method (2) can be naively understood by recognizing Equation (2) is approximately solved by $s_x \approx \frac{D_x}{P_x}$, where P_x is the spectral intensity of the LED source times the blaze and D_x is the spectral intensity of the star times the blaze. Both P_x and D_x describe their respective intensity at the given column x of the order. Therefore, the resultant continuum of a stellar spectrum extracted using method (2) is simply the continuum of the star divided by the spectrum of the LED source. Since both of these are slowly varying functions, we can model the extracted stellar continuum with a simple linear model for each order, calculated iteratively with 2σ outlier rejection for those pixels contained in absorption lines. See Figure 6 for examples of extracted spectra and the approximated continua.

In the EXPRES pipeline, we strictly use (2) to extract all data and only use (1) as a secondary check for our extracted RVs. However, we note that maximizing the S/N of the cross-correlation function described in Section 3.4 requires weights, which increase with signal-to-noise of the stellar continuum, to be assigned to different portions of the spectrum. Unfortunately, the posterior uncertainties of the extraction have too much scatter to provide these weights. We find that using the blaze function is a good proxy for a smooth weighting function; therefore, we approximate a blaze function model for each epoch using the α -hull method (Xu et al. 2019) applied to a science flat extracted with (1). We find that RV results obtained with (2), including the separately defined blaze function, are superior to those results from spectra extracted with (1).

Extraction method (2) is also the most appropriate choice for wavelength calibration spectra, such as those from the ThAr lamp and the LFC. The calibration methods used in Section 3.3 are linear and can make use of the scattered posterior uncertainties as weights when fitting each emission line without systematically shifting results as in cross-correlation. Additionally, light from the LED, ThAr lamp, and LFC all travel through approximately the same lengths of fiber, meaning removal of an instrumental response function (i.e., a calibration ‘‘continuum’’) should not be necessary when using method (2).

The S/N of a given observation is reported here as the maximum s_x/σ_{s_x} in echelle order 111. This is effectively the

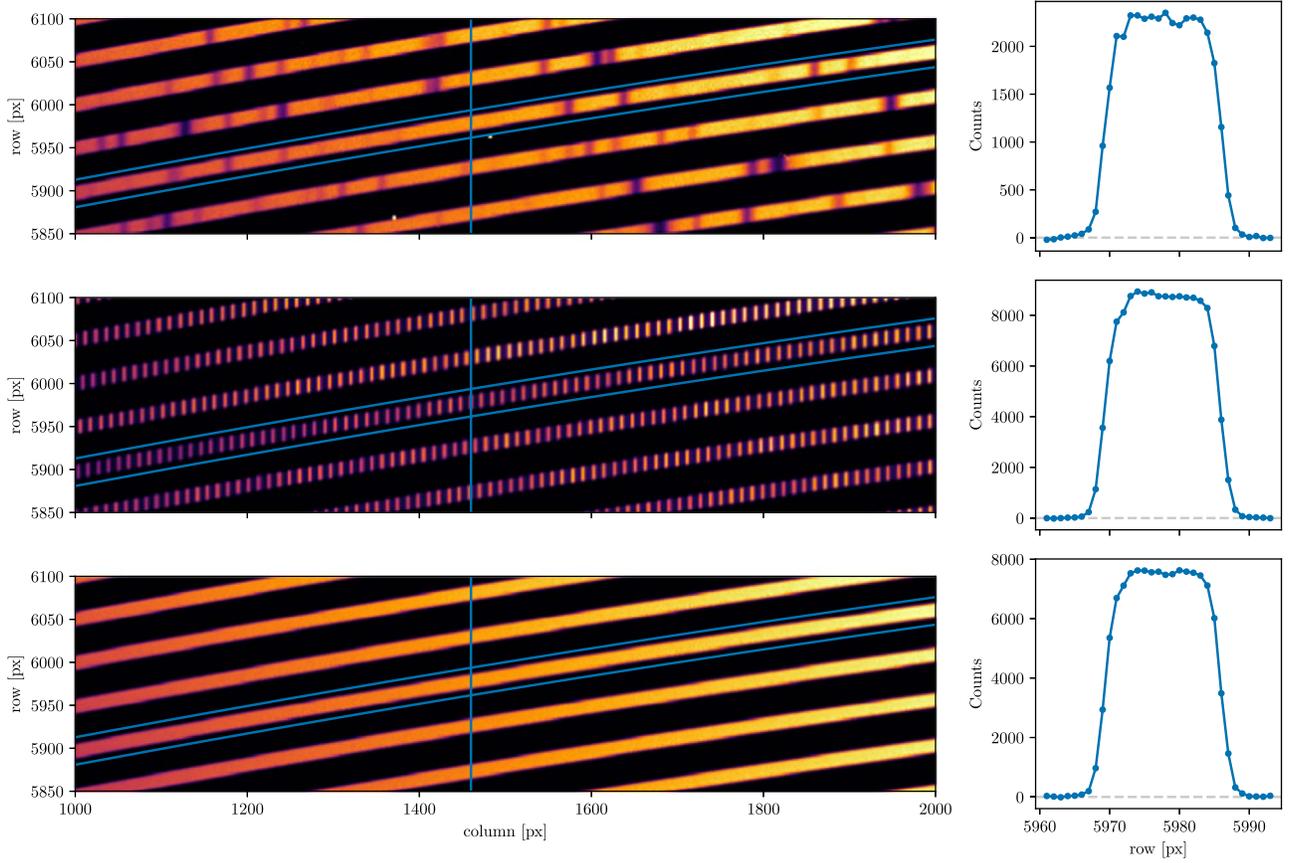


Figure 5. Section of raw images for a science exposure (above; HD 217014), laser frequency comb (middle), and calibration flat (below) taken with EXPRES on 2019 October 24. The extraction aperture for our optimal extraction of echelle order 100 is shown on the flat images as the intersection of the traced order (with vertical extent of 33 pixels) and the single-column slit at $x = 1460$, both shown in blue. The reduced counts in the extraction aperture for each image (i.e., removing hot pixels and QE variations) are shown in the right panels, as a function of pixel row position y .

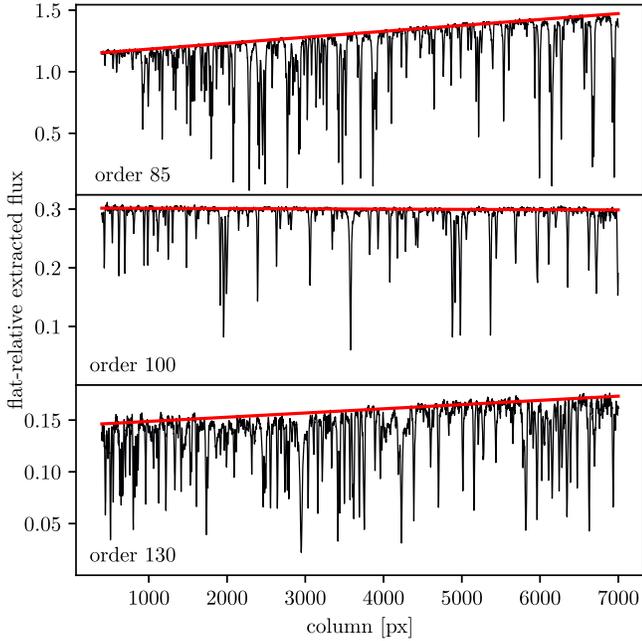


Figure 6. Flat-relative optimally extracted flux (black) and associated calculated continua (red) for three orders from HD 217014 observed on 2019 October 24. All continua were calculated using a linear model. Notice that the spectral flux relative to the flat-fielding LED source flux is significantly different between orders, but the resultant continuum across each order varies slowly enough to remain essentially linear.

per-pixel S/N at 550 nm, conforming to the metric used by Fischer et al. (2016) to compare many contemporaneous spectrographs. The resolution element of EXPRES contains approximately 4 pixels (Jurgenson et al. 2016); thus the per-resolution-element S/N is simply twice the per-pixel S/N.

3.3. Wavelength Calibration

EXPRES uses a laser frequency comb (LFC) as its primary calibration source, which generates a series of spectral lines evenly spaced in frequency, whose nominal frequencies ν_n satisfy the relation

$$\nu_n = \nu_{\text{rep}} \times n + \nu_{\text{offset}}, \quad (8)$$

for integers n . The repetition rate ν_{rep} and offset frequency ν_{offset} are referenced against a GPS-disciplined quartz oscillator, providing calibration stability corresponding to a fractional uncertainty of less than 8×10^{-12} for integration times greater than 1 s.

While it is possible to obtain an absolute calibration from the LFC once the free spectral range of the echellogram has been adequately characterized, the LFC suffers from poor throughput in very blue and very red orders. In particular, although our instrumental throughput is sufficient to permit order tracing and extraction from echelle orders 75 through 160 (for $3800 \text{ \AA} < \lambda < 8220 \text{ \AA}$), the LFC only sufficiently illuminates echelle orders 82 through 135 (for $4500 \text{ \AA} < \lambda < 7500 \text{ \AA}$). The photonic crystal fiber (PCF) of the LFC was then replaced

in 2019 July due to the decreasing stability of the LFC. With the replacement, the polarization of the LFC was switched, making the LFC redder and thus changing the orders illuminated by the LFC to echelle orders 82 through 130 (with the blue edge at 5300 Å). The polarization switch should significantly increase the lifetime and stability of the LFC. Consequently, we use ThAr lamp exposures taken at the beginning and end of each night as a secondary calibration source to provide well-constrained wavelength solutions for orders outside the range of the LFC.

Calibration triplets (3 LFC’s) are taken through the science fiber at roughly 15–30 minute intervals throughout observing, interwoven with science exposures. The exposure times of these calibration frames are chosen to match the target S/N of the science exposures. While EXPRES is equipped with a secondary square fiber to permit simultaneous wavelength calibrations, we choose to take calibrations through the science fiber so that our calibration data sample the same pixels and optical elements as the science exposures. This strategy aims to homogenize our exposures to pixel-level, uncalibratable systematic errors. Also, as shown by Blackman et al. (2020), the instrumental stability of EXPRES is such that sampling the LFC every 15–30 minutes provides enough information to correct for any instrumental changes throughout the night, as simultaneous calibration would. Calibration images are taken while the telescope is slewing and so typically cost little additional time (less than 2 minutes an hour).

A ThAr wavelength solution is generated from each ThAr exposure using the IDL code `thid.pro`, developed by Jeff Valenti, which identifies ThAr lines by matching lines in an exposure against a line atlas. A sixth-order, 2D polynomial is then fitted over pixel location x and the absolute echelle order m against the scaled wavelength $m\lambda$. Matching lines against an atlas is performed manually once at the beginning of each calibration epoch; otherwise, the wavelength solution from the immediately preceding ThAr exposure is used as an initial guess for the locations of atlas lines in a given ThAr exposure, allowing this process to be automated. Since the LFC lines are sparse relative to the precision of the ThAr calibration (1 LFC line every 10 pixels on average with the ThAr calibration accurate to the nearest pixel), this is sufficient to permit unambiguous mode identification for the LFC lines.

For any given LFC exposure, the locations of modes are identified by fitting Gaussians to each peak after a smooth background has been subtracted. An initial, trial wavelength solution is generated by linearly interpolating the ThAr solutions from the beginning and end of the night. These are used to determine the mode number n corresponding to the frequency of each mode. Once again, a 2D polynomial is fitted for $m\lambda$ as a function of m and x . Since the LFC produces a far denser set of lines (typically about 20,000 lines across 50 orders are identified in an LFC exposure, compared to about 4000 lines across 82 orders in a ThAr exposure), we use a 2D polynomial described by a 10×10 matrix of coefficients (ninth order in each dimension) for the fit. The locations of the ThAr lines are also included in the fit in order to constrain the behavior of this polynomial in echelle orders that are otherwise inaccessible to the LFC.

For each of the polynomial coefficients describing the wavelength solution, we fit a smooth function (a cubic polynomial) in time. We have found that interpolating the polynomial coefficients, rather than directly interpolating the

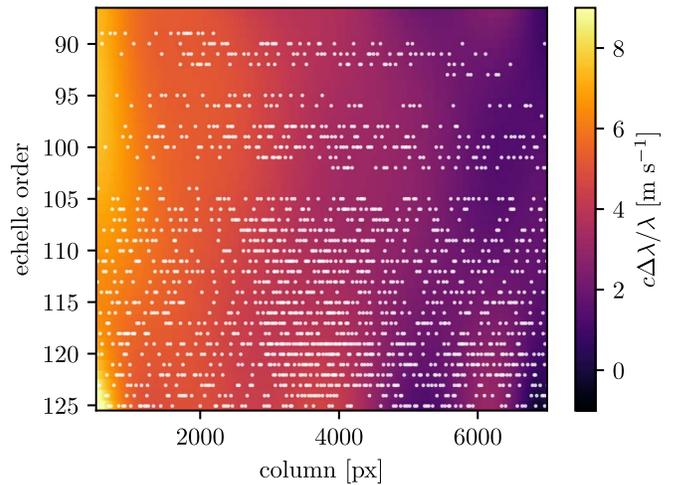


Figure 7. Calibration drift over the course of a single night (2019 October 24) for the LFC-constrained region of the echellogram, plotted in terms of the absolute echelle order number and the pixel column of the CCD. The average calibration drift for the whole night ($\sim 4.0 \text{ m s}^{-1}$) is of similar magnitude to local variations in the drift; therefore using a single average velocity offset would necessarily incur significant additional calibration error. In other words, the wavelength solution along the left and right edges of the shown spectral format would be offset by -4.0 and $+4.0 \text{ m s}^{-1}$, respectively, if a single velocity offset was used. Spectral lines used for radial-velocity solutions (shown as white dots drawn from the ESPRESSO G2 linelist) sample the detector in a nonuniform fashion and result in different overall velocity offsets, depending on spectral type.

pixel-wise wavelength solutions, is more robust to imperfections in individual calibration frames. This set of 100 functions is evaluated at the photon-weighted midpoint time of each science exposure to generate a wavelength solution. This differs from the standard practice at other spectrographs (e.g., HARPS and ESPRESSO), where a single velocity offset, rather than a time-dependent wavelength solution, is assigned to each science exposure. We choose to do this in order to accommodate time-dependent variations in the characteristics of the instrument, which may lead to calibration shifts that cannot be adequately described by a single average velocity offset.

We illustrate this in Figure 7, where we show the calibration drift over the course of a single night in the LFC-constrained region of the echellogram. Our radial-velocity solutions (see Section 3.4) sample spectral lines that are non-uniformly distributed across the detector; therefore, it is advantageous to characterize local variations to prevent overall systematic offsets. Of course, the precise differential velocity imparted by this calibration drift ultimately depends on the spectral type, as well as both barycentric and systemic velocity of the star under consideration.

As the final step in wavelength calibration, the EXPRES pipeline applies a barycentric correction to the wavelength solution of each stellar observation using the method described by Blackman et al. (2017). For each 1 s exposure of the exposure meter, a barycentric correction is calculated using BARYCORR (Wright & Eastman 2014). The photon-weighted average barycentric correction is calculated for each of eight wavelength bins of the spectra. A third-degree polynomial is then fit to these averages, yielding a smoothly varying wavelength-dependent barycentric correction $z_B(\lambda)$ for the observation. As shown by Tronsgaard et al. (2019), it is important to distinguish this from a photon-weighted midpoint

time used to calculate an overall chromatically dependent barycentric correction (e.g., Landoni et al. 2014), as this can impart a $\sim 10 \text{ cm s}^{-1}$ systematic error to the radial velocity, especially in cases of longer exposure times, high airmass, or poor seeing.

Finally, we apply z_B directly to the wavelength solution:

$$\lambda_{\text{bary}} = \lambda_{\text{lab}}^{(\text{vac})} (1 + z_B(\lambda_{\text{lab}}^{(\text{air})})), \quad (9)$$

where λ_{lab} is the LFC-generated lab-frame wavelength solution and λ_{bary} is the wavelength solution in the frame of the solar system barycenter. Because the EXPRES exposure meter is not in vacuum (as opposed to EXPRES itself), $z_B(\lambda)$ is measured using air-wavelengths. Therefore, λ_{lab} is converted from vacuum to air using the algorithm and parameters derived by Ciddor (1996) before applying the barycentric correction in Equation (9).

3.4. Radial-velocity Solutions

The data analysis pipeline of EXPRES employs two distinct computational techniques to independently extract radial velocities from stellar spectra:

1. A “cross-correlation function” method (CCF; see Baranne et al. 1979) is used to determine a rough estimate of the absolute radial velocity for each observation (Section 3.4.1).
2. A forward model based on a morphed NSO solar spectrum is used to derive a more precise relative radial-velocity curve (Section 3.4.2).

Both of these methods are currently implemented in the EXPRES pipeline for self-validation. Simultaneous results from both methods are presented in Section 4. The methods as implemented in the EXPRES pipeline are described as follows.

3.4.1. Cross-correlation

As the first step in our analysis, a CCF method estimates the absolute RV of EXPRES science targets precise to several tens of cm s^{-1} (depending on photon noise). We also use the CCF method to diagnose drifts and instabilities in our calibration sources, using line lists given by the comb parameters following Equation (8) for the LFC or a ThAr line atlas for the ThAr lamp.

The CCF is constructed from the input spectrum $f(\lambda)$ as well as a spectral-type linelist—a set of spectral lines at rest vacuum wavelengths $\{\lambda_i(0)\}$ associated with contrast weights $\{c_i\}$ and widths $\{h_i\}$. For a given trial radial velocity v , the wavelength of each line in the linelist is redshifted appropriately to

$$\lambda_i(v) = \lambda_i(0) \sqrt{\frac{c + v}{c - v}}. \quad (10)$$

The CCF is then computed as a numerical approximation to the integral

$$\text{CCF}(v) = \int d\lambda f(\lambda) \sum_i c_i w\left(\frac{\lambda - \lambda_i(v)}{h_i}\right), \quad (11)$$

where w is an arbitrary window function approximating a Dirac δ function and λ is with respect to the barycentric-corrected wavelength solution from Equation (9).

The CCF in Equation (11) is computed independently for each echelle order with a variety of trial velocities, and the CCFs for all relevant orders are co-added before a velocity model is fitted. This is a similar practice to other CCF-based RV pipelines (e.g., Brahm et al. 2017). When deriving extreme-precision radial velocities, we only include orders falling within the spectral range of the LFC, since in principle it affords considerably better sampling density and calibration stability than those regions covered by the ThAr lamp alone.

An appropriate functional model is then fitted against the co-added values of the CCF. The position parameter and posterior uncertainties of the fitted model are returned as the reported velocity and formal errors. Other quantities of astrophysical interest (e.g., rotational broadening width, bisector inverse slope) are also computed from the co-added CCF.

Our construction of the CCF incorporates the ability to use an arbitrary window function w . In the current iteration of the EXPRES pipeline, we use a cosine function, matching other contemporary CCF implementations (e.g., Freudling et al. 2013; Brahm et al. 2017; Modigliani et al. 2019). We also use a Gaussian functional model to fit the CCF for our reported radial velocities. Other possible combinations of window functions and CCF models are discussed in Section 3.4.3.

3.4.2. Forward Modeling

In addition to our CCF RV solution, we have developed a new forward-modeling technique by adapting algorithms developed for the iodine RV technique (Marcy & Butler 1992; Butler et al. 1996), as well as ideas from the “line-by-line” method developed by Dumusque (2018). Forward modeling from empirical stellar spectral templates is known to produce velocities with less statistical scatter than the CCF method, and typically measures relative rather than absolute radial velocities (Anglada-Escudé & Butler 2012). Our modeling process is simplified relative to the iodine method because the optical design of EXPRES was optimized to provide stability in the line spread function (LSF) of the spectrograph (Jurgenson et al. 2016; Blackman et al. 2020), eliminating the need to model the instrumental LSF with several free parameters. In addition, free parameters for wavelength solution and dispersion are eliminated since the barycentric wavelength solution (Equation (9)) is provided as part of the nightly optimal extraction.

First, we construct a spectral template for each stellar target. An ideal template will have very high S/N and will be a good spectral match to the program stars. Our starting point is to obtain a set of four consecutive spectra—each with S/N of about 250—providing an effective S/N of 500 per pixel or S/N of 1000 per resolution element. As described by Dumusque (2018), we prefer to use spectra with low barycentric velocities so that the program spectra shift around the approximate zero-point wavelengths. The telluric contamination is then modeled in each spectrum using SELENITE (Leet et al. 2019) and divided out. Finally, the set of spectra are co-added.

However, even the co-added spectrum will not provide a high enough S/N for a robust template. Therefore, we take the additional step of morphing the NSO solar spectrum (see Figure 8) with a native S/N $\sim 10,000$ and resolution $\sim 500,000$ to match the co-added, telluric-cleaned spectra for each of our program stars using the following procedure:

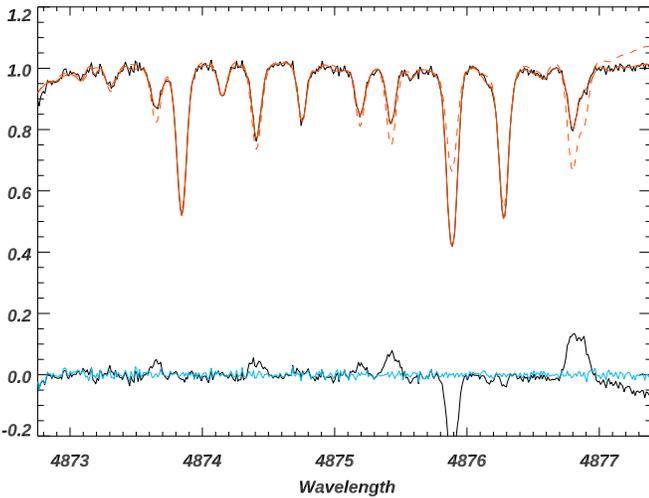


Figure 8. High-resolution, high-S/N NSO spectrum (red dashed line) is shifted and cubic-spline interpolated to the wavelength scale of a program observation (upper black line). The difference spectrum (bottom black line) is used to identify discrepancies between the spectra above the photon-noise threshold (bottom blue line). A Levenburg–Marquardt algorithm drives the growth of pseudo-lines until the NSO spectrum has morphed to match the spectrum of the program star (solid red line). This morphed spectrum is then used as a template for forward modeling.

1. The co-added program star spectrum is divided into ~ 2000 individual chunks that are 140 pixels ($\sim 2 \text{ \AA}$) wide within the orders of the spectrum covered by the LFC.
2. The barycentric wavelengths of the program star are used to extract a segment of the NSO spectrum with generous padding of 200 pixels for modeling shifts. This segment of the NSO spectrum is shifted to the barycentric frame of the co-added program star spectrum.
3. A Levenburg–Marquardt (L-M) least-squares algorithm is used to (i) determine the best-fit width for a Gaussian convolution kernel to rotationally broaden the NSO spectrum, (ii) refine the Doppler shift of the NSO spectrum, and (iii) apply a vertical shift to align the continuum of the NSO and the co-added spectrum.
4. The rotationally broadened and shifted NSO spectrum is cubic-spline interpolated onto the wavelength scale of the co-added spectrum.
5. A difference spectrum is calculated. Nodes are dropped down consecutively at points where the absolute value of the difference spectrum exceeds a threshold, characterized by the photon noise of the co-added spectrum. Pixels with the largest residuals in the difference spectrum are modeled first. The maximum number of nodes is 60, but depending on the chunk, there are typically about a dozen nodes required to model the NSO spectrum for each 130 pixel chunk.
6. At each node, a positive or negative Gaussian feature with a width characterized by the line spread function of EXPRES is used to perturb the NSO spectrum; the depth of the morphing feature is determined by L-M fitting of the residuals.
7. Iterative growth of the morphing lines stops when the residuals of the difference spectrum are consistent with photon noise.
8. Each chunk of the template is weighted according to the amount of spectral information, using the S/N and the

derivative of intensity I with pixel:

$$\sum \frac{\delta I}{c} \frac{\lambda}{\delta \lambda} \frac{1}{S/N}. \quad (12)$$

Once the template for each star has been generated, a L-M fit is executed for every 140 pixel chunk of the program star spectra. Telluric-affected pixels in each observation are assigned zero weight in the fit. There are only two free parameters for each chunk: a Doppler shift and continuum normalization scale factor. Thus, each chunk—45 LFC-calibrated orders each with about 47 chunks yielding ~ 2000 total chunks—provides an independent measurement of the relative RV for the star. The RV for each chunk is subsequently subtracted by the mean of that chunk over all observations, thus removing any offsets that might occur because of geometric anomalies in the detector while preserving the spread in RV variations.

Weights for each chunk are determined using empirical arguments, the χ^2 of the L-M fit, and a chunk-specific modifier based on its relative temporal scatter. Chunks that do not contain any absorption lines in the stellar template as well as chunks that yield relative velocities greater than $\pm 1000 \text{ m s}^{-1}$ are assigned zero weight. Moreover, chunks that have $\chi^2 > 5.0$ (typically occurring if an incorrect stellar template was used or a telluric line was missed, for example) and remaining chunks that are among those with the largest 3% of reduced χ^2 are all assigned zero weight.

Because some chunks have less spectral information, there will be more scatter in the RVs derived from these chunks. For example, chunks in the blue part of the spectrum typically have several absorption lines, but chunks in the red part of the spectrum may have only one spectral line, meaning the L-M fitting will not be well-constrained. Likewise, telluric contamination within a given chunk can manifest as large scatter in the RV over time. Therefore, the non-zero weight for given chunk i within an observation j is assigned as

$$w_{ij}^{-1} = \chi_{ij}^2 \frac{\sum_j^n (v_{ij} - \bar{v}_j)^2}{n - 1}, \quad (13)$$

where \bar{v}_j is the median velocity for all chunks of observation j and n is the total number of observations for a given stellar target. The reported RV measurement for each observation is thus a weighted mean of the individual chunk velocities, and the formal error is the corresponding standard error of the weighted mean.

3.4.3. Further Analysis

Once RVs have been derived, the extracted spectra and CCFs are passed down the pipeline for more sophisticated analysis. The spectral range of EXPRES is intended to permit characterization of stellar activity and planetary atmospheric absorption lines. For chromospheric activity in particular, we extract the Ca II line core emission ratio index S_{HK} (using the parametric model of Isaacson & Fischer 2010), calibrated to yield results consistent with the Mount Wilson Observatory catalog (Duncan et al. 1991).

We also aim to incorporate spectroscopic activity indicators directly into the RV solution methodology. For example, in addition to using a rectangular “box” function and truncated cosine in the CCF, which are implemented in other similar velocity analysis codes (e.g., Freudling et al. 2013; Brahm et al. 2017;

Modigliani et al. 2019), the EXPRES analysis code implements CCF computation using Gauss–Hermite window functions of the form

$$w(x) = \frac{1}{\sqrt{2^n n!} \sqrt{\pi}} H_n(x) \exp\left[-\frac{x^2}{2}\right], \quad (14)$$

where H_n is the n th (physicists’) Hermite polynomial. Computing higher-order CCFs as coefficients in a Hermite-functional decomposition, and more generally with respect to different orthogonal basis functions, will permit more sophisticated analysis of stellar activity (via a sparse description of variations in the CCF line profile) as an alternative parameterization to current derived observables (such as the CCF bisector inverse slope/FWHM). Alternatively, data-driven decorrelation of stellar activity from bulk radial velocities, or alternative template-based RV solution methodologies (P. Holzer et al. 2020, in preparation), may be possible once we have built up an archive of stellar spectra.

4. Initial Results—HD 217014

We now present velocities derived with this pipeline based on 47 observations of HD 217014 (51 Peg) over the 10-month span between the beginning of Epoch 4 and December 1, 2019 (Table 3). We do so to examine various sources of uncertainty and error in the velocimetric pipeline, while avoiding the known instrumental instabilities inherent in Epochs 1–3 (see Blackman et al. 2020; A. Szymkowiak et al. 2020, in preparation, for details), and to compare the two radial-velocity methods outlined in Section 3.4. Observations with an S/N less than 160 are not included in this analysis. We construct our CCF using the G2 line mask from the ESPRESSO pipeline (Freudling et al. 2013; Modigliani et al. 2019) and a cosine window function. We also fit all RVs with a single planet Keplerian model, constrained by the literature value of the orbital period (4.2308 days, Wang & Ford 2011). This Keplerian model is parameterized by the velocity semi-amplitude (K), the eccentricity (e), the argument of periastron, and a phase of periastron.

Figure 9 shows the resulting radial velocities from both the CCF and Forward Model methods, along with their respective orbital fits. Because the CCF uses a linelist with absolute wavelengths, the fit systemic velocity of $-33.2603(5) \text{ km s}^{-1}$ (in excellent agreement with Gaia Collaboration 2018) has been removed. For each of the RV epochs, we also fit an independent velocity offset relative to the overall offset. For epochs 4 and 5, these are $-1.5(4) \text{ m s}^{-1}$ and $1.2(5) \text{ m s}^{-1}$ when using the CCF method, and $-1.2(8) \text{ m s}^{-1}$ and $0.8(7) \text{ m s}^{-1}$ when using the forward model. These offsets account for modifications to the instrumental systematics, owing to the various fiber changes and realignments. The offsets differ slightly in magnitude between the two RV analysis methods because they intrinsically weigh regions of the detector differently, accentuating or mitigating certain instrumental systematics.

We show the values of the Keplerian fit parameters in Table 4, along with the rms residual and the median formal error (σ_v). The parameter uncertainties were derived by taking the square root of the product of the posterior variances and reduced χ^2 of the least-squares fit. By way of comparison, we also perform the same procedure with 8 years of archival velocities from the HIRES instrument on the Keck I telescope

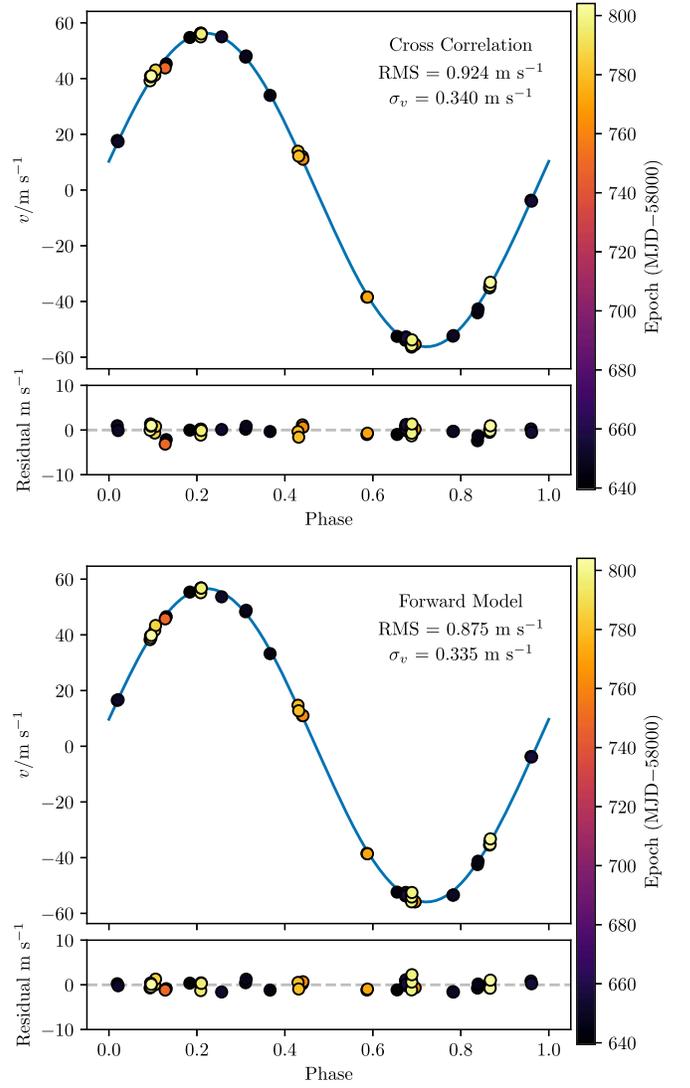


Figure 9. Phased radial velocities, Keplerian orbital fits, and residuals for EXPRES observations of 51 Peg b. The figure is labeled by the rms residual to the fitted orbital solution, as well as the median formal error σ_v of all data points shown.

Table 3
EXPRES Commissioning RVs of 51 Peg

BMJD	$V_{ccf}/\text{cm s}^{-1}$	$V_{fm}/\text{cm s}^{-1}$	RV Epoch	S/N
58639.45844	-3320697 ± 20	5739 ± 32	4	385
58641.45174	-3331424 ± 44	-5035 ± 42	4	179
58643.46218	-3321644 ± 34	4854 ± 34	4	225
58644.46095	-3322776 ± 35	3527 ± 34	4	233
58646.45596	-3330577 ± 39	-4045 ± 38	4	203
	⋮			

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

(corrected for instrumental systematics per Tal-Or et al. 2019) and 4 months of data from the HARPS DRS (Trifonov et al. 2020). The EXPRES orbital solution parameters for 51 Peg are consistent with those returned from these previous studies, but with higher precision due to the improved formal errors. The two EXPRES RV methods are also internally consistent, with the forward model producing a slightly more favorable rms.

Table 4
Fit Parameters for 51 Peg b

Instr.	$K/\text{m s}^{-1}$	e	RMS/ m s^{-1}	$\sigma_v/\text{m s}^{-1}$
EXPRES CCF	56.24 ± 0.14	0.000 ± 0.002	0.924	0.340
EXPRES FM	56.26 ± 0.13	0.007 ± 0.003	0.875	0.335
HARPS DRS	53.4 ± 1.6	0.062 ± 0.010	0.941	1.023
HIRES	56.7 ± 0.4	0.020 ± 0.007	2.74	1.169

Finally, we note that the rms of the EXPRES fit residuals approximately matches that of the HARPS DRS, with the EXPRES fit returning parameters more comparable with literature values (Wang & Ford 2011; Bedell et al. 2019; Wilson et al. 2019).

5. Discussion

Each step of the pipeline contains several free parameters—for instance, the degree of the fitting polynomials to use for the spatial wavelength solution fits and temporal smoothing, as well as various S/N thresholds for calibration line identification and order inclusion in the CCF. To assist community users of the instrument, we have opted, as far as possible, to preselect reasonable default values for most of these parameters, which may be overridden at runtime. In what follows, we document some nonobvious but critical aspects of these systematics, and we illustrate some aspects of the decision-making process for choosing our default values for some of these parameters.

5.1. Formal versus True Velocity Errors

Since the formal velocity errors returned from the CCF fitting procedure are constructed only from the co-added CCFs and their propagated errors, they do not account for effects like wavelength calibration error (inducing spurious velocity shifts) or time estimation error (via erroneous barycentric corrections). Instead, they mostly reflect velocity estimation error due to photon noise being propagated to the CCF.

On the other hand, the formal velocity errors returned from the forward model fitting procedure does include some information about relative uncertainty in certain regions of the EXPRES detector. For instance, a chunk that tends to have a telluric line will naturally incur more spread in the measured RV for that chunk. Therefore, even though this chunk is down-weighted by our analysis, its spurious effect still propagates to the RV error.

These assumptions are borne out in Figure 10, showing these formal errors as a function of the observation S/Ns. The CCF points depend essentially only on photon noise and potentially CCD readout and optimal extraction systematics—which we are confident of having adequately accounted for—up to some constant that may depend on, for example, the choice of CCF line mask or window function, or intrinsic astrophysical properties of the target. Conversely, the Forward Model formal errors contain much more scatter, which we believe folds in some uncertainty from telluric contamination and, potentially, stellar noise. Thus, our estimation of the true photon-noise limit of EXPRES is better described by the propagated errors of the CCF analysis, though the two analyses yield quite similar results. As shown in Figure 10, we define this limit to be 30 cm s^{-1} for a single observation at the EXPRES target S/N of 250.

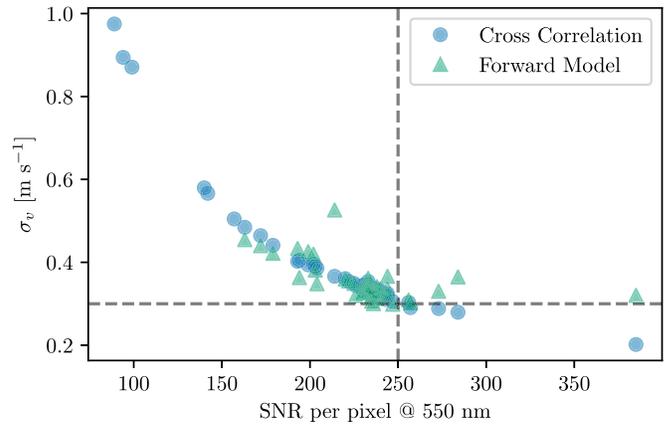


Figure 10. Formal RV errors returned from both RV fitting procedures, plotted against the per-pixel S/N (as a characterization of photon noise) for velocities from Section 4. EXPRES’s target S/N of 250 is shown with a vertical line and the approximate associated formal error of 30 cm s^{-1} is shown with a horizontal line. Note that data here with S/N less than 160 are not included in Figure 9.

Following our diagnosis and repair of the LFC beat frequency noted in Table 1, we measured the remaining uncalibratable velocity errors arising from wavelength calibration in particular to be relatively small—between 4 and 6 cm s^{-1} rms (A. Szymkowiak et al. 2020, in preparation). There also exist several other sources of error (e.g., from uncalibratable instrumental systematics and guiding errors) that constitute additional contributions to the RV error budget. To correctly estimate the velocimetric error, one needs to appropriately account for and then combine these error terms (e.g., by adding them in quadrature) with the formal value reported from the RV analysis. A detailed inventory of these error sources (Blackman et al. 2020) estimates the combined instrumental and guiding errors of EXPRES at $\sim 10 \text{ cm s}^{-1}$. Thus the single observation error of EXPRES is clearly dominated by the apparent photon noise.

5.2. Chromatic Dependences

Many of the novel techniques that we have adopted in the EXPRES pipeline involve the introduction of chromatic dependences into quantities that have previously been considered to be uniform with wavelength, such as calibration offsets and the barycentric offset velocity. It therefore behooves us to investigate possible chromatic effects that emerge at the end of the CCF velocity-solving and orbit-fitting procedure.

The CCF analysis in Section 4 was performed by co-adding CCFs derived from echelle orders 126 through 86 ($4850 \text{ \AA} < \lambda < 7150 \text{ \AA}$) before fitting an absorption-line model to derive a velocity. These orders are those for which at least $N_{\text{min}} = 19$ LFC lines are detected that pass both the S/N threshold and all quality checks imposed by our peakfinder. (N_{min} depends on the degree of the polynomial fitted to the wavelength solution, which is a free parameter in our code, as are the threshold values for these quality checks).

In Figure 11, we show the rms residuals from the orbital solutions that arise when we repeat these analyses while varying the range of echelle orders used when co-adding CCFs; our default parameter selections are indicated with the red dotted lines. In particular, this means extrapolating the wavelength solution and chromatic barycentric correction beyond the spectral range of the LFC,

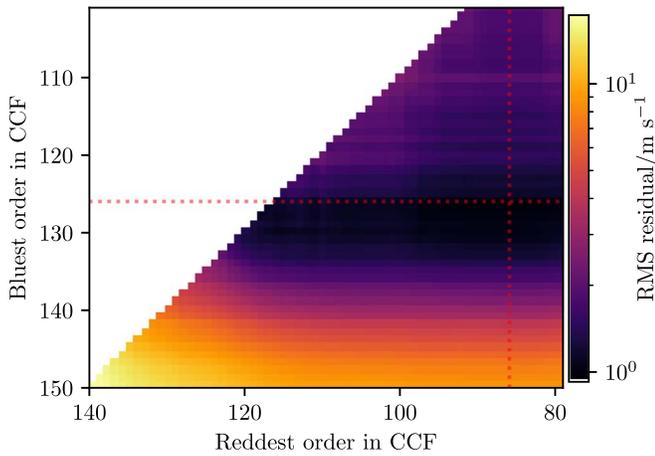


Figure 11. rms scatter from orbital solutions fitted to HD 217014 as a function of bluest and reddest echelle order included in the CCF computation. Points along each diagonal indicate sets of velocities computed with the same number of echelle orders. At least 10 echelle orders have been included in all CCF computations. Dotted red lines show the default parameters used in the preceding sections.

which covers echelle orders 130 through 82, and of the chromatic exposure meter, which covers orders 135 through 86 ($\sim 4650 \text{ \AA} < \lambda < 7150 \text{ \AA}$).

For this data set, we see that there is a sharp dependence on the bluest order co-added into the CCF. Moreover, we see that the introduction of orders redder than the LFC cutoff also slightly increases the rms error to the fit. Recalling that the wavelength solution outside of the LFC region is largely constrained by the ThAr lamp, this potentially implies a calibration offset between the LFC and the ThAr lamp, despite both sources illuminating the instrument through the same fiber.

On the other hand, we do not see any similarly sharp cutoff when extrapolating the chromatic barycentric correction to outside of the wavelength range covered by the exposure meter. This suggests that the wavelength dependences of our barycentric corrections (detailed more fully in Blackman et al. 2019) are generally smooth enough for robust extrapolation.

Finally, it is possible to choose narrower ranges of echelle orders that yield smaller rms errors than our default parameter selection. However, we note that this is potentially dependent on the specifics of the CCF line lists used in the computation and also possibly on underlying astrophysical properties of the science targets and other fortuitous factors. We have opted to include as many orders as can be accurate, so as to minimize photon noise, while still leaving the option to include fewer orders available to end users.

6. Conclusion

The commissioning process on the EXPRES instrument is essentially complete, along with the development of an optimal extraction pipeline that we have been using for preliminary RV analysis through both CCF and forward modeling techniques. Within the instrumental back-end (i.e., limiting ourselves to the calibration unit and the spectrograph proper), we have determined our photon-noise-limited RV errors to be approximately 0.3 m s^{-1} for a single observation with S/N of 250. With EXPRES’s current observing strategy of four observations per night per target, this result implies a nightly

measurement error of only 0.15 m s^{-1} . While our on-star measurement error appears to be $\sim 0.9 \text{ m s}^{-1}$ —based on residual rms to an orbital fit of 51 Peg b—we must also note that our RV analysis pipeline does not fully address photospheric velocity sources, telluric contamination, or longer-term instrumental errors. These other sources of RV scatter are beyond the scope of this paper, although we are actively investigating them. The RV precision presented in this paper, therefore, represents our worst-case scenario in the absence of further improvements. We anticipate upcoming hardware improvements and more sophisticated RV solution methodologies to only enhance our measurement precision and long-term instrumental stability.

Moreover, there are multiple parts of the pipeline that remain under active development. We are investigating the use of spectro-perfectionism (Bolton & Schlegel 2010; Cornachione et al. 2019) as an alternative to optimal extraction. We are also exploring a hierarchical, non-parametric wavelength solution that takes advantage of the low degrees of freedom allowed in a stabilized instrument and the density of lines offered by new wavelength calibrators (L. L. Zhao et al. 2020, in preparation). There are also plans to implement a more data-driven approach (as in Bedell et al. 2019, with modifications to permit chromatic barycentric corrections) as yet a third RV analysis technique.

These caveats notwithstanding, we have demonstrated that the technical innovations that have been invested into the development of novel instrumentation and software analysis techniques for EXPRES have largely paid off—they have permitted us to unambiguously attain $\text{sub-} \text{m s}^{-1}$ on-sky radial-velocity precision. Presently, this makes EXPRES the most precise EPRV spectrograph in the Northern Hemisphere.

Finally, we hope the lessons we have learned in the process of commissioning this instrument, and the techniques we have developed, to be of some value to the community of EPRV instrument builders moving forward.

We thank the anonymous referee for comments that significantly improved this paper. We gratefully acknowledge Yale University, the Heising-Simons Foundation, and an anonymous Yale alumni donor for providing telescope time for EXPRES used to obtain the data in this paper. We also acknowledge the extensive work that went into the RePack spectral extraction code, written by Lars Buchhave; this code helped benchmark the EXPRES optimal extraction described in this paper. J.O. thanks X. Dumusque for illuminating and productive discussions. We especially thank the NSF for funding that allowed for precise wavelength calibration and software pipelines through NSF ATI-1509436 and NSF AST-1616086 and for the construction of EXPRES through MRI-1429365. This material is based on work supported by the National Science Foundation Graduate Research Fellowship under grant No. DGE1122492 (RRP, LLZ, ABD).

Facility: LDT.

Software: SciPy library (Virtanen et al. 2020), NumPy (Oliphant 2006; van der Walt et al. 2011), Astropy (Astropy Collaboration et al. 2013; Price-Whelan et al. 2018).

ORCID iDs

Ryan R. Petersburg  <https://orcid.org/0000-0003-2168-0191>
 J. M. Joel Ong (王加冕)  <https://orcid.org/0000-0001-7664-648X>

Lily L. Zhao  <https://orcid.org/0000-0002-3852-3590>

Ryan T. Blackman  <https://orcid.org/0000-0002-0303-3276>
 John M. Brewer  <https://orcid.org/0000-0002-9873-1471>
 Lars A. Buchhave  <https://orcid.org/0000-0003-1605-5666>
 Samuel H. C. Cabot  <https://orcid.org/0000-0001-9749-6150>
 Allen B. Davis  <https://orcid.org/0000-0002-5070-8395>
 Colby A. Jurgenson  <https://orcid.org/0000-0002-5018-7761>
 Christopher Leet  <https://orcid.org/0000-0001-5847-9147>
 David Sawyer  <https://orcid.org/0000-0001-8483-4667>
 Mikhail Sharov  <https://orcid.org/0000-0003-3047-5368>
 René Tronsgaard  <https://orcid.org/0000-0003-1001-0707>
 Andrew E. Szymkowiak  <https://orcid.org/0000-0002-4974-687X>
 Debra A. Fischer  <https://orcid.org/0000-0003-2221-0861>

References

- Anglada-Escudé, G., & Butler, R. P. 2012, *ApJS*, **200**, 15
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Baranne, A., Mayor, M., & Poncet, J. L. 1979, *VA*, **23**, 279
- Bedell, M., Hogg, D. W., Foreman-Mackey, D., Montet, B. T., & Luger, R. 2019, *AJ*, **158**, 164
- Blackman, R. T., Fischer, D. A., Jurgenson, C. A., et al. 2020, *AJ*, in press
- Blackman, R. T., Ong, J. M. J., & Fischer, D. A. 2019, *AJ*, **158**, 40
- Blackman, R. T., Szymkowiak, A. E., Fischer, D. A., & Jurgenson, C. A. 2017, *ApJ*, **837**, 18
- Bolton, A. S., & Schlegel, D. J. 2010, *PASP*, **122**, 248
- Brahm, R., Jordán, A., & Espinoza, N. 2017, *PASP*, **129**, 034002
- Butler, R. P., Marcy, G. W., Williams, E., et al. 1996, *PASP*, **108**, 500
- Ciddor, P. E. 1996, *ApOpt*, **35**, 1566
- Cornachione, M. A., Bolton, A. S., Eastman, J. D., et al. 2019, *PASP*, **131**, 124503
- Dumusque, X. 2018, *A&A*, **620**, A47
- Duncan, D. K., Vaughan, A. H., Wilson, O. C., et al. 1991, *ApJS*, **76**, 383
- Fischer, D. A., Anglada-Escudé, G., Arriagada, P., et al. 2016, *PASP*, **128**, 066001
- Freudling, W., Romaniello, M., Bramich, D. M., et al. 2013, *A&A*, **559**, A96
- Gaia Collaboration 2018, *yCat*, **1345**, 0
- Home, K. 1986, *PASP*, **98**, 609
- Isaacson, H., & Fischer, D. 2010, *ApJ*, **725**, 875
- Jurgenson, C., Fischer, D., McCracken, T., et al. 2016, *Proc. SPIE*, **9908**, 99086T
- Landoni, M., Riva, M., Pepe, F., et al. 2014, *Proc. SPIE*, **9147**, 91478K
- Leet, C., Fischer, D. A., & Valenti, J. A. 2019, *AJ*, **157**, 187
- Marcy, G. W., & Butler, R. P. 1992, *PASP*, **104**, 270
- Modigliani, A., Sownowska, D., & Lovis, C. 2019, ESPRESSO Pipeline User Manual, <https://www.eso.org/sci/software/pipelines/espresso/espresso-pipe-recipes.html>
- Oliphant, T. 2006, NumPy: A guide to NumPy (Spanish Fork, UT: Trelgol Publishing), <http://www.numpy.org/>
- Pepe, F., Cristiani, S., Rebolo, R., et al. 2013, *Msngr*, **153**, 6
- Piskunov, N. E., & Valenti, J. A. 2002, *A&A*, **385**, 1095
- Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, *AJ*, **156**, 123
- Probst, R. A., Lo Curto, G., Ávila, G., et al. 2016, *Proc. SPIE*, **9908**, 99086A
- Schwab, C., Rakich, A., Gong, Q., et al. 2016, *Proc. SPIE*, **9908**, 99087H
- Steinmetz, T., Wilken, T., Araujo-Hauck, C., et al. 2008, *Sci*, **321**, 1335
- Tal-Or, L., Trifonov, T., Zucker, S., Mazeh, T., & Zechmeister, M. 2019, *MNRAS*, **484**, L8
- Trifonov, T., Tal-Or, L., Zechmeister, M., et al. 2020, *A&A*, in press
- Tronsgaard, R., Buchhave, L. A., Wright, J. T., Eastman, J. D., & Blackman, R. T. 2019, *MNRAS*, **489**, 2395
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, **17**, 261
- Wang, J., & Ford, E. B. 2011, *MNRAS*, **418**, 1822
- Wilson, M. L., Eastman, J. D., Cornachione, M. A., et al. 2019, *PASP*, **131**, 115001
- Winn, J. N., & Fabrycky, D. C. 2015, *ARA&A*, **53**, 409
- Wright, J. T., & Eastman, J. D. 2014, *PASP*, **126**, 838
- Xu, X., Cisewski-Kehe, J., Davis, A. B., Fischer, D. A., & Brewer, J. M. 2019, *AJ*, **157**, 243
- Zechmeister, M., Anglada-Escudé, G., & Reiners, A. 2014, *A&A*, **561**, A59