The Astronomical Journal, 161:252 (12pp), 2021 June https://doi.org/10.3847/1538-3881/abef68
© 2021. The American Astronomical Society. All rights reserved.

Broadband Stability of the Habitable Zone Planet Finder Fabry–Pérot Etalon Calibration System: Evidence for Chromatic Variation

Ryan C. Terrien[1], Joe P. Ninan[2,3], Scott A. Diddams[4,5], Suvrath Mahadevan[2,3], Samuel Halverson[6], Chad Bender[7], Connor Fredrick[4,5], Fred Hearty[2,3], Jeff Jennings[4,5], Andrew J. Metcalf[8], Andrew Monson[2], Arpita Roy[9,10], Christian Schwab[11], and Guđmundur Stefánsson[12,13]

[1] Carleton College, One North College Street, Northfield, MN 55057, USA; rterrien@carleton.edu
[2] Department of Astronomy & Astrophysics, 525 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA
[3] Center for Exoplanets and Habitable Worlds, 525 Davey Laboratory, The Pennsylvania State University, University Park, PA 16802, USA
[4] National Institute of Standards and Technology, 325 Broadway, Boulder, CO 80305, USA
[5] Department of Physics, University of Colorado, 2000 Colorado Avenue, Boulder, CO 80309, USA
[6] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA
[7] Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721, USA
[8] Space Vehicles Directorate, Air Force Research Laboratory, 3550 Aberdeen Avenue SE, Kirtland AFB, NM 87117, USA
[9] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
[10] Department of Physics and Astronomy, Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA
[11] Department of Physics and Astronomy, Macquarie University, Balaclava Road, North Ryde, NSW 2109, Australia
[12] Princeton University, Department of Astrophysical Sciences, 4 Ivy Lane, Princeton, NJ 08540, USA
Received 2020 October 6; revised 2021 March 8; accepted 2021 March 15; published 2021 May 5

Abstract

The comblike spectrum of a white light–illuminated Fabry–Pérot etalon can serve as a cost-effective and stable reference for precise Doppler measurements. Understanding the stability of these devices across their broad (hundreds of nanometers) spectral bandwidths is essential to realizing their full potential as Doppler calibrators. However, published descriptions remain limited to small bandwidths or short time spans. We present an ∼6 month broadband stability monitoring campaign of the Fabry–Pérot etalon system deployed with the near-infrared Habitable Zone Planet Finder (HPF) spectrograph. We monitor the wavelengths of each of ∼3500 resonant modes measured in HPF spectra of this Fabry–Pérot etalon (free spectral range = 30 GHz, bandwidth = 820–1280 nm), leveraging the accuracy and precision of an electro-optic frequency comb reference. These results reveal chromatic structure in the Fabry–Pérot mode locations and their evolution with time. We measure an average drift on the order of 2 cm s^{-1} day^{-1}, with local departures up to ±5 cm s^{-1} day^{-1}. We discuss these behaviors in the context of the Fabry–Pérot etalon mirror dispersion and other optical properties of the system and the implications for the use of similar systems for precise Doppler measurements. Our results show that this system supports the wavelength calibration of HPF at the $\lesssim$10 cm s^{-1} level over a night and the $\lesssim$30 cm s^{-1} level over ∼10 days. Our results also highlight the need for long-term and spectrally resolved study of similar systems that will be deployed to support Doppler measurement precision approaching ∼10 cm s^{-1}.

Unified Astronomy Thesaurus concepts: Near infrared astronomy (1093); Radial velocity (1332); Spectroscopy (1558); Fabry-Perot interferometers (524); Exoplanet detection methods (489)

1. Introduction

High-precision Doppler radial velocity (RV) measurements have progressed significantly in the past decade. Leveraging a suite of new technologies and analysis techniques (Fischer et al. 2016), current state-of-the-art RV measurements have reached the ∼1 m s^{-1} precision regime on bright stars and are pushing toward the ∼10 cm s^{-1} precision (i.e., one part in 3×10^9) required to detect an Earth-mass planet orbiting in the habitable zone (Kopparapu et al. 2013) of a Sun-like star (Damasso et al. 2020; Petersburg et al. 2020). New instruments and surveys striving for this level of precision include EXPRES (Blackman et al. 2020), ESPRESSO (Pepe et al. 2010), and NEID (Schwab et al. 2016). These and others are motivated by an array of timely scientific objectives and opportunities, including the drive for precise exoplanetary mass constraints to test hypotheses about exoplanet formation and evolution (e.g., Lee 2019), as well as to inform plans for future space-based studies of exoplanetary atmospheres (Kempton et al. 2018; Fortenbach & Dressing 2020).

Approaching the ∼10 cm s^{-1} goal is a significant and multifaceted challenge (Gaudi et al. 2019) that requires exquisite control of the spectrograph optomechanical platform (Pepe et al. 2010; Stefansson et al. 2016; Robertson et al. 2019) and instrumental illumination (Roy et al. 2014; Halverson et al. 2015; Kanodia et al. 2018; Sirk et al. 2018; Blackman et al. 2020), as well as robust analysis techniques (Pepe et al. 2010; Petersburg et al. 2020). Even with state-of-the-art RV systems, instrumental drifts[14] at the level of tens of meters per second or larger are routinely observed (e.g., Bauer et al. 2020). Careful wavelength calibration provides a path toward reducing the systematic noise floor imposed by these spectrograph drifts by tracking a well-understood calibration source simultaneously with stellar observations and using these to correct the instrumental drifts and isolate the stellar signal.

Historically, wavelength references such as molecular iodine cells and thorium–argon (Th–Ar) lamps have provided stable and well-measured reference spectra for RV measurements and

[13] Henry Norris Russell Fellow.

[14] The term "drift" is used here to describe the uncontrolled variations in a spectrograph or calibration system, which can manifest as changes in the measured spectra or derived quantities.

enabled RV measurements with a precision of $\sim$1 m s^{-1} (e.g., Butler et al. 2017; Udry et al. 2019). However, the limited spectral coverage and irregular spectral line separations and amplitudes, as well as operational issues related to lifetime, chemical purity (e.g., Nave et al. 2018), and commercial availability, have driven interest in new calibration sources. Among such systems are those based on laser frequency combs (LFCs; Murphy et al. 2007) and white light–illuminated Fabry–Pérot etalons (FPs; Wildi et al. 2010), both of which yield stable optical spectra of regularly spaced, approximately equal intensity, narrow spectral lines (resonant modes) that can be used for calibration. There are, however, substantial differences in the nature and utility of these two types of calibration sources.

The mode spectrum of an LFC is tied to precisely known atomic radio frequency standards; the result is a well-defined optical frequency for each spectral line, with each line frequency known to a fractional absolute[15] precision limited only by the frequency standard. For the best laboratory standards, a precision of 1×10^{-18} has been achieved (Diddams et al. 2020). However, for practical implementations with a GPS reference at an astronomical observatory, the typical precision is 1×10^{-11} or better (Metcalf et al. 2019a). This level of absolute precision corresponds to an RV error of $<$1 cm s^{-1}, making these systems promising for detection of Earth-mass planets over extended temporal periods.

Although recent years have seen the deployment of a number of commercial and custom astronomical LFC systems (Probst et al. 2014; Metcalf et al. 2019a; Donati et al. 2020; Petersburg et al. 2020; Hirano et al. 2020), these systems remain costly and operationally complicated. This is due largely to the confluence of two stringent requirements for use as Doppler calibrators: the need for broad (many hundreds of nanometers) spectral coverage and the need to resolve individual LFC modes. A principal challenge is that maintaining flat illumination over such broad bandwidths—particularly toward blue ($<$500 nm) wavelengths—is difficult when coupled with the high repetition rate ($\gtrsim$10 GHz) required to generate resolvable comb modes at a typical Doppler spectrograph resolution ($R = \lambda/\Delta\lambda$ typically between 50,000 and 100,000). For a detailed discussion of these and other challenges, see Probst (2015) and McCracken et al. (2017).

In comparison, FPs are simpler and less costly, so they offer an attractive alternative to LFCs. The core mechanism of an FP is an optical resonator formed by two partially reflective mirrors that is characterized by a spectrum of resonant longitudinal modes that are approximately equally spaced in frequency (Perot & Fabry 1899). This resonator can act as a stable filter of injected continuum, transmitting a regular, comblike spectrum. Unlike in an LFC system, the center frequency of a given line in the spectrum of an FP is not known a priori, as the frequency of each line depends on the precise optical details of the FP, including its size and geometry, the characteristics of the mirror reflective coating, and how light is coupled in and out. Moreover, because the mode spectrum of an FP is tied to physical properties and not atomic frequency standards, the frequencies of its modes may be influenced by environmental conditions (e.g., temperature and pressure; Jennings et al. 2017), as well as slow material drift (Berthold et al. 1977).

Nonetheless, astronomical FP systems may be actively or passively stabilized; the stability and high information content of their spectra have been demonstrated to support $\lesssim$1 m s^{-1} RV precision in several systems (Halverson et al. 2014; Cersullo et al. 2017, 2019; Stürmer et al. 2017; Das et al. 2018). Further, an extensive history of the use of FP systems in the field of optical frequency metrology indicates that suitably designed FP systems can support narrow-bandwidth measurements ($<$1 nm) at a level of stability and precision far exceeding 10 cm s^{-1} over many years (e.g., Dubé et al. 2009). However, astronomical FPs are distinguished by their comparatively broad bandwidths ($\gtrsim$500 nm) and wide resonances;[16] their use is therefore subject to different systematics. The long-term and spectrally resolved broadband performance of astronomical FP systems at the $\lesssim$10 cm s^{-1} level remains mostly undetermined.

The FP analyzed in this work is a broadband (820–1280 nm), passively stabilized astronomical FP. Our earlier laboratory study of this FP (Jennings et al. 2020, hereafter J20) revealed significant deviations in the drift behavior of three FP resonant modes (at 780, 1064, and 1319 nm) compared to expectations if these drifts are caused by changes in the FP mirror separation (described in more detail in Section 2.4). Confirmation of this behavior in situ at the observatory, as well as more spectrally extensive measurement, is essential if this and other FP systems are to be used to calibrate the long-term and potentially complicated drifts of the RV spectrographs targeting Earth-mass exoplanets. As an example, one possible implementation of an astronomical FP calibrator relies on the independent tracking of a small number of FP modes in order to track the drift of the entire FP mode spectrum (e.g., using Rb reference gas cells; Reiners et al. 2014; Schwab et al. 2018). This approach may be limited by the extent to which the behavior of the full population of FP modes can be parameterized by the behavior of this small set of modes.

To help inform the use of this and other astronomical FP calibrators, we present here measurements extending over 6 months of the frequency stability of a broadband near-infrared (NIR) FP system for the Habitable Zone Planet Finder (HPF) spectrograph (Mahadevan et al. 2014). These measurements leverage a dedicated LFC calibration system (Metcalf et al. 2019a) to probe the behavior of each of $\sim$3500 FP modes in the NIR. In Section 2, we provide descriptions of the spectrograph, LFC, and FP systems we use and discuss the basic equations describing each. In Section 3, we describe our data set and methodology, highlighting the use of the LFC for wavelength calibration and the precision achieved in tracking the FP resonant modes. In Section 4, we describe the results of our measurements, showing the wavelength dependence of the FP mode spectrum, the precision possible with the "bulk" stability of the FP, and the significant chromatic structure in the drift rates of the FP resonant modes. In Section 5, we discuss the potential causes and implications of our results for Doppler spectrograph calibration. We summarize our key findings in Section 6.

2. Background

2.1. The HPF

The HPF is a stabilized, fiber-fed, cross-dispersed echelle spectrograph at the Hobby–Eberly Telescope (HET) at

[15] That is, relative to the SI second.

[16] Equivalently, low values of the FP finesse.

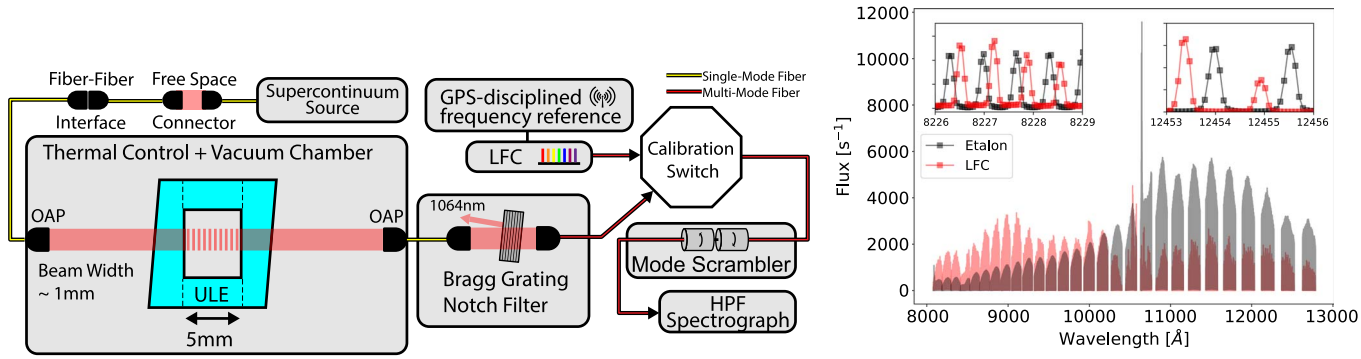

(a) (b)

Figure 1. (a) Schematic of the HPF FP calibration system. An NKT supercontinuum light source is passed through a short free-spacing connection and a fiber–fiber interface before illuminating the resonator, which is enclosed in a temperature-controlled vacuum chamber. Light is coupled into and out of the fiber via off-axis parabolic collimators (OAPs). The resonator spacer and mirrors are made from ULE glass. The mirrors are wedged to suppress undesired parasitic etalons and have an antireflection coating to suppress undesired reflection. The output light is filtered by a narrowband Bragg grating filter, which suppresses the strong pump laser line from the NKT SuperK source to prevent saturation of the HPF detector. The FP light is one of several options that may be selected by the calibration switch to feed the calibration fiber of HPF. Alternately, the mode spectrum of the HPF LFC, traceable to the SI second via GPS, can be selected. The calibration light is passed through a mode scrambler prior to injection into the HPF spectrograph. (b) Example raw spectra of the HPF FP and LFC. The spectral variations result from a combination of the intrinsic source profile and the HPF throughput, and the shape within each order is dominated by the spectrograph blaze function. As shown with the rescaled spectra in the insets, both the LFC and FP modes are narrow and well separated; there are $\gtrsim$100 modes in each of the 28 HPF echelle orders.

McDonald Observatory in Texas, optimized for Doppler RV exoplanet detection around low-mass (M dwarf) stars. It observes in the NIR (820–1280 nm) over 28 echelle orders at high resolution ($R \approx 55{,}000$). The HPF has demonstrated a thermal stability of $\sim$1 mK (Stefansson et al. 2016) over $\sim$1 month timescales, and, with a custom-built LFC calibration system (described in Section 2.2), it has been shown to be capable of a Doppler precision of $\sim$1.5 m s^{-1} on an M dwarf star over many months (Metcalf et al. 2019a).

The HPF is fed by three fibers: one corresponding to the science target, one corresponding to blank sky (to correct for emission lines in Earth's atmosphere), and one for the calibration source. For optimal drift correction during stellar observations, the calibration source may be observed simultaneously with the stellar target. Using the LFC to illuminate both the science and the calibration fibers has shown that these fibers track each other to $\lesssim$10 cm s^{-1} (Metcalf et al. 2019a).

Since its installation in 2017, the HPF has been in regular use as a facility instrument at the HET, where its observations have enabled studies of nearby M dwarf planet companions (Metcalf et al. 2019a; Cañas et al. 2020; Kanodia et al. 2020; Stefansson et al. 2020a, 2020b, 2020c), planetary atmospheres (Ninan et al. 2020), stellar activity (Robertson et al. 2020), and precise Doppler spectrograph systematics (Ninan et al. 2019).

2.2. The HPF LFC System

A custom-designed electro-optic LFC calibration system is deployed with HPF. This system is described in detail in Metcalf et al. (2019a, 2019b), but we briefly outline its properties here. A seed laser at 1064 nm is intensity- and phase-modulated at 30 GHz and broadened in a silicon nitride waveguide to generate a spectral comb with 30 GHz mode spacing spanning the HPF bandpass. The output spectrum is flattened with a combination of passive and active filters. This system provides reliable calibration light for HPF on a nightly basis and can be used simultaneously with stellar observations. The routine usage of the HPF LFC system is enabled by its exceptional reliability; the LFC was operational >97% of the time from 2017 May to 2019 November. The accuracy and precision provided by LFC spectra are our primary reference for measuring the daily and long-term drift of HPF and thus for determining the HPF wavelength solution at any given time (discussed in Section 3.2) that underlies the FP results presented here.

2.3. The HPF FP System

The HPF FP is described in detail in Jennings et al. (2020). Briefly, the FP resonator is constructed from two planar mirrors separated by an $L \approx 5$ mm ultralow-expansion (ULE) glass spacer (built by Light Machinery). This etalon is enclosed in a temperature-controlled vacuum chamber, and light is coupled in and out with single-mode optical fibers collimated by off-axis parabolic mirrors (housing and stabilization system built by Stable Laser Systems). The system is illuminated with an NKT SuperK supercontinuum source. The transmitted comb-like spectrum is once again collimated and then filtered by a narrowband Bragg grating notch filter (built by Optigrate, with a thickness of $\approx$2 mm, FWHM $\approx$0.5 nm, and 10^{-3} attenuation over $\sim$1 nm spectral width) in order to suppress the bright NKT SuperK pump laser wavelength near 1064 nm. The filtered light is coupled into a multimode fiber and injected into the HPF calibration switch. From there, it passes through a mode scrambler (built by Giga Concepts) and finally into the HPF spectrograph. A schematic sketch of the system is shown in Figure 1.

The fundamental parameters of the optical spectrum of this FP system are its free spectral range and finesse. The free spectral range is approximately 30 GHz, corresponding to approximately six resolution elements of HPF, and the free spectral range is explored in much more detail below. The targeted (coating-limited) finesse across the HPF bandpass is $\approx$40; measured (in the lab at NIST) finesse values are 33 (780 nm), 43 (1064 nm), and 41 (1319 nm; Jennings et al. 2020). These correspond to an intrinsic line width of approximately 800 MHz, a factor of $\sim$7 narrower than the instrumental profile (IP) of HPF. Due to the uncertainties involved in deconvolving these effects, an independent

measurement of the finesse (or the stability thereof) is beyond the scope of this work.

2.4. Theoretical FP Description

We discuss here the salient theoretical aspects of a broadband planar FP with mirrors separated by a distance L. Collimated light is coupled in through the partially transmissive mirrors, and successive reflections lead to interference and resonance for certain frequencies. Notably, the dielectric mirrors on the inside of the resonator in the HPF FP are designed for reflectivity of ≈92% across the entire 820–1280 nm band (corresponding to a finesse of ≈40). These mirrors achieve high reflectivity by interference off the different layers in the dielectric stack; maintaining such a reflectivity over the full 820–1280 nm band requires many layers, so the full mirror coating has a nonnegligible thickness. Solutions to Maxwell's equations for this situation are well studied; modifications to the resonance condition (as in Devoe et al. 1988) yield a frequency-dependent expression for the FP longitudinal mode spacing (free spectral range, $\Delta f_{\rm FSR}$),

$$\Delta f_{\rm FSR}(f) = \frac{c}{2nL + \frac{c}{\pi}\frac{\partial\phi(f)}{\partial f}}, \tag{1}$$

where f is the frequency of the light, n is the refractive index of the propagation medium, $\phi(f)$ is the frequency-dependent phase delay from the mirror reflection, and c is the speed of light.

An often-used simplification of Equation (1) is notable: if phase delays from the mirrors are neglected and the mirrors are in a medium with a refractive index n (presumably a slow function of frequency), $\Delta f_{\rm FSR}$ is simply

$$\Delta f_{\rm FSR} = \frac{c}{2nL}. \tag{2}$$

For mode index m, the corresponding frequency would then be

$$f_m = m\ \Delta f_{\rm FSR}. \tag{3}$$

The most fundamental parameter describing the FP mode spectrum is the mirror separation L; Equations (2) and (3) suggest that changes in L should drive the mode frequency changes, and the ratio of such changes between two modes should be equivalent to the ratio of the mode indices (see discussion in J20).

3. Methods

We measure and track the centroids of each of ∼3500 modes of the FP spectrum in HPF spectra obtained over ∼1000 epochs spanning 6 months. An accurate wavelength calibration (described in Section 3.2) allows us to translate these centroids to wavelengths, correcting for the spectrograph drift. The resulting data provide a spectral and temporal baseline to assess the wavelength-dependent stability of the FP spectrum.

3.1. Data Set

The HPF FP is part of a suite of routinely used calibration sources that also includes the LFC and hollow cathode lamps. Daily morning and evening calibration sequences each include between one and four observations of the FP, with a uniform integration time of ≈160 s, resulting in a total number of recorded counts that is consistent to $\lesssim$10% across observations. These FP observations are typically taken immediately following observations of the LFC, resulting in a minimal spectrograph drift for cross-calibration between the LFC and FP.

The one-dimensional spectra used here are optimally extracted (Horne 1986) from two-dimensional images from the H2RG detector array (Blank et al. 2012) using a custom HPF data reduction pipeline, which provides flux and variance measurements for each pixel and is described in detail in the supplemental material of Metcalf et al. (2019a) and in Ninan et al. (2018). Each FP spectrum has an average signal-to-noise ratio (S/N) per extracted pixel at mode peak of ≈490 and an approximate total photon-limited RV precision (Bouchy et al. 2001) of 4.9 cm s^{-1}.

We select for this study a period spanning 2018 November through 2019 May, which avoids spectrograph and LFC maintenance events and spans a period of uniform FP drift behavior. Later epochs do show some deviations from the behavior presented here that are under study. Including all of the FP spectra taken in the 2018 November–2019 May time frame yields 1050 spectra.

3.2. Wavelength Calibration of FP Spectra

The HPF LFC provides an absolute calibration of the entire HPF spectrum (i.e., an accurate and precise mapping of pixel to wavelength or frequency in an extracted spectrum; the wavelength solution) at any given time. However, the LFC is not continuously observed and cannot be observed simultaneously with the FP, so the determination of the wavelength solution for an arbitrary epoch hinges on a reliable understanding of the drift of the HPF wavelength solution. Careful daily measurements of the LFC and FP over ∼2 yr have enabled the construction and refinement of an accurate multiparameter model for this drift, the dominant terms of which correspond to a uniform pixel offset, a magnification term, and offsets affecting each readout channel of the H2RG detector. Here we briefly outline the relevant aspects of the wavelength calibration for FP spectra; a complete description of the HPF wavelength calibration algorithm and performance will be published in a later paper.

The HPF wavelength solution and drift measurement begins with a high-S/N LFC template, which is created by combining multiple LFC spectra that are close in time. These LFC frames are combined using an estimate of the underlying drift that is iteratively refined. The frequencies of the LFC modes in this template are known from the comb equation,

$$f_m = mf_r + f_0, \tag{4}$$

where m is the LFC mode index, f_r is the LFC repetition frequency, and f_0 is the LFC offset frequency (Diddams et al. 2020). In the HPF LFC system, both f_r and f_0 are radio frequencies that are tracked with respect to a GPS-disciplined reference oscillator. The mode index m is easily determined by coarse comparison of the template to the spectrum of a hollow cathode lamp. Thus, each mode frequency in the template spectrum is absolutely determined. The corresponding pixel position is determined by a least-squares fit of a Gaussian function to each LFC mode in this template, which provides ∼100 anchor points for the wavelength solution across each of

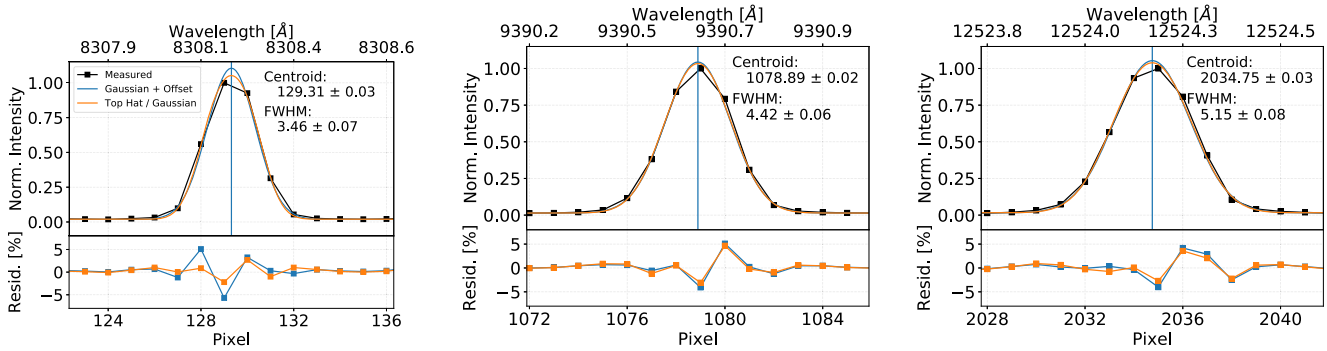

(a) (b) (c)

Figure 2. Example fits for three modes across the HPF bandpass. We tested the performance of a simple Gaussian + offset model, as well as the convolution of a top hat and Gaussian, and found only marginal reductions in the residuals with the more complicated model. We elect to use the simple Gaussian + offset, which has typical residuals of $\lesssim$5%.

the 28 echelle orders of HPF. The wavelength solution for this template is then taken as the least-squares fit of a fifth-order polynomial, fit independently to these anchor points in each spectral order. The residuals to these fits are largely normally distributed with a scatter of $\sim$10 m s^{-1}, although for approximately a third of the orders, there is some correlated noise at a level of 10–20 m s^{-1} at scales of tens of modes.

This template wavelength solution can then be appropriately transformed by a drift model to account for the instrumental drift to any epoch of interest. This model has been iteratively refined over time to represent the drift of HPF with the fewest possible parameters to maximize the S/N of the drift measurement. Each new LFC observation is compared to the LFC template in the context of this model to determine the appropriate value for each drift model parameter. The drift model addresses only change in the "dispersion" sense, i.e., those that are accessible in the one-dimensional extracted spectra.

Our HPF drift model comprises a second-order polynomial function of pixel index within each order, along with discrete offsets for the four readout channels of the H2RG detector, which were found to drift apart slowly. Over short (daily) timescales, the dominant effect is seen in the zeroth-order coefficient of the polynomial. Changes in this coefficient are equivalent to a linear offset of the spectrum projected onto the H2RG detector array. This coefficient shows a daily sawtooth pattern with a peak-to-peak amplitude corresponding to approximately 15 m s^{-1} (as shown in Figure A1 of Stefansson et al. 2020a). This pattern is thought to be driven by the change in mechanical loading from the daily fill and subsequent evaporation of the liquid nitrogen coolant used by HPF, which is stored in a tank suspended underneath the HPF optical bench. The highly repeatable shape of this sawtooth allows us to accurately interpolate the appropriate value for the zeroth-order coefficient at any epoch.

The first- and second-order coefficients in our drift model polynomial vary smoothly and slowly, on a timescale of weeks. The offsets for each readout channel also vary slowly, on an even longer timescale. This behavior allows the accurate interpolation of the relevant parameters in our drift model for any given epoch. Combined with the dominant term above, this establishes a precise determination of the wavelength solution for any single spectrum.

3.3. Mode Measurement

Within each FP spectrum, we use the corresponding wavelength solution to measure the central wavelength of each individual FP resonant mode. We measure the centroid of each mode independently using a least-squares fit of a Gaussian with a constant offset, a robust technique that provides easy-to-interpret outputs. A representative set of these fits is shown in Figure 2. The fit is performed on a 16 pixel window around each mode peak, which encompasses the full width of the $\sim$5 pixel line width and surrounding baseline. We note that these line widths are dominated by the IP, as the intrinsic FP line width is about seven times smaller. The residuals are typically $<$5%.

As indicated by the structure in the residuals of Figure 2, the FP line profiles are not perfectly well represented by a Gaussian. The fibers illuminating the HPF are masked by a slit; the line profile is therefore expected to be slightly more flat-topped than a Gaussian. Further, the line profile contains some asymmetries resulting from optical aberrations impacting the HPF response function. Similar fits to the LFC spectra (whose modes are intrinsically even narrower than the FP modes) reveal equivalent structure in the residuals, indicating that this behavior is not unique to the FP.

To probe the sensitivity of our results to the choice of fitting function, we repeated our analysis with a more complicated model corresponding to a convolution of a Gaussian and a top-hat function. While the residuals were reduced, the improvement was generally marginal and did not justify the increased complexity of the model. We therefore favor the simpler offset Gaussian model and are careful to measure mode characteristics and behavior across groups of modes in order to smooth over variations that may relate to the aberration-related asymmetries (which correlate across echelle orders) and limited sampling of individual modes.

We further mitigate the potential impact of the unique systematics for any individual mode fit by focusing in this work on the patterns that are coherent across tens to hundreds of modes and across HPF spectral orders. All of our results below involve averaging over tens to hundreds of modes.

3.4. Centroid Precision

An accurate estimate of the measurement precision on a single FP mode is necessary for determining the significance of the spectral or temporal variations observed. To help

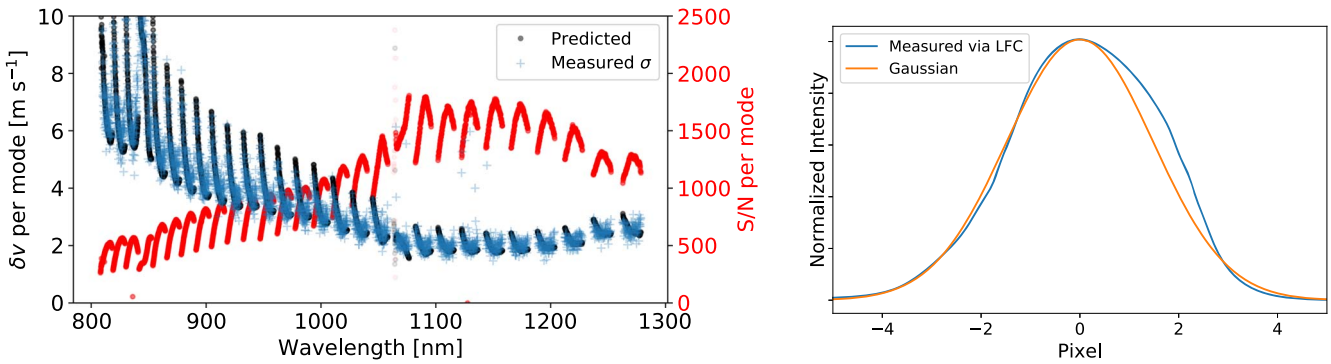

(a) (b)

Figure 3. (a) The S/N and corresponding predicted mode centroid precision for all modes in a single 160 s observation of the HPF FP spectrum, as derived using Equation (5). The S/N values are photon noise–dominated, and spectral variations in the overall FP flux distribution are dominated by both the output spectrum of the SuperK source and the instrumental efficiency of HPF. Also shown are measured standard deviations (σ) for each mode centroid spanning 100 epochs from a typical period from 2019 April 24 through 2019 May 14. Points within 1 Å of the 1064 nm pump laser are deemphasized because they are subject to instability resulting from the intensity variations of the pump laser. (b) Two example profiles used to measure the scaling factor A in Equation (5): a simple Gaussian (the functional form used for the fits described in Section 3.3) and one measured directly using the HPF LFC, corresponding to the red edge of echelle order index 5 (approximately 8765 Å). Both profiles yield an estimate of $A \approx 0.4$.

understand the precision obtained with our data, we consider the fundamental photon-limited precision (δv) on a single mode, which can be estimated as in Brault (1987),

$$\delta v \approx A \frac{\mathrm{FWHM}}{\mathrm{S/N} \times \sqrt{n}}, \tag{5}$$

where A is a (line shape–dependent) constant of order unity (Murphy et al. 2007), FWHM is the FWHM of the mode, S/N is the peak S/N[17] in the mode, and n is the number of pixels sampling the mode.

We follow Murphy et al. (2007) to confirm the value for A for FP modes in HPF spectra; we repeatedly simulate and recover the centroid of a mode with appropriate sampling and noise characterized by Poisson statistics and a small amount (∼five electrons) of detector read noise. Performing this numerical experiment with a Gaussian line at a variety of widths and S/N values, we confirm the estimate in Murphy et al. (2007) that $A \approx 0.4$. We repeated this experiment with a high-resolution line profile measured by scanning the HPF LFC (similar to, e.g., Murphy et al. 2012) as in Ninan et al. (2019). The tested profile corresponds to the edge of the HPF detector, where optical aberrations are maximized, and is shown in Figure 3. The resulting estimate for A was still approximately 0.4, confirming that this value is appropriate across the HPF spectrum of the FP.

We estimate the S/N for a given mode using the flux and variance outputs of our HPF data reduction pipeline. A typical distribution of S/N values for a single 160 s exposure and the resulting photon-limited uncertainties (related through Equation (5)) are shown in Figure 3, along with the measured standard deviation (σ) for each mode centroid across 100 epochs during a typical period confirming that our measurements are mostly photon noise–limited for a single epoch.

4. Results

4.1. Static Description

The starting point for understanding the temporal behavior of the HPF FP mode spectrum is a precise description of its properties at a single point in time. Besides defining a reference point for measuring the temporal behavior of the FP mode positions, examining the measured Δf_{FSR} may provide clues as to whether the HPF FP mode spectrum is measurably affected by phase-related deviations, as suggested by Equation (1).

We establish an accurate "single-epoch" description of the spectrum by averaging over the results of our mode centroid measurements from the representative period from 2019 April 24 through 2019 May 14 (the same period used in Figure 3). The resulting Δf_{FSR} measurements are shown in Figure 4. The measured Δf_{FSR} values cluster near the specified value of 30 GHz, but variations of approximately 50 MHz are clearly detected over large (tens to hundreds of modes) ranges.

We note that short-range (mode-to-mode) structure is also detected and can be seen as the scatter in the "raw" points plotted in Figure 4, which is elevated above the approximately megahertz-level uncertainty expected for a pair of mode centroids (see Section 3.4). The origin of this short-range structure is difficult to trace because any individual mode fit is subject to unique systematics due to detector effects (e.g., persistence and quantum efficiency variations) and the varying mismatch between the "true" line profile and the uniform Gaussian used (see Section 3.3). We focus instead in this work on the behaviors that are coherent across tens to hundreds of modes and HPF spectral orders, which is more interpretable with our current understanding of HPF systematics.

The spectral shape of the FSR variations shown in Figure 4 can be connected to the phase delay curve of the FP mirror using Equation (1). Using this equation and the vendor-provided phase curve for the FP mirror (shown in the inset of Figure 4), we perform a least-squares fit to the measured Δf_{FSR} values and find $L = 4.99231 \pm 0.00003$ mm.[18] The good correspondence between the Δf_{FSR} values predicted by this L (which sets the bulk offset of the Δf_{FSR}) and the phase curve (which imposes the spectral shape) indicates that most of the observed variation results from the mirror phase.

[17] For FP modes measured with HPF, the ratio between the integrated and peak S/N is approximately 0.5.

[18] This fit assumes no uncertainty in the phase curve and should not be understood as an accurate determination of L. Subtleties resulting from mirror dispersion in the use of FP for precise length measurements are discussed in detail in Lichten (1985, 1986).

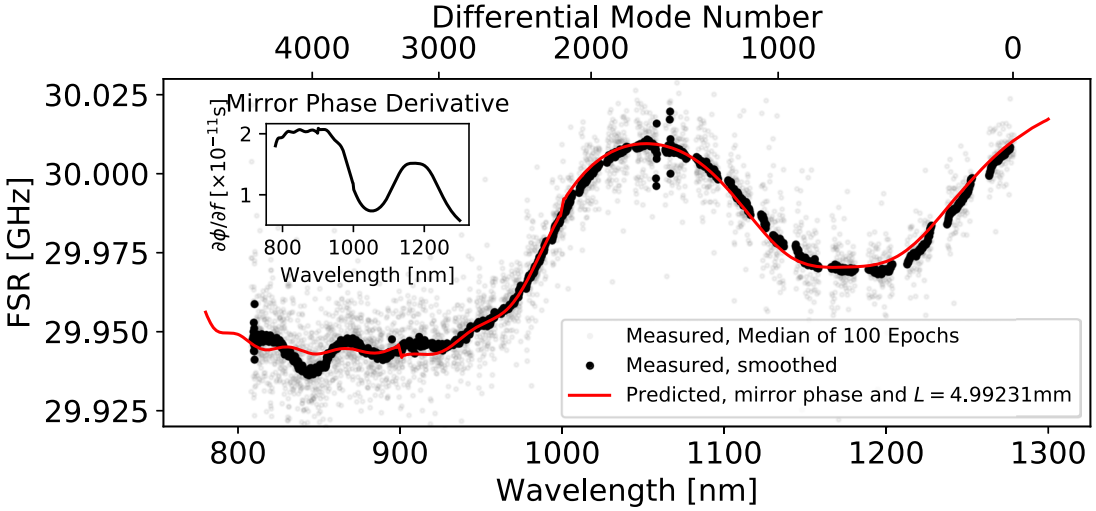

Figure 4. The FP free spectral range ($\Delta f_{\rm FSR}$) derived from the representative 100 epochs spanning 2019 April 24–2019 May 14. The FP modes are identified by their wavelengths and "differential mode number" relative to the lowest mode index in our spectrum. Short-range variations are smoothed with a 31-point rolling (spectral) mean. The "bulk" value of approximately 30 GHz primarily results from the cavity spacing, while spectral variations can be understood as resulting from wavelength-dependent mirror phase delays. Projecting the (vendor-provided) phase delays for our mirrors through Equation (1) suggests that the $\Delta f_{\rm FSR}$ variations are largely related to the mirror properties.

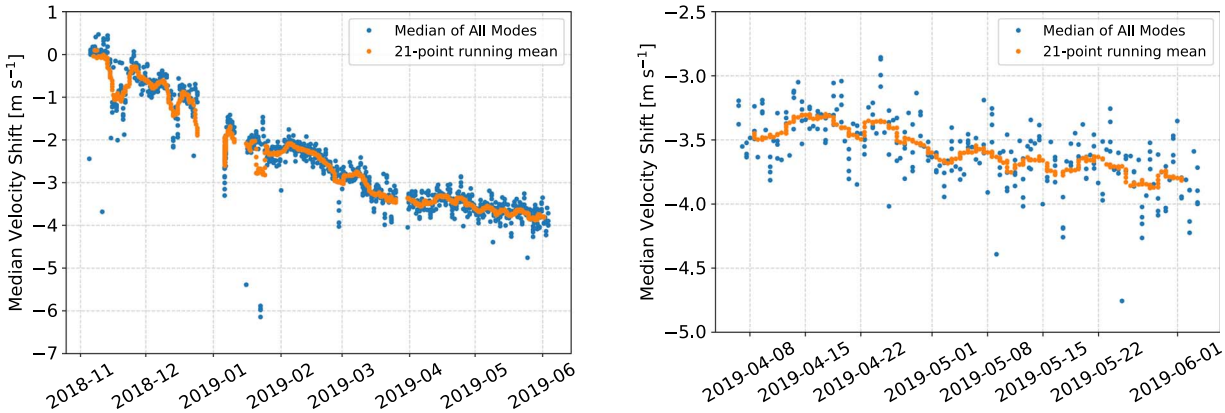

(a) (b)

Figure 5. Absolute drift of the FP, represented as a bulk velocity shift. Panel (a) shows the full time span presented in this work, including points where the stability and illumination of the FP may have caused short-term deviations. Panel (b) highlights an ∼2 month period without significant deviations.

4.2. Temporal Monitoring

We traced the drifts of each of the mode centroids over a period of ∼6 months. For comparison with other systems and as a simple examination of the FP stability, we measure the median drift behavior of all of the measured resonant modes, shown in Figure 5 as a velocity shift. By taking this median, we collapse the wavelength-dependent behavior, which is explored in more detail below. The long-term behavior of the median shift is smooth, although a number of short-term deviations are apparent. Many of these are related to periods of instability in the environment or illumination of the FP; we defer the detailed study of these short-term variations for a future work.

However, considering only the median velocity shift belies the wide spectral variation in long-term drift behavior we have measured. A few examples showing the range of these drifts are presented in Figure 6, represented again as velocity shifts. The spectral variation in both sign and magnitude of the measured velocity shifts is surprising and suggests an origin beyond a simple long-term change in the FP mirror separation L, as discussed in Section 2.4.

In order to probe the nature of these wavelength-dependent (chromatic) drifts, we note that the long-term drift is generally linear and fit for the slope (m s^{-1} day^{-1}) at each wavelength, shown in Figure 7. A clear wavelength-dependent pattern is apparent, with the drift rates varying in both sign and magnitude across the spectrum. Figure 7 also shows the measured drift rates of this FP from J20 and makes clear that current measurements of the drift rates for this FP are substantially smaller; this difference is discussed in Section 5.2.

5. Discussion

5.1. Bulk Shift

The median drift of the FP during the study period amounts to around 4 m s^{-1} over 7 months (≈2 cm s^{-1} day^{-1}), which compares favorably with similar systems (Stürmer et al. 2017; Bauer et al. 2020). To estimate the performance of this system as a calibrator for pure velocity shifts, we explore the scatter of the median drift of the FP at different timescales using the 2

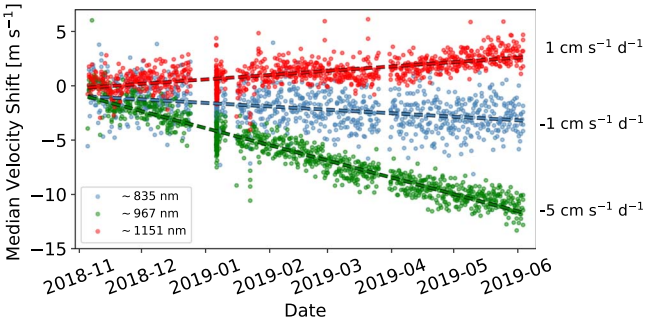

Figure 6. Long-term drift behavior of three sets of 30 FP modes, grouped around 835, 967, and 1151 nm, showing examples of spectrally resolved drifts, in contrast with the median drift shown in Figure 5. Each group is near the center of an HPF spectral order. Measured centroid shifts for individual etalon modes, as measured via the Gaussian fitting technique shown in Figure 2, are shown as Doppler shifts. The range of behavior spans Doppler shifts of varying signs and magnitudes. The long-term behavior can be approximated as linear in each case. The variation in scatter for each group results from the S/N with which these centroids were measured (see Figure 3).

month period of stable operation indicated in Figure 5. We bin the median velocity shifts at varying times and calculate the resulting standard deviations (sigma clipping at 5σ). The results, shown in Figure 8, indicate that this scatter within a night is $\lesssim$10 cm s^{-1} and $\lesssim$30 cm s^{-1} over several days. This can be understood as the expected quality of the calibration provided by the FP in the case where the spectrograph drift is dominated by a velocity-like term. These results are promising for the use of FP systems as calibrators in this context, but we caution that there are still unexplained short- and long-term variations in the FP drift behavior that are under study.

5.2. Difference from J20 Measurement

The bottom panel of Figure 7 shows that the present drift rates differ significantly from those measured in J20. The latter were based on repeated temporal scans of three FP modes (at 780, 1064, and 1319 nm) with tunable continuous-wave lasers. The FP transmission profiles were recorded, and the laser wavelength at any given time was recorded by measuring the heterodyne beat frequency generated when the laser was mixed with a stabilized, self-referenced 250 MHz frequency comb system.

The origin of the difference between the present measurements and those from J20 is presently unknown, but there are several important differences between the two measurements. One is that the J20 measurements were performed during early 2018, approximately 1 yr before those presented here and closer in time to when the system was assembled and placed in vacuum for the first time. Stabilized FP are known to undergo some relaxation related to the spacer materials and stresses in the optical contacts between the mirrors and the spacer (Berthold et al. 1977; Dubé et al. 2009). The result is typically that the overall length of the FP shrinks over time, with an asymptotically decreasing shrink rate. Further, between the J20 measurements and those presented here, the system was shipped from the NIST lab in Colorado to the HET in Texas. It is possible that during this shipment, the system underwent a kind of "annealing" in which mechanical stresses were rapidly dissipated. A combination of the FP relaxation and potential mechanical "annealing" during shipping could explain the differences in drift rates measured.

It is also important to compare the two measurement techniques, both of which achieve excellent absolute accuracy by referencing to stabilized LFCs. The J20 measurements achieve high sampling and precision on the centroids of three individual FP resonant modes and involve illuminating the FP only selectively with the scanning lasers. The present measurements are much less accurate for any individual mode but leverage an ensemble of ~3500 modes across the HPF bandpass and involve illuminating the FP across its full bandpass (including the supercontinuum pump laser, which contains substantial optical power). These two techniques are therefore subject to different systematic errors. For example, environmentally sensitive low-amplitude parasitic etalons might hamper the centroid measurement of an individual FP mode in a way that averages out over larger groups of nearby modes. On the other hand, the present measurements may be subject to effects from the varying optical power in the supercontinuum pump laser that are not present in the J20 measurements.

5.3. The Origins of the Wavelength-dependent Drift

The most important result of this study is the specific wavelength-dependent FP drift behavior observed in Figure 7. To our knowledge, this is the most detailed long-term chromatic study of a broadband astronomical FP yet published. Importantly, the observed behavior deviates from the commonly held assumption that the long-term drift behavior of such an FP is dominated by a slow relaxation of the FP spacer material and interfaces, resulting in a gradual shrinkage of the mirror separation (and shifting of the modes toward higher frequencies/smaller wavelengths). Indeed, the average drift measured here (Figure 5) is small in magnitude (a few centimeters per second per day) and in the appropriate direction (blueshift) and is similar to the drift measured in Stürmer et al. (2017) that is attributed to the slow material relaxation. However, the present measurements confirm the result of J20 that the drift of the FP modes is not related purely by their mode indices (as described in Section 2.4) and therefore is not driven by a change in the cavity mirror separation L alone.

There are a number of possible physical origins for this effect. Noting the oscillatory nature of the variations in Figure 7, one possibility is a low-amplitude "parasitic" etalon introduced by reflections within the optical path of the FP system. Such an etalon might be introduced by, for example, reflections within the mirror substrates, off fiber terminations, or within the Bragg grating notch filter. Thus exposed to environmental or other variations, the effect on the measured spectrum might amount to a gradually changing "peak-pulling" effect that manifests as the variations we measure.

A single such parasitic etalon (or its peak-pulling effect) would (to first order) be described by Equation (2). To the extent that the material refractive index is constant, the $\Delta f_{\rm FSR}$ should be roughly constant, so any peak-pulling effect should also be roughly periodic in frequency. For reference, a parasitic etalon with $\Delta f_{\rm FSR} \approx 20$ THz would correspond to $L \approx 10\ \mu$m, and a parasitic etalon "beating" with the FP resonant modes at a frequency of ~20 THz would correspond to an L very nearly equal to the few millimeter spacing of the FP itself. Inspection of the frequency axis in Figure 7 reveals that the period of the effect varies by a factor of 3. For a single parasitic etalon, this implies a change in the $\Delta f_{\rm FSR}$ by a similar factor, which we consider unlikely. We therefore consider the impact of a single

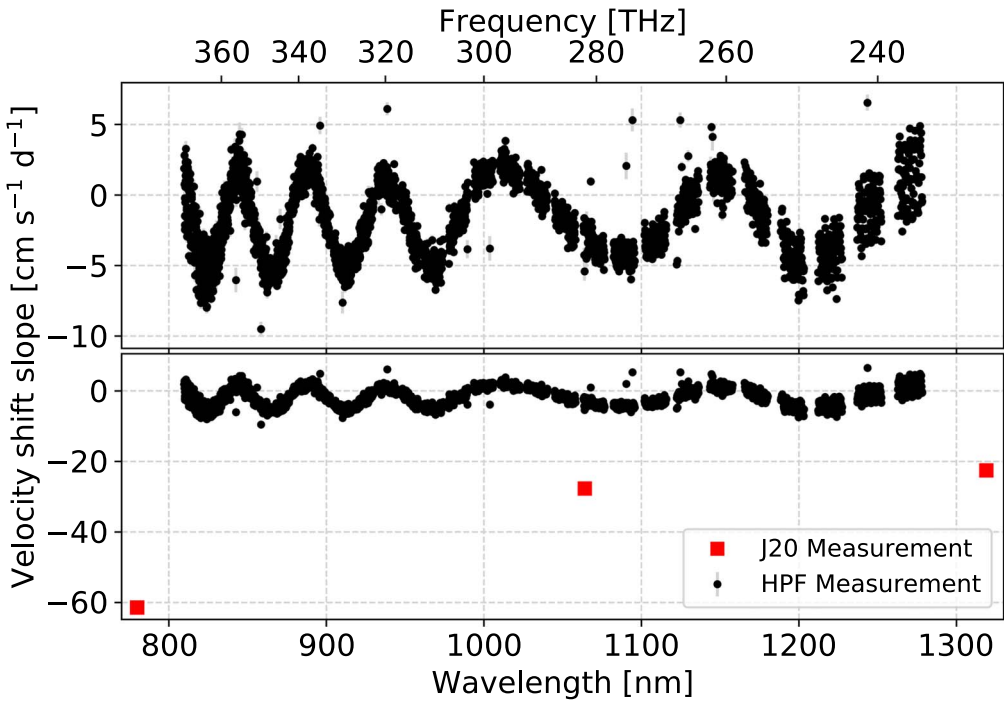

Figure 7. (Top) Fitted linear slopes of the FP mode drifts over ∼1000 epochs from 2018 November through 2019 May. The long-term linear drifts show a clear wavelength dependence and vary in both sign and magnitude. (Bottom) The same measurements but plotted along with the results from J20, showing the significantly smaller magnitude of the measured drifts of this system in its present configuration.

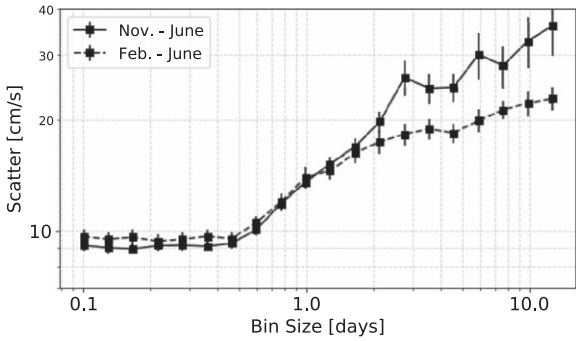

Figure 8. Measured scatter in the median velocity shift for the entire FP spectrum for a range of bin sizes over the stable period identified in Figure 5, showing the stability of the spectrum as a calibrator in the case where spectrograph drift is dominated by a velocity-like term. The dashed line shows only the second half of our full date range, which emphasizes a period where the FP drift behavior was somewhat smoother.

parasitic etalon (resulting from, for example, the Bragg grating) to be an unlikely explanation. A more complicated combination of such effects cannot be ruled out, however.

Another possibility relates to how the in-coupled light is incident on the FP. The FP is illuminated by broadband supercontinuum light, which is injected through a single-mode optical fiber and collimated by an off-axis parabolic mirror. The numerical aperture and mode pattern exiting the input fiber are wavelength-dependent, and this could, in principle, translate to different distributions of incident angles onto the FP. Coupled with long-term mechanical settling of the optical mounts, this might yield changes that manifest as the measured chromatic drifts. A nonnormal angle of incidence θ onto an etalon can be parameterized by simply substituting $L\cos\theta$ for L in Equation (1). A change in θ with this simple parameterization would cause a uniform offset of the $\Delta f_{\rm FSR}$ values, which would manifest as a (spectrally) smooth change in $\Delta f_{\rm FSR}$ or Doppler shift. The linear changes shown in Figure 7 are clearly not well described by such a change, and producing the observed pattern with an incident angle mechanism would appear to require a complicated illumination pattern that varies non-monotonically with wavelength. We consider that a change in the angle of incidence onto the FP is therefore unlikely to be the main cause of the chromatic drift behavior observed.

Other potential explanations for this drift behavior involve the FP mirror coating itself. We have shown that phase delays resulting from this coating are likely playing a measurable role in determining the FP mode spectrum (Section 4.1), so it is possible that changes in this mirror coating are driving changes in the FP mode spectrum as well. The mechanism of this hypothesized effect could involve changes in either the thicknesses or the refractive indices of the layers in the dielectric stack. An exploration of the details of this type of effect will be important for refining the design of future FP systems used for long-term high-precision Doppler spectroscopy, but this is beyond the scope of this work.

Although we do not aim for a comprehensive list of potential "ultimate" causes for the changes described above, we note that the incident optical power to the FP from the supercontinuum source is variable, being turned off and on regularly to extend the usable lifetime of the FP system. When active, the supercontinuum source illuminates the FP with ∼1 mW of optical power, and this could, in principle, impose thermal transients on downstream optical components (e.g., the FP mirrors).

Tests of such optical "heating" with a narrower-band illumination at similar power (Jennings et al. 2020) suggest that $\lesssim$1 mW illumination does not impact the mode centroid positions at the few hundred kilohertz level. However, as noted in Section 5.2, it is not clear how equivalent the present FP configuration at the HET is to the testing configuration of Jennings et al. (2020). These configurations differ in their

illumination source profile, the timing of the illumination, and potentially in the optical alignment of elements in the FP system. It is also nontrivial to accurately determine the spectral profile of the resonant optical power in the FP in the presence of the evolving spectral profiles of the illumination source and the FP itself. Thus, while we cannot yet trace the impact of optical heating in our measurements, this effect remains an important area of investigation for this system.

Finally, it is important to note the potential role of subtle changes in the HPF IP in the precise long-term monitoring of the FP drift. Even in a highly stabilized spectrograph like HPF, the IP is expected to evolve over the long term. This evolution could, in principle, drive the appearance of wavelength-dependent centroid drifts in the FP monitoring because this analysis assumes a simple and unchanging line profile.

Indeed, information from the HPF LFC and FP have been central to measuring the long-term drift of the HPF IP, and the dominant terms of the drift correction address a slow evolution in the camera magnification and a relaxation of the HPF interpixel capacitance (R. C. Terrien et al. 2021, in preparation), which both appear as IP changes in the non-drift-corrected data. Importantly, these effects in the non-drift-corrected data are highly correlated with the position of a particular wavelength in the HPF echelle format due to the dependence of optical aberration or detector behavior on focal plane position. Because no such correlation remains in the wavelength-dependent drifts we present here (e.g., Figure 7), we consider it unlikely that the evolution of the IP is the source of these drifts. We cannot rule out, however, all potential IP changes, and efforts are underway to leverage the unique flexibility of the LFC to measure these changes directly.

5.4. Implications for Spectrograph Calibration

The chromatic structure in the behavior of this FP system has significant implications for the calibration of Doppler spectrographs that aim at $\sim$10 cm s^{-1} precision. With an FP calibrator, such systems may indeed be limited by the interplay of these chromatic variations and other instrumental systematics related to the detector or spectrograph drift. Our results show that the extraction of the highest-precision Doppler shifts will require extensive characterization of the spectral and temporal characteristics of the spectrograph and the FP calibration system. The present FP calibration system clearly does not conform to the expectation that the time evolution of the FP mode spectrum can be predicted simply by projecting the evolution of a single mode (or a small number of modes) under the assumption of the dominance of the change in mirror separation L, so we caution against relying on this assumption in other systems.

At the level of $\sim$1 m s^{-1}, the results in Section 5.1 show that this FP system performs well as a stand-alone calibrator for time periods up to at least $\sim$10 days. This system is in use as a "holdover" calibrator at HPF for times when the LFC is unavailable and is used to derive the drift model parameters as explained in Section 5.4.1. We note that the drift behavior of this FP and of HPF are sufficiently independent that the process of refining the FP mode measurements resulted in the discovery and subsequent implementation of the linear and readout channel terms in our HPF drift model.

Our results also emphasize the need for similar in situ characterization of other systems, in view of the significant differences between our present findings and those presented in J20. We also note that there are a wide variety of astronomical FP calibrator designs in use already, many of which do not (yet) have the benefit of an accompanying LFC to evaluate their behavior in situ. A nonexhaustive list of the differences in design choices among these systems includes different illumination sources (laser-driven light sources, supercontinuum sources), different coupling (single-mode, multimode), different resonator spacer and mirror substrate materials (ULE, Zerodur), and different mirror coating properties (soft or hard coatings, low or high finesse). This variety in the face of a similar goal—wavelength calibration of high-resolution Doppler spectrographs—suggests that these designs can be optimized and improved. The results presented here are pertinent to our specific design; further study will be necessary to establish the extent to which other designs are subject to this type of behavior, but this work serves as a cautionary note about assuming a simple relationship relating the drift of different FP modes.

5.4.1. Calibration of HPF Using the FP

The HPF FP has been used as a "holdover" calibration source for brief (one to a few days) periods when the LFC is unavailable. Here we briefly discuss our current method for using the FP in this context.

As described in Section 3.2, our HPF drift model comprises a second-order polynomial (per order), along with discrete offsets for each readout channel. In the short term (days to weeks), the drift is dominated by the zeroth-order term (i.e., pixel offset) of the polynomial, and the higher-order terms vary more slowly on a timescale of weeks to months.

Due to the wavelength-dependent drift of the FP we uncovered in this study, as well as the small time baselines when the FP is needed, we presently only rely on the FP to estimate the zeroth-order term of the polynomial. To estimate the value of this term at a given epoch, we follow a similar procedure to that described in Section 3.2. We first create a high-S/N FP template spectrum by combining multiple epochs of wavelength-calibrated FP spectra. This template spectrum is then fit via least-squares optimization to a target epoch, with a single pixel offset as the only free parameter. The resulting pixel shift term constrains the zeroth-order term of the drift model for a single point in time.

This constraint is subject to a slow drift, as indicated by the FP drift measurements in Figure 5. Figure 9 shows the difference between the zeroth-order term in the drift model (LFC c_0) and that derived directly from the FP (FP c_0), as measured in our wavelength calibration. A smooth (second-order) polynomial fit (also shown in Figure 9) is sufficient to correct the FP-based offset measurement for this long-term drift.

6. Conclusion

Broadband FP offer an attractive alternative to classical RV calibration sources, with a rich spectrum of sharp features at regular frequency spacing. These devices have many distinct advantages over atomic emission lamps and molecular absorption cells and are intrinsically significantly simpler than the "gold standard" of an LFC. However, their use for calibration at the $\sim$10 cm s^{-1} level, and as long-term calibration sources in their own right, remains largely untested.

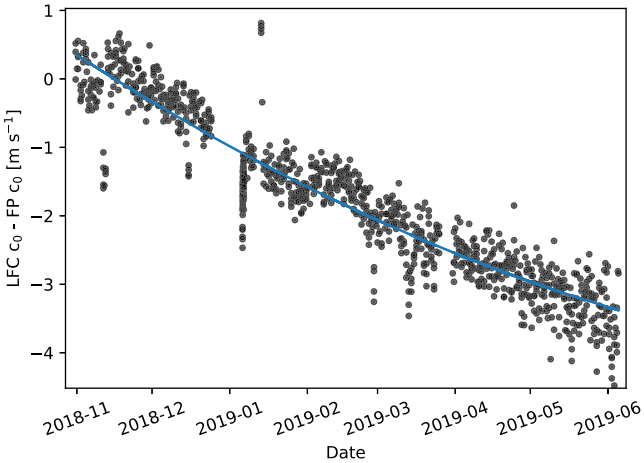

Figure 9. Difference between the zeroth-order term in the LFC drift model and the corresponding term estimated from the FP spectrum. As in Figure 5, there are a small number of times when environmental or illumination changes induced short-term deviations. A second-order polynomial approximation that is used to convert from pixel drift from the FP measurement to the zeroth-order term in the drift model is shown by the blue curve. For clarity, the *y*-axis is shown in units of velocity (m s^{-1}) instead of the scaled pixel units.

We performed an extensive spectral study of the stability of a FP calibration system for the HPF spectrograph. We leveraged the long-term accuracy and precision provided by the dedicated HPF LFC calibration system. Using the LFC-based wavelength calibrations, we located and tracked the centroids of each of ∼3500 FP resonant modes (Figure 2) over a period of 6 months, confirming that we were photon-limited at any one epoch (Figure 3). Our main results for this FP system are as follows.

1. Chromatic differences in the single-epoch mode spectrum of the FP are traceable to the resonator mirror coating (Figure 4).
2. The median drift of the FP supports a calibration precision of $\lesssim$30 cm s^{-1} over $\lesssim$10 days and $\lesssim$10 cm s^{-1} over a night (Figures 5 and 8).
3. The drift rate of modes across the FP spectrum is a complicated function of wavelength, varying in both sign and magnitude (Figures 6 and 7). These drift rates do not conform to the expectation of a simple model (a gradually shrinking resonator mirror separation) that has been suggested to explain long-term drifts in other systems.

We discussed potential origins for this chromatic behavior, including parasitic etalons, illumination angle changes, and changes in the FP mirror coating. We also discussed potential origins for the significant differences from measurements of this system presented in J20.

These measurements represent an important step forward in characterizing FP etalons for precise RV spectrograph calibration. Our results highlight the need for extensive characterization of FP systems over long timescales and suggest that a multiwavelength study of the FP stability in situ may be required for enabling the highest levels of RV precision.

The authors wish to thank the anonymous referee for insightful comments that substantially improved this work. The authors also wish to thank Light Machinery for graciously sharing the resonator mirror phase-curve data. The authors also thank Mark Notcutt and Mike Grisham at Stable Laser Systems for sharing their knowledge about FP systems and their insight into the design of this system. R.C.T. thanks Bruce Duffy for his advice and support with managing computational resources at Carleton.

This work was partially supported by funding from the Center for Exoplanets and Habitable Worlds. The Center for Exoplanets and Habitable Worlds is supported by the Pennsylvania State University, the Eberly College of Science, and the Pennsylvania Space Grant Consortium. These results are based on observations obtained with the Habitable Zone Planet Finder spectrograph on the HET. We acknowledge support from NIST and the NIST-on-a-Chip Program; NSF grants AST 1006676, AST 1126413, AST 1310875, AST 1310885, ATI 1006676, ATI 2009889, ATI 2009982, and ATI 2009955; NSF OD OIA MRI1626251, and the NASA Astrobiology Institute (NNA09DA76A) in our pursuit of precision radial velocities in the NIR. We acknowledge support from the Heising-Simons Foundation via grant 2017-0494. The HET is a joint project of the University of Texas at Austin, the Pennsylvania State University, Ludwig-Maximilians-Universität München, and Georg-August Universität Gottingen. The HET is named in honor of its principal benefactors, William P. Hobby and Robert E. Eberly.

The HET collaboration acknowledges the support and resources from the Texas Advanced Computing Center. We thank the resident astronomers and telescope operators at the HET for the skillful execution of our observations with HPF.

Computations for this research were performed on the Pennsylvania State University's Institute for Computational and Data Sciences' Roar supercomputer, including the CyberLAMP cluster supported by NSF grant MRI1626251.

Facility: HET(HPF).

Software: `astropy` (Astropy Collaboration et al. 2013, 2018), `numpy` (van der Walt et al. 2011), `scipy` (Virtanen et al. 2020), `HxRGproc` (Ninan et al. 2018), `matplotlib` (Hunter 2007), `GNU Parallel` (Tange 2011).

ORCID iDs

Ryan C. Terrien https://orcid.org/0000-0002-4788-8858
Joe P. Ninan https://orcid.org/0000-0001-8720-5612
Scott A. Diddams https://orcid.org/0000-0002-2144-0764
Suvrath Mahadevan https://orcid.org/0000-0001-9596-7983
Samuel Halverson https://orcid.org/0000-0003-1312-9391
Chad Bender https://orcid.org/0000-0003-4384-7220
Connor Fredrick https://orcid.org/0000-0002-0560-1433
Jeff Jennings https://orcid.org/0000-0002-7032-2350
Andrew J. Metcalf https://orcid.org/0000-0001-5000-1018
Andrew Monson https://orcid.org/0000-0002-0048-2586
Arpita Roy https://orcid.org/0000-0001-8127-5775
Christian Schwab https://orcid.org/0000-0002-4046-987X
Guđmundur Stefánsson https://orcid.org/0000-0001-7409-5688

References

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Bauer, F. F., Zechmeister, M., Kaminski, A., et al. 2020, A&A, 640, A50
Berthold, J. W., Jacobs, S. F., & Norton, M. A. 1977, Metro, 13, 9
Blackman, R. T., Fischer, D. A., Jurgenson, C. A., et al. 2020, AJ, 159, 238
Blank, R., Anglin, S., Beletic, J. W., et al. 2012, Proc. SPIE, 8453, 845310

Bouchy, F., Pepe, F., & Queloz, D. 2001, A&A, 374, 733
Brault, J. W. 1987, AcMik, 1, 215
Butler, R. P., Vogt, S. S., Laughlin, G., et al. 2017, AJ, 153, 208
Cañas, C. I., Stefansson, G., Kanodia, S., et al. 2020, AJ, 160, 147
Cersullo, F., Coffinet, A., Chazelas, B., Lovis, C., & Pepe, F. 2019, A&A, 624, A122
Cersullo, F., Wildi, F., Chazelas, B., & Pepe, F. 2017, A&A, 601, A102
Damasso, M., Sozzetti, A., Lovis, C., et al. 2020, A&A, 642, A31
Das, T., Banyal, R. K., Kathiravan, S., Sivarani, T., & Ravindra, B. 2018, Proc. SPIE, 10702, 107026A
Devoe, R. G., Fabre, C., Jungmann, K., Hoffnagle, J., & Brewer, R. G. 1988, PhRvA, 37, 1802
Diddams, S. A., Vahala, K., & Udem, T. 2020, Sci, 369, eaay3676
Donati, J. F., Kouach, D., Moutou, C., et al. 2020, MNRAS, 498, 5684
Dubé, P., Madej, A. A., Bernard, J. E., Marmet, L., & Shiner, A. D. 2009, ApPhB, 95, 43
Fischer, D. A., Anglada-Escude, G., Arriagada, P., et al. 2016, PASP, 128, 066001
Fortenbach, C. D., & Dressing, C. D. 2020, PASP, 132, 054501
Gaudi, S., Blackwood, G., Howard, A., et al. 2019, BAAS, 51, 232
Halverson, S., Mahadevan, S., Ramsey, L., et al. 2014, PASP, 126, 445
Halverson, S., Roy, A., Mahadevan, S., et al. 2015, ApJ, 806, 61
Hirano, T., Kuzuhara, M., Kotani, T., et al. 2020, PASJ, 72, 93
Horne, K. 1986, PASP, 98, 609
Hunter, J. D. 2007, CSE, 9, 90
Jennings, J., Halverson, S., Terrien, R., et al. 2017, OExpr, 25, 15599
Jennings, J., Terrien, R., Fredrick, C., et al. 2020, OSA Continuum, 3, 1177
Kanodia, S., Cañas, C. I., Stefansson, G., et al. 2020, ApJ, 899, 29
Kanodia, S., Mahadevan, S., Ramsey, L. W., et al. 2018, Proc. SPIE, 10702, 107026Q
Kempton, E. M. R., Bean, J. L., Louie, D. R., et al. 2018, PASP, 130, 114401
Kopparapu, R. K., Ramirez, R., Kasting, J. F., et al. 2013, ApJ, 765, 131
Lee, E. J. 2019, ApJ, 878, 36
Lichten, W. 1985, JOSAA, 2, 1869
Lichten, W. 1986, JOSAA, 3, 909
Mahadevan, S., Ramsey, L. W., Terrien, R., et al. 2014, Proc. SPIE, 9147, 91471G
McCracken, R. A., Charsley, J. M., & Reid, D. T. 2017, OExpr, 25, 15058
Metcalf, A. J., Anderson, T., Bender, C. F., et al. 2019a, Optica, 6, 233
Metcalf, A. J., Fredrick, C. D., Terrien, R. C., Papp, S. B., & Diddams, S. A. 2019b, OptL, 44, 2673
Murphy, M. T., Locke, C. R., Light, P. S., Luiten, A. N., & Lawrence, J. S. 2012, MNRAS, 422, 761
Murphy, M. T., Udem, T., Holzwarth, R., et al. 2007, MNRAS, 380, 839
Nave, G., Kerber, F., Den Hartog, E. A., & Lo Curto, G. 2018, Proc. SPIE, 10704, 1070407
Ninan, J. P., Bender, C. F., Mahadevan, S., et al. 2018, Proc. SPIE, 10709, 107092U
Ninan, J. P., Mahadevan, S., Stefansson, G., et al. 2019, JATIS, 5, 041511
Ninan, J. P., Stefansson, G., Mahadevan, S., et al. 2020, ApJ, 894, 97
Pepe, F. A., Cristiani, S., Rebolo Lopez, R., et al. 2010, Proc. SPIE, 7735, 77350F
Perot, A., & Fabry, C. 1899, ApJ, 9, 87
Petersburg, R. R., Ong, J. M. J., Zhao, L. L., et al. 2020, AJ, 159, 187
Probst, R. A. 2015, Dissertation, Laser Frequency Combs for Astronomy, Ludwig-Maximilians-Universität München, https://edoc.ub.uni-muenchen.de/18701/1/Probst_Rafael.pdf
Probst, R. A., Lo Curto, G., Avila, G., et al. 2014, Proc. SPIE, 9147, 91471C
Reiners, A., Banyal, R. K., & Ulbrich, R. G. 2014, A&A, 569, A77
Robertson, P., Anderson, T., Stefansson, G., et al. 2019, JATIS, 5, 015003
Robertson, P., Stefansson, G., Mahadevan, S., et al. 2020, ApJ, 897, 125
Roy, A., Halverson, S., Mahadevan, S., & Ramsey, L. W. 2014, Proc. SPIE, 9147, 91476B
Schwab, C., Feger, T., Stürmer, J., et al. 2018, Proc. SPIE, 10702, 1070272
Schwab, C., Rakich, A., Gong, Q., et al. 2016, Proc. SPIE, 9908, 99087H
Sirk, M. M., Wishnow, E. H., Weisfeiler, M., et al. 2018, Proc. SPIE, 10702, 107026F
Stefansson, G., Cañas, C., Wisniewski, J., et al. 2020a, AJ, 159, 100
Stefansson, G., Hearty, F., Robertson, P., et al. 2016, ApJ, 833, 175
Stefansson, G., Kopparapu, R., Lin, A., et al. 2020b, AJ, 160, 259
Stefansson, G., Mahadevan, S., Maney, M., et al. 2020c, AJ, 160, 192
Stürmer, J., Seifahrt, A., Schwab, C., & Bean, J. L. 2017, JATIS, 3, 025003
Tange, O. 2011, ;login:, 36, 42, https://www.usenix.org/system/files/login/articles/105438-Tange.pdf
Udry, S., Dumusque, X., Lovis, C., et al. 2019, A&A, 622, A37
van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
Wildi, F., Pepe, F., Chazelas, B., Lo Curto, G., & Lovis, C. 2010, Proc. SPIE, 7735, 77354X