



Solar Contamination in Extreme-precision Radial-velocity Measurements: Deleterious Effects and Prospects for Mitigation

Arpita Roy^{1,2,3,15} , Samuel Halverson^{4,5,16} , Suvrath Mahadevan^{2,3} , Gudmundur Stefansson^{2,3,17} , Andrew Monson^{2,3}, Sarah E. Logsdon⁶ , Chad F. Bender⁷ , Cullen H. Blake⁸ , Eli Golub⁶, Arvind Gupta^{2,3}, Kurt P. Jaehnig⁹, Shubham Kanodia^{2,3} , Kyle Kaplan⁷, Michael W. McElwain¹⁰, Joe P. Ninan^{2,3} , Jayadev Rajagopal⁶, Paul Robertson¹¹ , Christian Schwab¹² , Ryan C. Terrien¹³ , Sharon Xuesong Wang¹⁴, Marsha J. Wolf⁹, and Jason T. Wright^{2,3} 

¹ Department of Astronomy, California Institute of Technology, Pasadena, CA 91125, USA; aroy@caltech.edu

² Department of Astronomy & Astrophysics, 525 Davey Laboratory, Pennsylvania State University, University Park, PA, 16802, USA

³ Center for Exoplanets and Habitable Worlds, 525 Davey Laboratory, Pennsylvania State University, University Park, PA, 16802, USA

⁴ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

⁵ Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA, USA

⁶ NSF's National Optical-Infrared Astronomy Research Laboratory, Tucson, AZ 85726, USA

⁷ Steward Observatory, The University of Arizona, 933 N. Cherry Ave, Tucson, AZ 85721, USA

⁸ Department of Physics and Astronomy, University of Pennsylvania, 209 S. 33rd Street, Philadelphia, PA 19104, USA

⁹ Department of Astronomy, University of Wisconsin, Madison, WI 53706, USA

¹⁰ Exoplanets and Stellar Astrophysics Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

¹¹ Department of Physics & Astronomy, The University of California, Irvine, Irvine, CA 92697, USA

¹² Department of Physics and Astronomy, Macquarie University, Balaclava Road, North Ryde, NSW 2109, Australia

¹³ Carleton College, One North College Street, Northfield, MN 55057, USA

¹⁴ The Observatories of the Carnegie Institution of Washington, 813 Santa Barbara Street, Pasadena, CA 91101, USA

Received 2019 October 29; revised 2020 February 3; accepted 2020 February 18; published 2020 March 18

Abstract

Solar contamination, due to moonlight and atmospheric scattering of sunlight, can cause systematic errors in stellar radial velocity (RV) measurements that significantly detract from the $\sim 10 \text{ cm s}^{-1}$ sensitivity required for the detection and characterization of terrestrial exoplanets in or near habitable zones of Sun-like stars. The addition of low-level spectral contamination at variable effective velocity offsets introduces systematic noise when measuring velocities using classical mask-based or template-based cross-correlation techniques. Here we present simulations estimating the range of RV measurement error induced by uncorrected scattered sunlight contamination. We explore potential correction techniques, using both simultaneous spectrometer sky fibers and broadband imaging via coherent fiber imaging bundles, that could reliably reduce this source of error to below the photon-noise limit of typical stellar observations. We discuss the limitations of these simulations, the underlying assumptions, and mitigation mechanisms. We also present and discuss the components designed and built into the NEID (NN-EXPLORE Exoplanet Investigations with Doppler spectroscopy) precision RV instrument for the WIYN 3.5 m telescope, to serve as an ongoing resource for the community to explore and evaluate correction techniques. We emphasize that while “bright time” has been traditionally adequate for RV science, the goal of 10 cm s^{-1} precision on the most interesting exoplanetary systems may necessitate access to darker skies for these next-generation instruments.

Unified Astronomy Thesaurus concepts: [Exoplanet astronomy \(486\)](#); [Radial velocity \(1332\)](#); [High resolution spectroscopy \(2096\)](#); [Sky brightness \(1462\)](#); [Astronomy data analysis \(1858\)](#)

1. Introduction

The field of radial velocity (RV) exoplanet detection is engaged in an ongoing battle for ever higher precision measurements, motivated by the urge to detect Earth-mass planets in the habitable zone (Kopparapu et al. 2013). Such terrestrial-mass planets orbiting Sun-like stars induce Doppler reflex signals of $\sim 10 \text{ cm s}^{-1}$, which nominally defines the precision goal for many next-generation instruments and surveys. This realm of “extreme” precision spectroscopy requires substantial advances in both instrument design and analysis techniques (Fischer et al. 2016) because previously negligible sources of error become significant at these levels of scrutiny. A number of heretofore uncharacterized phenomena that affect spectral line profiles must now be extensively studied

and mitigated. Several next-generation extreme-precision spectrographs (see Wright & Robertson 2017 for a complete list), including the GMT Consortium Large Earth Finder (G-CLEF; Podgorski et al. 2014), NN-EXPLORE Exoplanet Investigations with Doppler spectroscopy (NEID; Halverson et al. 2016), and the Keck Planet Finder (KPF; Gibson et al. 2018) are thus being built with a bottom-up systems engineering approach based on comprehensive RV error budgets, which can be valuable tools for predicting instrument performance.

Spectral contamination remains an area of grave concern in formulating RV error estimates because it significantly detracts from our ability to measure intrinsic stellar line profiles (Pepe & Lovis 2008). The worst case of contamination occurs when the superimposed light contains spectral features, which can both pollute individual line shapes and induce a spurious velocity signal across the ensemble of lines. This is exacerbated by the fact that relative shifts between contaminant and source spectra are often introduced by the motion of the star and the

¹⁵ Robert A. Millikan Postdoctoral Fellow.

¹⁶ NASA Sagan Fellow.

¹⁷ NASA Earth and Space Science Fellow.




Figure 1. Toy model for illustration of “peak pulling” of the measured cross-correlation function (CCF) due to scattered solar contamination. Top: A relatively low amplitude “sky” solar spectrum (colored curves) is added to a synthetic G2 dwarf “target” spectrum (gray) at different absolute velocity offsets. In this illustration, we exaggerate “sky” to be 10% of the target brightness to highlight the impact. Spectral lines used to construct the aggregate CCF are highlighted in red. Bottom: The effective velocity offset of the combined spectrum relative to the target spectrum is calculated by fitting a Gaussian to the peak of the CCF. As the cross-correlation mask encounters solar features, spurious peaks are introduced into the CCF that can lead to significant velocity offsets if left uncorrected (see inset figures).

Earth-based observatory over time. In the case of background sky brightness interfering with RV measurements, there are no hardware solutions to entirely eliminate this effect for seeing-limited instruments. In fact, the magnitude of the contaminant must often be diagnosed from the same observations that need to be corrected for the effect. Mitigating this by trying to observe only in dark time is possible but impractical, given competition with other subfields like extragalactic and cosmological studies.

Blended light from companions or inconvenient background objects are known to be capable of producing deceptive false positives, deserving careful treatment (e.g., Wright et al. 2013). Unlike the incidental nature of these contaminants, however, we focus here on scattered or reflected sunlight, which is essentially always present in observed spectra from ground-based instruments at some level. The magnitude of this contamination depends on several factors, including sky brightness, target–moon separation, lunar phase, ecliptic latitude, zenith angle, and phase of the solar cycle (Krisciunas 1997). Thus far, this error source has largely been mitigated by avoiding (1) twilight observations, (2) observations during a full moon or when the target is close to the moon, and (3) observations during cloudy skies or cirrus clouds (Pepe & Lovis 2008; Seager 2010). Ensuring that the sky background is faint (>7 – 10 magnitudes

fainter as a rule of thumb) has been largely adequate for current instruments operating in the $\geq 1 \text{ m s}^{-1}$ precision regime (Pepe & Lovis 2008).

The feasibility of performing solar contamination correction is yet to be explored for Doppler measurements at the 10 cm s^{-1} level. The presence of scattered sunlight in observed spectra can be a two-fold source of error in weighted mask-based (Pepe et al. 2002) or high signal-to-noise ratio (S/N) template-based cross-correlation for RV measurement, affecting both “peak pulling” and the introduction of more complex and time-variable structure in the cross-correlation function (CCF) or χ^2 space (Figure 1).

Here we carefully consider the impact of solar light contamination on RV measurements and explore the mitigation possible with a simultaneous sky fiber. Several upcoming instruments, for example, NEID (Schwab et al. 2016) and KPF (Gibson et al. 2018), include a dedicated sky fiber for the correction of telluric absorption and emission lines and the removal of scattered sunlight. However, the sky fiber is highly spectrally dispersed in these instruments, making faint levels of solar contamination difficult to measure even when there is a consequential impact on RV precision.

For cases of very faint sky with low S/N in the sky fiber, we examine the viability of substituting the sky fiber with broadband




Figure 2. NEID high-resolution (HR) mode port adapter fiber head that collects light from the telescope for transportation to the spectrometer. Left: image of the as-built fiber head being installed with NEID, taken through a Zygo Optical Profilometer, which uses white light interferometry to obtain a high-precision image. Middle: diagram showing the relative placement of the coherent fiber bundles (CFBs) around the science fiber, as well as the position of the sky fiber (unused sky fibers provide redundancy). Right: Each CFB is composed of several thousand fiber cores, 3500 for NEID specifically. Magnified CFB image from the manufacturer (Schott).

sky images from coherent fiber bundles (CFBs). CFBs are included in the NEID fiber head primarily for target acquisition purposes, but present an interesting alternative path for sky brightness measurements. We attempt to keep our simulation fairly general across seeing-limited next-generation spectrographs. Where specification is necessary, we base our parameters on the NEID spectrograph, currently being commissioned on the 3.5 m WIYN telescope at Kitt Peak. As a result, these calculations are readily adaptable for other instruments with adjustments for telescope aperture, fiber size and shape, instrument bandpass and resolution, and expected sky conditions.

We examine the impact of sky brightness on RV measurements in Section 2, including descriptions of our simulations, measurement technique, and the magnitude of the uncorrected error. We present three correction strategies in Section 3 and detail their inclusion in the NEID design in Section 4. We discuss our results and some nuanced instrument-specific effects in Section 5, and we conclude in Section 6.

2. Impact of Sky Brightness on RV Measurements

Here we set out to predict the deleterious effects of solar contamination on RV precision and prepare our mitigation techniques that leverage the sky fiber, in close anticipation of detailed tests on NEID. For this exercise, it is important to create realistic spectra that accurately model the absolute flux and noise inherent to different levels of sky brightness. Instruments like NEID use a dedicated sky fiber to directly sample the solar contamination spectrum at a resolution identical to the target spectrum. Since the science and sky fibers are the same size and at relatively small separations (22'' on-sky in the case of NEID, Figure 2), they should contain very similar amounts of sky flux. Both fibers should also have comparable throughput or be flux calibrated to high degrees of confidence. This similarity is quantified by measurement on NEID, but must be verified for individual instruments.

2.1. Simulating Science and Sky Fiber Spectra

We simulate both “sky” and “star” spectra using a synthetic high-resolution solar spectrum ($R \sim 500,000$; $T_{\text{eff}} = 5800$ K; $\log(g) = 4.5$; $[\text{Fe}/\text{H}] = 0.0$) computed with the PHOENIX

stellar atmosphere models (Husser et al. 2013). We deliberately opt to use the same synthetic spectrum for both sources (i.e., the case when observing a Sun-like star) because Doppler errors are maximized when the contaminant is similar to the observed target, thus rendering a worst-case scenario. To ensure accurate photon counts corresponding to different stellar magnitude levels, we scale the model spectrum using the absolute flux distribution of solar analogs based on spectrophotometry from the *Hubble Space Telescope* Faint Object Spectrograph (Colina & Bohlin 1997). The spectrum is then attenuated for atmospheric extinction at Kitt Peak and converted to flux incident on the 3.5 m WIYN telescope aperture.

Wavelength-dependent seeing losses through an octagonal fiber are calculated using a median WIYN point-spread function (PSF) with pODI (the precursor to ODI, the One Degree Imager at WIYN) in the r -band provided by NSF’s National Optical-Infrared Astronomy Research Laboratory, ahead of extensive characterization with NEID. This PSF is highly non-Gaussian and shows an extended aureole, as expected from Kolmogorov theory, and we model it with a 3-Moffat function as recommended by Racine (1996). To this we apply the complete NEID spectrometer throughput model, which accounts for all subsequent losses in the system.

These steps are repeated for the “sky” spectrum, after converting a desired sky brightness (in mag arcsec $^{-2}$) to a V-band magnitude for the 0.92 NEID fibers for the high-resolution mode, keeping in mind that the sky fiber does not suffer seeing losses, due to the uniform illumination on the fiber face. While the sky fiber is immune to atmospheric dispersion effects, note that it is important to have considered performance requirements on the atmospheric dispersion corrector to trace photon incidence rates through the science fiber and ensure reliable chromatic count rates across the instrument bandpass (Logsdon et al. 2018; Wehbe et al. 2020).

The synthetic spectra are convolved with a Gaussian kernel to match the NEID instrument spectral resolution ($R = 100,000$) and binned to the approximate point-spread function sampling (5 pixels FWHM). The spectra are further distributed into 68 echelle orders (3800–6800 Å) based on the predicted NEID order mapping. Even though NEID extends to 9300 Å, we limit ourselves to ≤ 7000 Å because this is the region that contains

most of the Doppler information for Sun-like stars, and to which our best cross-correlation masks are currently limited (from instruments classically calibrated with ThAr lamps). Each order is multiplied by an artificial blaze function derived from measured grating response curves of an R4 echelle. This ensures that the edges of the orders are realistically downweighted with respect to the centers and the associated noise levels are generated correctly. To calculate the photons in each pixel of the extracted spectrum, we assume the 2D flux profile across the fiber trace is collapsed in the cross-dispersion direction (fiber image has a width of ~ 5 pixels) and scale by a standard exposure time of 20 minutes for all simulations. Poissonian photon noise and Gaussian read noise ($4e^-$ per pixel based on the performance of the NEID CCD, allotting gain = 1) are added to both the target and sky spectra. The inclusion of both photon noise and read noise is critical because at $R \sim 100,000$ the photon flux rate in the sky fiber quickly becomes comparable to the read noise as sky brightness decreases. We simulate the science fiber by adding the “scattered sunlight” and the stellar spectra at realistically varying velocity separations based on the extent of barycentric motion ($\pm 30 \text{ km s}^{-1}$) to mimic a range of observational epochs. The sky fiber contains only the scattered sunlight component.

2.2. Measuring RV Errors in Simulated Spectra

The simulated science and sky fiber spectra are passed through our RV analysis code, which has been extensively vetted at the $< 1 \text{ m s}^{-1}$ level on both PRL Advanced Radial-velocity Abu-sky Search (PARAS) and High Accuracy Radial-velocity Planet Searcher (HARPS) data (Roy et al. 2016). CCFs are calculated using a standard stellar mask-based technique (Baranne et al. 1996; Pepe et al. 2002; Chakraborty et al. 2014) for each order, then summed to form an aggregate CCF. Effective velocity offsets are measured by fitting a Gaussian to the peak of the aggregate CCF. We fit $\pm 20 \text{ km s}^{-1}$ around the expected peak value, similar to our implementation in the PARAS pipeline, since fitting smaller intervals around the peak increases RV error.

Note that echelle blaze functions are not removed during processing because they are static in this idealized case and preserve the true photon counts associated with noise generation. This also means that there is no necessity for order reweighting in this simulation. In reality, of course, changes in illumination due to clouds, seeing, telescope focus drifts, airmass, and other observing parameters render a slightly variable blaze response and order weighting that must be either removed or corrected during spectral processing (e.g., Anglada-Escudé & Butler 2012).

Beyond the proven performance of our code on real data from current RV instruments, here we want to be sensitive to much smaller errors corresponding to the precision of next-generation spectrographs. Hence, we begin by checking our algorithmic error on “perfect” noiseless simulations. We find that we can both introduce RV offsets and recover these injected velocities at the $< 10^{-4} \text{ cm s}^{-1}$ level. This ensures that we can robustly simulate and investigate contamination errors at the $\ll 10 \text{ cm s}^{-1}$ level of interest, while treating algorithmic error as negligible.

2.3. Uncorrected Solar Contamination Errors

Our set of simulated scenarios encompass stars with $V_{\text{mag}} = 6\text{--}12$ and sky brightnesses from 16 to 22 mag arcsec $^{-2}$. This is




Figure 3. Radial-velocity error caused by solar contamination, as a function of velocity separation, for a given stellar V_{mag} and sky brightness. Top panels: The contaminant sky spectrum is shifted across the full range of possible barycentric velocities ($\pm 30 \text{ km s}^{-1}$) against the target star spectrum in its rest frame. Here we show the resulting CCFs for star and sky; the sky CCF is multiplied by 10 for visibility. Bottom: The RV error is low when the star is well separated from the sky spectrum in velocity space. RV error is highest when the separation between star and sky is $\sim 4 \text{ km s}^{-1}$ for NEID, as the contaminant peak passes through the wings of the primary peak. If the star and sky spectra are perfectly aligned, there is essentially no RV impact (although note that any scrutiny of line shape metrics might be affected).

based on typical targets for extreme-precision RV measurements on NEID and the range of conditions at Kitt Peak National Observatory (KPNO). KPNO dark-time sky brightness is $V \sim 21.45 \text{ mag arcsec}^{-2}$ at zenith, but increases rapidly to $V \sim 20.9 \text{ mag arcsec}^{-2}$ (airmass = 2) in the direction of Tucson (Massey & Foltz 2000; Neugent & Massey 2010). During “gray time” at the WIYN site (> 3 days from full moon), additional moon contamination can add up to $1.8 \text{ mag arcsec}^{-2}$ in the V-band (Walker 1991). We simulate this full range of overall sky brightness, even though some of it is certainly from airglow and city lights, which interact with the target spectrum differently than the scattered sunlight spectrum.

The contaminating spectra are shifted in velocity relative to the target spectrum to span the full range of potential velocity offsets due to the barycentric motion of the Earth, approximately $\pm 30 \text{ km s}^{-1}$. As shown in Figure 3, the contaminant affects maximum peak pulling (and correspondingly RV error) when it begins to blend into the wings of the primary CCF peak. The separation that causes the highest RV error is set principally by the resolution of the instrument, and to a lesser degree by choices like mask width and the extent of the peak being fitted during RV measurement. For NEID this is at approximately $\pm 4 \text{ km s}^{-1}$, close to the FWHM of the CCF. However, this is not the only velocity separation of concern, because there are also smaller ancillary peaks from the sky spectrum interacting with the target spectrum that cause errors at the few percent level of the worst case (e.g., at $\pm 20 \text{ km s}^{-1}$ in Figure 3), implying that contamination correction cannot be easily dismissed at any velocity separation. Since this is a decidedly non-Gaussian error, and for the sake of estimating the worst-case impact, here we focus on the maximum absolute RV error for each combination of star and sky brightness (i.e., the maximum values in Figure 3, bottom panel). We generate 1000 realizations of spectra with realistic noise properties for each scenario.

If left uncorrected, solar contamination can cause substantial RV variations that overwhelm other aspects of the error budget. Figure 4 shows the maximum RV error caused by a range of




Figure 4. Uncorrected radial-velocity errors introduced by the presence of contaminating moonlight or scattered sunlight, for a range of stellar V -band magnitudes and sky brightness levels. Note that these are worst-case errors, calculated at the velocity separation between the star and the sky that causes maximum RV displacement. Top: simulations with realistic noise, including photon noise, read noise, and solar contamination, showing the full measured RV error. The full uncorrected error is separated into two categories in the bottom panels. Bottom left: isolating the error due to solar contamination. The dashed white line illustrates the error incurred following the traditional practice of observing when the sky is ~ 10 magnitudes fainter than the target. Bottom right: contribution of all other error sources, illustrating the achievable noise limit for each stellar brightness.

sky brightnesses for typical NEID targets. The top panel shows the full uncorrected error (mean $+1\sigma$ uncertainty over 1000 realizations), which is the realistic limit imposed by the combined effect of solar contamination, photon noise, read noise, and the mask cross-correlation technique. We separate these errors into two categories in the bottom panel.

In the bottom right panel, we isolate the error caused by solar contamination alone, which contributes the bulk of the mean error in the top panel. In the bottom left panel, we show the achievable noise floor for each stellar magnitude as derived from 1000 injection recovery runs on uncontaminated stellar spectra. This encapsulates all other error sources in the simulation. It is dominated by photon and read noise, but algorithmic errors are also naturally included during RV measurement.

The traditional practice of observing when the sky is ~ 10 magnitudes fainter than the target adds contamination errors at the 0.5 m s^{-1} level, as illustrated by the white dashed diagonal line across the bottom left plot in Figure 4. This is consistent with overall operations at the 1 m s^{-1} level of precision. However, it is inadequate for instruments hoping to detect

terrestrial-mass planets with actual RV semiamplitudes of 10 cm s^{-1} , because solar contamination has the potential to overwhelm that limit across most of our parameter space. Without correction, a magnitude differential of >12 seems necessary to limit the isolated error from solar contamination to $<10 \text{ cm s}^{-1}$. This becomes problematic when we consider that bright time is typically allotted for exoplanet RV science. Achieving sensitivity at the 10 cm s^{-1} level on-sky poses stringent additional requirements on observing conditions that must be factored in during time allotment on these next-generation instruments (Section 6).

3. Correcting for Solar Contamination

Solar contamination correction will need to be an essential algorithmic part of next-generation RV pipelines, optimizing the use of the sky fiber and other ancillary data across the observing conditions of the telescope site. Here we explore methods to correct for this effect, the level of correction possible, and inherent limitations.




Figure 5. Overview of the direct sky fiber cross-correlation CCF subtraction technique, as described in Section 3.1. The CCF of the sky fiber is computed by cross-correlation with a numerical mask (top) and subtracted from the CCF of the science fiber spectrum (middle). The final target RV is derived from the resultant corrected target CCF (bottom). This technique has the advantage of simplicity, although it does inherently rely on the sky and science fiber PSFs being similar. Note that the spectrum and CCF continuum levels are matched for illustrative purposes only; in reality, the sky fiber has a much lower flux level.

3.1. Direct Sky Fiber CCF Subtraction

In the field of precision RV measurements, sky correction has been previously attempted by directly subtracting the CCF of the sky fiber spectrum from the CCF of the science fiber spectrum. An overview of this technique is shown in Figure 5. Direct CCF subtraction is a simple but powerful technique that leverages the inclusion of a sky fiber with an identical fiber size and cross section (and hence very similar line spread function) as the science fiber. With a real instrument, the two fibers can point at slightly different parts of the sky and have different throughputs and aberrations, necessitating a relative scaling between the CCFs before subtraction. This method has been successfully implemented for the SOPHIE spectrograph, although at precisions of 20–100 m s⁻¹ it was mainly intended as a coarse correction of strong moonlight contamination to salvage certain observations (Barge et al. 2008; Hébrard et al. 2008; Pollacco et al. 2008; Bonomo et al. 2010; Santerne et al. 2011a, 2011b). Bonomo et al. (2010) also note that peak pulling from moonlight contamination is immediately evident in the CCF bisectors, a method we often use to detect spectral contamination (Wright et al. 2013) or stellar activity issues (Robertson et al. 2014) in RV measurements.

We test this correction method by subtracting the sky fiber CCF from the science fiber CCF for our simulated scenarios in Figure 6 (left). This reduces the contamination effect considerably, bringing the worst RV errors down from ~ 100 m s⁻¹ to ~ 10 m s⁻¹. This residual error is due to both a shift in the mean velocity measured after CCF subtraction (residual peak pulling from sky) and an increase in the standard deviation of the measured noise distribution (due to noise in the subtracted sky fiber CCF). In fact, as the sky gets darker, the sky fiber CCF becomes dominated by read noise and visibly indistinguishable from noise and does not actually apply any correction. That effect, however, is obscured by the fact that contamination errors are already low when the sky is dark. Since these small shifts are

difficult to measure, even on simulated spectra, irregularities persist in the contours even after running 1000 realizations for each combination of star and sky.

For the brightest NEID targets ($V_{\text{mag}} < 8.5$), this simple correction brings the errors down to < 1 m s⁻¹ for the full range of expected sky brightnesses at Kitt Peak. For fainter stars, this residual error can grow to ~ 10 m s⁻¹. Even though the residual error from solar contamination is lower than the noise floor across these observing conditions (bottom left of Figure 4), it is not negligible in the context of 10 m s⁻¹ long-term on-sky precision and must be accounted for in a comprehensive error budget.

3.2. Model Sky Subtraction

A different method of contamination removal involves the subtraction of a noiseless (or very high S/N) model sky spectrum from the science fiber spectrum, based on the flux recorded in the sky fiber (see Figure 7). The continuum level of the sky fiber CCF is directly related to the absolute sky brightness:

$$\text{Sky Brightness} = 2.5 \times \log\left(\frac{\text{constant}}{\text{CCF continuum}}\right). \quad (1)$$

In this method, we fit the sky fiber CCF with a Gaussian to measure both the continuum level and RV position of the scattered sunlight relative to the target star. Since the sky spectrum is approximately at the barycentric velocity, the CCF center can be fixed or heavily constrained for the peak fitting process. The estimated sky brightness is used to scale a noiseless model sky spectrum at the correct flux level. In this idealized scenario, we use the same synthetic spectrum that was used to generate the injected sky spectrum. In reality, this model should be drawn from a library of high-S/N observations of twilight sky at that particular site, preferably observed through the science fiber itself to facilitate better line-spread function (LSF) matching.

The model sky spectrum is subtracted from the science fiber spectrum, and a final CCF of the corrected target spectrum is generated. The results of this correction are shown in Figure 6 (right). This technique also reduces the solar contamination effect considerably, bringing the worst RV errors down from ~ 100 m s⁻¹ to ~ 1 m s⁻¹. It does not, however, work as well as simple subtraction of the sky fiber CCF because the extra steps of measurement, translation, and interpolation of spectra can introduce small errors.

This technique could be particularly useful for configurations where the sky and science fibers have different geometries or spectral resolutions, where direct CCF subtraction is not appropriate. Using the sky spectrum CCF to scale the template solar spectrum, rather than the recorded sky spectrum itself, may be advantageous for low-sky background levels where the continuum is difficult to identify. However, there is an additional burden in ensuring that the model spectrum matches the true sky spectrum based on site, airmass, moon phase, and other observing variables.

3.3. Limitations of the Sky Fiber

Both of the solar contamination correction techniques described above rely on the use of a simultaneous sky fiber. The CCF subtraction method further relies on the similarity of the science and sky fibers (or the ability to transform the LSF




Figure 6. Residual radial-velocity errors from solar contamination after applying two different correction methods. Details of the simulations are the same as in Figure 4. Left: subtracting the sky fiber CCF directly from the science fiber CCF, for a range of stellar V -band magnitudes and sky brightness levels. Right: subtracting a model sky spectrum (constructed based on the flux in the simulated sky fiber) from the science fiber spectrum. These results should be compared to the uncorrected solar contamination in Figure 4 (bottom right) and do not include the photon or read noise contributions.




Figure 7. Overview of the model sky subtraction technique, as described in Section 3.2. The CCF of the sky fiber is computed using a numerical mask and used to scale a synthetic or high-S/N solar spectrum to the estimated sky brightness (top). The scaled spectrum is then subtracted from the science spectrum, leaving only the target spectrum (middle). The final target radial velocity is then computed from the target CCF (bottom). Note that the spectrum and CCF continuum levels are matched for illustrative purposes only; in reality, the sky fiber has a much lower flux level.

and chromatic variability between the two). However, there may be times when a sky fiber is not available (in certain instruments one has to choose between calibration light or sky illumination in the simultaneous fiber) or has contamination issues itself.

There are two main locations that introduce contamination into the sky fiber. The first occurs at the port adapter fiber head

(Figure 2) and is caused by the wings of the stellar PSF extending into the area covered by the sky fiber ($22''$ separation). For a $V_{\text{mag}} = 12$ star and a sky brightness of $17 \text{ mag arcsec}^{-2}$, the sky fiber receives 0.001% extra contaminating flux; as the sky dims to $21.45 \text{ mag arcsec}^{-2}$ (nominal for WIYN dark time at zenith), this increases to 0.08%. Irrespective of instrument, some contamination at the port fiber




Figure 8. Overview of the simultaneous broadband correction technique, as described in Section 3.4. An image of the sky in one or more photometric bands is used to scale a solar template spectrum as a proxy for the background contamination. The scaled solar spectrum is then subtracted from the science spectrum, leaving only the target spectrum (middle). The final target radial velocity is then computed from the target CCF (bottom). In the case of NEID, a coherent fiber imaging bundle is used to sample the background sky in the vicinity of the science target. Note that the spectrum and CCF continuum levels are matched for illustrative purposes only; in reality, the sky fiber has a much lower flux level.

head is unavoidable given the desire to sample the same (or adjacent) patch of sky. This effect is naturally mitigated in better seeing, but the fact that it varies with observing conditions adds further complexity to modeling efforts.

The second, and more serious, source of concern is cross-contamination between fiber traces on the detector, caused by low-level wings of the cross-dispersion profile. Leaking of starlight into the sky fiber of the same echelle order (intraorder contamination), expected to be between 10^{-3} and 10^{-5} levels, is not hugely problematic in the CCF correction method, because it essentially results in a slightly depleted primary CCF peak. However, attempts to gauge sky brightness from the sky fiber are strongly affected as a function of stellar brightness. On the red end of the NEID bandpass, interorder contamination becomes much more significant as the spatial separation between the sky fiber and the neighboring calibration order diminishes to ~ 7 pixels (compared to ~ 80 pixels at the blue end). Attempted measurements of the sky fiber spectrum get quickly overwhelmed by light from the calibration source (laser frequency comb, etalon, or emission lamp). To fully utilize the sky fiber for NEID, calibration light will need to be filtered in the redder orders, although we will also test modeling and removal of cross-contamination on real data.

3.4. Simultaneous Broadband Correction

Given the low flux levels and concerns of contamination in the sky fiber, simultaneous broadband measurements can provide important redundancy and might become the only viable option for the darkest skies. In this method, the level of scattered solar irradiance is estimated by sampling the

background sky via broadband imaging, preferably in reasonably close proximity to the target object. This yields an independent estimate of the solar contamination and does not rely on the potentially noisy high-resolution sky spectrum. An overview of the technique is shown in Figure 8.

CFBs, with several thousand fiber cores, are flexible image guides that enable large fields of view relative to typical spectrograph fibers and are often used in astronomical observatories for target acquisition, guiding, and maintaining focus (Ramsey et al. 2003; Newman et al. 2004; Smee et al. 2013; Yan et al. 2016). For NEID, three CFBs will surround the science fiber in the port adapter fiber head. This design is itself adapted from the Habitable-Zone Planet Finder (Metcalf et al. 2019) fiber head design (Kanodia et al. 2018), which in turn is modeled on a design proposed by P. McQueen (private communication) for the Hobby–Eberly telescope’s upgraded high-resolution spectrometer.

Figure 2 shows the relative placement of the CFBs with respect to the science and sky fibers in NEID, as well as a magnified image of a single CFB. The CFBs ($6''/3$ on sky) will be used to regularly calibrate the port tip/tilt module and triangulate the precise center of the science fiber in the focal plane. This triangulation system provides a robust measurement of the location of the target star relative to the science fiber. All CFBs will be imaged onto a single commercial FLI MicroLine CCD and continue to collect data throughout the night. The CFBs will thus be gathering large amounts of sky light simultaneously with science exposures. The presence of three bundles creates additional redundancy for bright stars drifting into their fields of view. While contamination from stellar light




Figure 9. Predicting the relationship between the sky fiber spectrum and broadband images from the CFB. Left: mean photons per resolution element in the NEID sky fiber spectrum, for a 20 minute exposure, as a function of sky brightness. Middle: number of photons per pixel in a CFB V-band image of the sky for the same exposure time. For the sake of direct comparison, we are not applying the saturation limit of the detector. Right: relation between the counts in the broadband CFB image and the flux in the spectrum, showing a clean linear scaling that can be easily extended to fainter skies.




Figure 10. On-sky proof of concept that a relationship can be established between the sky fiber spectrum and broadband images from the CFB, as predicted in Figure 9. Data were taken as the Sun set to collect exposures at different sky brightnesses. Left: mean photons per extracted pixel in the echelette sky spectrum; all exposures have been scaled to 1 minute. Middle: number of photons per pixel in a CFB V-band image of the sky; all exposures are scaled to 1 second. Right: relation between the flux in the spectrum and the broadband CFB image, proving a clean linear scaling that can be extended to fainter skies. Note that the CFB counts are binned to 1 minute for direct comparison.

at the port adapter fiber head can also affect the CFBs (15'' separation from science fiber), they are immune to concerns of cross-contamination on the detector, making them an important complement to the sky fiber.

The CFBs provide a new and potent avenue for the correction of solar contamination using simultaneous broadband sky imaging. By analyzing CFB images in tandem with the science fiber observing the relatively bright sky, a scaling relation can be set up between the CCF continuum and the absolute sky brightness in the V-band. Figure 9 shows the simulated relative efficiency in sky flux collection between the NEID fiber and each CFB, and the predicted linear scaling relation between the two. This scaling can be extended to darker skies, allowing the option to circumvent the sky fiber when it becomes less reliable. The sky brightness level measured from the CFB can consequently be used to scale a noiseless model sky spectrum and subtract it from the science spectrum, as described above.

One caveat to note here is that the CFB image is not dispersed and hence cannot communicate chromatic variations of the sky spectral energy distribution (SED) as the weather changes during the observation. While the sky fiber technically contains this information, in both previous correction methods we ultimately measure the CCF, which collapses wavelength in favor of S/N, thereby forsaking this knowledge as well. Chromatic yet efficient sky measurements would require, for example, different filters on the three CFBs, which are not feasible with the current NEID port design. The CFBs also make it difficult to distinguish between fluctuations in scattered sunlight versus airglow or city lights; for NEID we try to limit these issues with custom filters (Section 4).

3.4.1. On-sky Demonstration

As proof of concept, we conducted an on-sky experiment at Penn State to show that the broadband image from the CFB can, in fact, be used as a reliable proxy for the sky fiber (Figure 10). A commercial 24 inch PlaneWave corrected Dall-Kirkham (CDK)




Figure 11. On-sky instrumental setup for scattered sunlight measurement. A small fraction of light collected by the 24 inch PlaneWave is diverted to a secondary image plane formed at the focus of the manual off-axis guider assembly. A CFB and $300\ \mu\text{m}$ fiber are placed at the focus of the MOAG. The CFB output is imaged onto an FLI CCD, while the $300\ \mu\text{m}$ fiber is inserted into a compact high-resolution spectrometer.

telescope was used to feed a dedicated optical fiber that directly measured the scattered night sky spectrum. We recorded spectra from evening to after dusk, mimicking the change in sky brightness for a variety of observing conditions and lunar separations. Figure 11 shows our on-sky measurement scheme using the 24 inch telescope. A fraction of light from the telescope pupil is picked off using an off-axis mirror within the manual off-axis guider (MOAG) assembly. The MOAG produces an image of the night sky that simultaneously illuminates a $300\ \mu\text{m}$ optical fiber and a 1.4 mm coherent fiber imaging bundle (CFB) cable. The $300\ \mu\text{m}$ optical fiber is fed into a small commercial echelette spectrometer, which produces a cross-dispersed, moderate-resolution ($R \sim 10,000$) optical spectrum (400–600 nm) of the night sky. The output of the CFB is imaged onto a 1024×1024 FLI CCD using a simple lens relay system that includes a Johnson V-band filter. Figure 12 shows the extracted spectra and a sample image of the CFB. The FLI CCD was experiencing a shutter lag during these exposures because of its orientation with respect to the telescope (a known issue with these cameras), leading to streaking in the images. This was corrected using darks and an interpolation algorithm to estimate the underlying flux from the shutter bleed.

The goal of this experiment was to explore the feasibility of using an integrated V-band flux as an independent method for estimating the solar contamination level. Ideally, the broadband measurement from the CFB can be used to reliably scale a template solar spectrum, which would then yield an estimate of the absolute solar contamination level. To test the validity of this technique, we compare the measured spectrum from the optical

fiber, under a variety of different pointings and conditions, to the scaled solar spectrum from the broadband CFB image (Figure 10). This follows the behavior predicted in Figure 9, with time as a surrogate for sky brightness. Flux scaling of the CFB image thus has the advantage of providing an independent verification of contamination levels, as well as potentially yielding a significantly higher S/N measurement of the solar contamination than the recorded high-resolution spectrum from the single optical fiber.

4. Building Correction Methods into NEID Design

NEID includes a dedicated CFB imaging camera, primarily used to aid in optical registration of the telescope focal plane with the NEID science and sky fiber (Logsdon et al. 2018; Schwab et al. 2018). The CFBs are mechanically mounted onto the same fiber connector as the NEID science and sky fibers and, when not in use as target registration fiducials, will provide independent samples of the background sky.

The outputs of the CFBs are imaged through a filter wheel onto a commercial FLI CCD package using a pair of telecentric imaging lenses. In addition to some standard neutral-density filters for the target registration function of the CFBs, we include high-performance off the shelf (Semrock) spectral filters for monitoring sky brightness. Instead of limiting the sky measurement to a standard V-band filter (which is included for target centering via CFB triangulation), we have two tailored filters centered at 488 nm and 510 nm to ensure bright airglow lines are avoided in the final recorded images (Figure 13). We are particularly careful to avoid the strong OI feature at $5577\ \text{\AA}$ and the Na I D lines and




Figure 12. Simultaneous spectra and images taken of the sky as the Sun set. Left: moderate-resolution ($R \sim 10,000$) optical spectrum of the sky taken with a small commercial echelette spectrometer calibrated with a Hg lamp. Right: a coherent fiber bundle (CFB) imaged on an FLI CCD.




Figure 13. Spectrum of the night sky at zenith above Kitt Peak from Massey & Foltz (2000), overlaid with the NEID CFB filter curves and a typical V-band filter. Emission features from both natural and human-made sources are highlighted. High-pressure sodium (HPS) lamps cause the broad feature extending from ~ 5500 to 6500 \AA . The two filters from Senrock, centered on 488 nm and 510 nm, were carefully selected to avoid bright (and variable) emission lines that could confound measurement of overall sky brightness, as well as good out-of-band suppression of the light. Filter transmission begins after the HgI 4358 feature and cuts off before the strong OI feature at 5577 Å and the Na D lines.

ensure good suppression outside the desired filter window. Avoiding highly variable, bright airglow lines is essential for ensuring the CFBs are tracing scattered solar contamination, and not variability due to other atmospheric effects.

In addition to the CFBs, NEID also includes a chromatically dispersed exposure meter to aid in performing accurate barycentric corrections. While the exposure meter does simultaneously record low-resolution spectra from all three NEID fibers, it only contains a very small fraction ($<1\%$) of the light in the main spectrograph, leading to extremely small levels of flux in the sky fiber trace. In combination with the fact that it is read out very fast to carefully track the flux accumulation during an exposure, the exposure meter is not a viable source of additional sky brightness monitoring for NEID.

5. Discussion

While our calculations show that RV errors due to scattered solar contamination may be mitigated using several different techniques, there are a certain site-based and instrument-specific

effects that we exclude for the sake of generalization. As listed below, these must all be considered when applying solar contamination correction to real observations and can add further complexity to the process:

1. Any sky fiber will be spatially separated from the science fiber in the spectrometer focal plane, leading to different aberrations and a fundamentally different PSF compared to the science fiber. The degree to which this will affect different correction techniques will vary based on the spectrometer design, though we aim to probe this with NEID in the near future.
2. We have assumed the science and sky fibers have identical (or perfectly characterized) throughputs including focal ratio degradation. This will also be system-specific, since fibers naturally have different transmission properties depending on manufacturing.
3. The sky is assumed to be uniform, both spatially and spectrally, between the science and sky fibers, which may not be the case depending on angular separation on-sky.

4. The SED of the sky contamination is assumed to be identical to the Sun, though in reality atmospheric conditions will modulate the solar spectrum chromatically to some degree (e.g., Brine & Iqbal 1983).
5. We assume that the sky fiber does not suffer significant contamination from the science or calibration fiber. While this is instrument dependent, it is unlikely to be strictly true for any instrument. The wings of stellar PSF overlapping with the sky fiber at the telescope fiber injection unit, scattering from the cross-disperser, scattering from optics like the primary echelle grating (e.g., the Rowland Ghost), and intra- and interorder cross-contamination between fibers on the detector will all contribute to stellar or calibration light contamination in the sky fiber.

6. Conclusion

We present an analytical study estimating the Doppler RV error due to scattered sunlight contamination in RV observations using next-generation spectrometers. This subtle effect, while low in absolute amplitude, can begin to inhibit high-precision Doppler measurements below the m s^{-1} level if not explicitly corrected. Brighter stars are naturally better inured against typical levels of scattered solar contamination than fainter stars, which show orders of magnitude larger RV errors.

To begin with, we show that this effect *must* be acknowledged when scheduling observations and that one cannot observe in all conditions of sky brightness, moon separation, and moon phase without considering the solar contamination error. If overlooked, this error source can easily grow to several m s^{-1} in magnitude and dominate the meager error budgets of next-generation instruments. We also demonstrate that the presence of a sky fiber is ideal, particularly if it has identical spectral properties (resolution, PSF, throughput) to the science fiber. If no sky fiber is available or the sky is too faint to reliably measure from the sky fiber, simultaneous broadband sampling of the sky shows great promise. Lastly, we emphasize that some level of sky correction is going to be an absolute necessity in the limit of faint targets (e.g., for follow-up of TESS targets with $V_{\text{mag}} > 12$) or highest desired precision (potential small rocky planets around bright stars).

We present three prospective correction methods: (1) subtracting the CCF of a dedicated sky fiber with the same spectral properties (resolution, throughput) as the target fiber, (2) using the sky fiber CCF to subtract a scaled model solar spectra, and (3) using simultaneous broadband measurements of the background sky to scale a template solar spectrum, using CFBs in the case of NEID. While simple CCF subtraction works best in the case of the sky fiber having the same spectral properties as the science fiber, subtracting model spectra presents a comparable alternative. The latter technique has the added benefit of not requiring the same PSF properties for both the science and sky fibers.

However, our presentation of the sky fiber usage here is in many ways idealized, since we do not account for complex effects like cross-contamination between fibers and other sources of temporal variability between the sky and science fiber. As an alternative to using the sky fiber, we present a third option of using a broadband source like the NEID coherent fiber bundles, with superior collecting power, to gain a coarse sky brightness measurement instead. The CFBs in NEID

are installed primarily for target acquisition but naturally observe the sky during all science observations. Recognizing the potential for a high-S/N proxy for the sky fiber early on, we have built toward this new functionality and anticipate demonstrating the required suppression of solar contamination with NEID.

Minimizing the impact of solar contamination to very low levels requires a combination of strategies. Careful planning of observations takes into account moon phase, RV and barycentric correction of the star, hardware solutions like sky fibers and CFBs to accurately measure the level of sky contamination, and algorithms to perform accurate corrections. Our intent in this paper is to present and model the deleterious impact of solar contamination and show that it can, at least in principle, be corrected to well below 10 cm s^{-1} using a combination of strategies. Future experiments with the publicly available NEID RV platform on the 3.5 m WIYN telescope will let us truly probe the effectiveness of these strategies in the face of real on-sky and instrument systematics, as well as explore the path to mitigating these issues further as we gather longer baselines of data at extreme precision.

On-sky tests were conducted with the Planewave CDK 24" Telescope operated by the Penn State Department of Astronomy & Astrophysics at Davey Lab Observatory. NEID is funded by NASA through JPL by contract 1547612. 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. This work was performed by S.P.H. (in part) under contract with the Jet Propulsion Laboratory (JPL) funded by NASA through the Sagan Fellowship Program executed by the NASA Exoplanet Science Institute. This work was supported by NASA Headquarters under the NASA Earth and Space Science Fellowship Program through grant NNX16AO28H. This research has made use of NASA's Astrophysics Data System Bibliographic Services.

Facility: DAVEY:0.6m.

ORCID iDs

Arpita Roy  <https://orcid.org/0000-0001-8127-5775>
 Samuel Halverson  <https://orcid.org/0000-0003-1312-9391>
 Suvrath Mahadevan  <https://orcid.org/0000-0001-9596-7983>
 Gudmundur Stefansson  <https://orcid.org/0000-0001-7409-5688>
 Sarah E. Logsdon  <https://orcid.org/0000-0002-9632-9382>
 Chad F. Bender  <https://orcid.org/0000-0003-4384-7220>
 Cullen H. Blake  <https://orcid.org/0000-0002-6096-1749>
 Shubham Kanodia  <https://orcid.org/0000-0001-8401-4300>
 Joe P. Ninan  <https://orcid.org/0000-0001-8720-5612>
 Paul Robertson  <https://orcid.org/0000-0003-0149-9678>
 Christian Schwab  <https://orcid.org/0000-0002-0091-7105>
 Ryan C. Terrien  <https://orcid.org/0000-0002-4788-8858>
 Jason T. Wright  <https://orcid.org/0000-0001-6160-5888>

References

- Anglada-Escudé, G., & Butler, R. P. 2012, *ApJS*, 200, 15
 Baranne, A., Queloz, D., Mayor, M., et al. 1996, *A&AS*, 119, 373
 Barge, P., Baglin, A., Auvergne, M., et al. 2008, *A&A*, 482, L17
 Bonomo, A. S., Santerne, A., Alonso, R., et al. 2010, *A&A*, 520, A65

- Brine, D., & Iqbal, M. 1983, *SoEn*, **30**, 447
- Chakraborty, A., Mahadevan, S., Roy, A., et al. 2014, *PASP*, **126**, 133
- Colina, L., & Bohlin, R. 1997, *AJ*, **113**, 1138
- Fischer, D. A., Anglada-Escude, G., Arriagada, P., et al. 2016, *PASP*, **128**, 066001
- Gibson, S. R., Howard, A. W., Roy, A., et al. 2018, *Proc. SPIE*, **10702**, 107025X
- Halverson, S., Terrien, R. C., Mahadevan, S., et al. 2016, *Proc. SPIE*, **9908**, 99086P
- Hébrard, G., Bouchy, F., Pont, F., et al. 2008, *A&A*, **488**, 763
- Husser, T.-O., Wende-von Berg, S., Dreizler, S., et al. 2013, *A&A*, **553**, A6
- Kanodia, S., Mahadevan, S., Ramsey, L. W., et al. 2018, *Proc. SPIE*, **10702**, 107026Q
- Kopparapu, R. K., Ramirez, R., Kasting, J. F., et al. 2013, *ApJ*, **765**, 131
- Krisciunas, K. 1997, *PASP*, **109**, 1181
- Logsdon, S. E., McElwain, M. W., Gong, Q., et al. 2018, *Proc. SPIE*, **10702**, 1070267
- Massey, P., & Foltz, C. B. 2000, *PASP*, **112**, 566
- Metcalfe, A. J., Anderson, T., Bender, C. F., et al. 2019, *Optica*, **6**, 233
- Neugent, K. F., & Massey, P. 2010, *PASP*, **122**, 1246
- Newman, P. R., Long, D. C., Snedden, S. A., et al. 2004, *Proc. SPIE*, **5492**, 533
- Pepe, F., Mayor, M., Galland, F., et al. 2002, *A&A*, **388**, 632
- Pepe, F. A., & Lovis, C. 2008, *PhST*, **130**, 014007
- Podgorski, W., Bean, J., Bergner, H., et al. 2014, *Proc. SPIE*, **9147**, 91478W
- Pollacco, D., Skillen, I., Collier Cameron, A., et al. 2008, *MNRAS*, **385**, 1576
- Racine, R. 1996, *PASP*, **108**, 699
- Ramsey, L. W., Engel, L. G., Sessions, N., et al. 2003, *Proc. SPIE*, **4841**, 1036
- Robertson, P., Mahadevan, S., Endl, M., & Roy, A. 2014, *Sci*, **345**, 440
- Roy, A., Chakraborty, A., Mahadevan, S., et al. 2016, *Proc. SPIE*, **9908**, 99086R
- Santerne, A., Bonomo, A. S., Hébrard, G., et al. 2011a, *A&A*, **536**, A70
- Santerne, A., Díaz, R. F., Bouchy, F., et al. 2011b, *A&A*, **528**, A63
- Schwab, C., Liang, M., Gong, Q., et al. 2018, *Proc. SPIE*, **10702**, 1070271
- Schwab, C., Rakich, A., Gong, Q., et al. 2016, *Proc. SPIE*, **9908**, 99087H
- Seager, S. 2010, *Exoplanets* (Tucson, AZ: Univ. Arizona Press)
- Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *AJ*, **146**, 32
- Walker, A. 1991, *PASP*, **103**, 1033
- Wehbe, B., Cabral, A., Martins, J., et al. 2020, *MNRAS*, **481**, 3515
- Wright, J. T., & Robertson, P. 2017, *RNAAS*, **1**, 51
- Wright, J. T., Roy, A., Mahadevan, S., et al. 2013, *ApJ*, **770**, 119
- Yan, R., Bundy, K., Law, D. R., et al. 2016, *AJ*, **152**, 197