EXOPLANET TRANSIT SPECTROSCOPY USING WFC3: WASP-12 b, WASP-17 b, AND WASP-19 b

AVI M. MANDELL[1], KOREY HAYNES[1,2], EVAN SINUKOFF[3], NIKKU MADHUSUDHAN[4], ADAM BURROWS[5], AND DRAKE DEMING[6]

[1] Solar System Exploration Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA; Avi.Mandell@nasa.gov
[2] School of Physics, Astronomy, and Computational Sciences, George Mason University, Fairfax, VA 22030, USA
[3] Institute for Astronomy, University of Hawaii, Honolulu, HI 96822, USA
[4] Yale Center for Astronomy and Astrophysics, Yale University, New Haven, CT 06511, USA
[5] Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA
[6] Department of Astronomy, University of Maryland, College Park, MD 20742, USA

ABSTRACT

We report an analysis of transit spectroscopy of the extrasolar planets WASP-12 b, WASP-17 b, and WASP-19 b using the Wide Field Camera 3 (WFC3) on the *Hubble Space Telescope* (*HST*). We analyze the data for a single transit for each planet using a strategy similar, in certain aspects, to the techniques used by Berta et al., but we extend their methodology to allow us to correct for channel- or wavelength-dependent instrumental effects by utilizing the band-integrated time series and measurements of the drift of the spectrum on the detector over time. We achieve almost photon-limited results for individual spectral bins, but the uncertainties in the transit depth for the band-integrated data are exacerbated by the uneven sampling of the light curve imposed by the orbital phasing of *HST*'s observations. Our final transit spectra for all three objects are consistent with the presence of a broad absorption feature at 1.4 μm most likely due to water. However, the amplitude of the absorption is less than that expected based on previous observations with *Spitzer*, possibly due to hazes absorbing in the NIR or non-solar compositions. The degeneracy of models with different compositions and temperature structures combined with the low amplitude of any features in the data preclude our ability to place unambiguous constraints on the atmospheric composition without additional observations with WFC3 to improve the signal-to-noise ratio and/or a comprehensive multi-wavelength analysis.

Key words: planetary systems – techniques: photometric – techniques: spectroscopic

Online-only material: color figures

1. INTRODUCTION

Over the past decade there has been significant progress in characterizing exoplanets orbiting a wide variety of nearby stars, including the first detections of light emitted by an exoplanet (Charbonneau et al. 2005; Deming et al. 2005), the first spectrum of an exoplanet (Richardson et al. 2007; Grillmair et al. 2007; Swain et al. 2008), the first phase curve for an exoplanet (Knutson et al. 2007), the first detection of haze in an exoplanetary atmosphere (Pont et al. 2008), and tentative constraints claimed for the water, methane, carbon monoxide, and carbon dioxide abundances in several exoplanetary atmospheres (Grillmair et al. 2008; Swain et al. 2008, 2009b, 2009a; Madhusudhan & Seager 2009; Madhusudhan et al. 2011). Almost 100 transiting exoplanets with $V_{star} < 12$ have been discovered to date, many with multi-band photometry from both space- and ground-based observatories. We are firmly in the era of exoplanet characterization, and yet the sparse data available for each planet has resulted in more questions than answers.

The Wide Field Camera 3 (WFC3) on the *Hubble Space Telescope* (*HST*) provides the potential for spectroscopic characterization of molecular features in exoplanet atmospheres, a capability that has not existed in space since the demise of NICMOS on *HST* and the Infrared Spectrograph on *Spitzer*. WFC3 is an optical/NIR camera capable of slitless grism spectroscopy, with wavelength coverage in the IR spanning between 0.8 and 1.7 μm. Studies of exoplanets have focused on using the G141 grism, the long-wavelength dispersion element on the infrared channel that covers the wavelength range 1.1–1.7 μm at a maximum resolving power of 130 at 1.4 μm (Dressel 2012). This region spans both the major bands of water between 1.3 and

1.5 μm as well as another water band at 1.15 μm, and bands of a few other molecular species. Observations measuring flux within NIR water bands are impossible from the ground due to the extinction and variability caused by water vapor in Earth's atmosphere; WFC3 therefore represents the only current platform for measuring absorption and/or emission from water in exoplanet atmospheres.

In this paper, we present WFC3 observations of three transiting "hot Jupiter" exoplanets—WASP-12 b, WASP-17 b, and WASP-19 b—during transit of the host star. Two of these data sets, for WASP-17 b and WASP-19 b, were observed as part of a large *HST* program to examine single transits and occultations from 16 hot Jupiters (P.I.: D. Deming), while the data for the transit of WASP-12 b were taken as part of a single-object campaign (P.I.: M. Swain) and first analyzed in Swain et al. (2013). All three planets orbit extremely close to their parent star and have large atmospheric scale heights, making them excellent targets for transmission spectroscopy. WASP-12 b and WASP-17 b (as well as WASP-19 b to a lesser extent) belong to a class of "bloated" or "inflated" planets, which have significantly larger radii than would be predicted from traditional evolutionary models (Burrows et al. 2000; Guillot & Showman 2002). WASP-17 b is also in a retrograde orbit compared to the rotation of its host star (Anderson et al. 2010; Bayliss et al. 2010; Triaud et al. 2010), while WASP-12 b and WASP-19 b appear to be in prograde orbits (Albrecht et al. 2012; Hellier et al. 2011). Retrograde orbits have commonly been interpreted as evidence that the planet was forced into a highly inclined and eccentric orbit through planet–planet scattering (Rasio & Ford 1996; Weidenschilling & Marzari 1996) or the Kozai mechanism (Fabrycky & Tremaine 2007), and was subsequently

re-circularized through dissipation of orbital energy by tides (Jackson et al. 2008). The extremely short orbit of WASP-19 b also argues for tidal decay after scattering (Hellier et al. 2011). In the tidal decay scenario the large radii of the planets could be due to internal dissipation of tidal energy during orbital circularization (Bodenheimer et al. 2001). However, based on recent models by Ibgui & Burrows (2009), Anderson et al. (2011) conclude that any transient tidal heating produced during circularization of the orbit of WASP-17 b would have dissipated by the time the planet reached its current orbit, making the planet's large radius unsustainable. Other theories for the misalignment of the stellar rotation and the planet's orbit do not require a previous eccentric orbit and tidal re-circularization (Rogers et al. 2012), and a number of other theories for the heating mechanisms required to produce large planetary radii have been proposed, including "kinetic heating" due to the dissipation of wind energy deep in the atmosphere (Guillot & Showman 2002) and Ohmic dissipation (Batygin & Stevenson 2010). Therefore, the dynamical origin of these extremely hot and inflated giant planets is still highly uncertain.

In principle, understanding the atmospheric composition of hot Jupiters can help constrain their formation and dynamical histories. Unfortunately, observational studies have produced conflicting results regarding the atmospheric compositions of several hot Jupiters, including WASP-12 b and WASP-19 b. Madhusudhan et al. (2011) first raised the possibility of a non-solar abundance in the atmosphere of WASP-12 b using occultation measurements in four *Spitzer* photometric bands (Campo et al. 2011) and three ground-based NIR photometric bands (Croll et al. 2011) to constrain the carbon-to-oxygen (C/O) ratio to super-solar values, possibly greater than unity. Similar *Spitzer* and ground-based measurements for WASP-19 b were consistent with both solar and super-solar C/O models (Anderson et al. 2013), raising the possibility of a population of carbon-rich hot Jupiters. However, Crossfield et al. (2012) recently re-analyzed the *Spitzer* data for WASP-12 b in light of the discovery of a faint candidate companion imaged by Bergfors et al. (2013), concluding that the dilution-corrected *Spitzer* and ground-based photometry can be fit by solar-metallicity models with almost isothermal temperature structures.

While transmission spectroscopy only weakly constrains the overall temperature structure of a transiting exoplanet, it can place strong constraints on the presence of molecular features in absorption through the limb of the planet, thereby constraining the atmospheric composition. Models by Madhusudhan (2012) suggest that spectral features of H_2O and hydrocarbons (e.g., CH_4, HCN, and C_2H_2) will change drastically with different C/O values, and the WFC3 bandpass covers several of these features. In this paper we present our data reduction and analysis of the three transits, including our analysis of contamination from nearby sources and our strategy to compensate for the significant instrumental systematics in much of the WFC3 data, and conclude with preliminary constraints on the atmospheric composition and structure of the three planets.

2. OBSERVATIONS

The observations of WASP-17 and WASP-19 analyzed here were conducted between June and July of 2011, while the observations of WASP-12 were obtained in April of 2011. Observation dates and exposure information are listed in Table 1. The observations were taken with the G141 grism on WFC3's infrared channel, providing slitless spectra covering the wave-

Table 1
Observation Parameters

	WASP-12	WASP-17	WASP-19
Date of observation	2011 Apr 12	2011 Jul 8	2011 Jul 1
Integration time	7.624	12.795	21.657
Sub-array mode	256	512	128
CALWF3 version	2.7	2.3	2.3
NSamp	3	16	5
Timing sequence	SPARS10	RAPID	SPARS10
Peak pixel value[a]	38000	64000	73000

Note. [a] The number of electrons recorded at the peak of the spectral distribution in a single exposure.

length range 1.1 –1.7 μm at a maximum resolving power of 130 at 1.4 μm (Dressel 2012). Dithering was avoided to minimize variations in pixel-to-pixel sensitivity. The "spatial scanning" mode suggested as a strategy to increase efficiency and decrease persistence for bright objects (McCullough & MacKenty 2012) was not used since it had not been developed at the time of observation. Each target was allocated 4–5 *HST* orbits, each lasting 90 minutes followed by 45 minute gaps due to Earth occultations of the telescope. This was sufficient to cover a single transit while including some out-of-transit data as well.

The IR channel of the WFC3 instrument uses a 1024 × 1024 pixel detector array, but smaller sub-arrays can be downloaded to decrease the readout time and increase the exposure cadence. Additionally, there are two possible sampling sequences: RAPID sampling, which reads as quickly as possible (limited only by the readout time per sub-array) in order to maximize sampling for short exposures of bright targets, and SPARS sampling, which takes two quick reads and then spaces reads linearly, to allow "sampling up the ramp," or SUTR. RAPID sampling naturally has shorter readout times for each sub-array size but imposes a maximum integration time, while the SPARS10 sampling sequence has a minimum exposure time of ~7 s but no maximum.

Observations of WASP-17 were taken using the 512 × 512 sub-array with 16 non-destructive reads per exposure and sampled using the RAPID sampling sequence. This resulted in a total integration time of 12.795 s per exposure and 27 exposures per orbit, with a total of 131 exposures taken over five *HST* orbits. Observations of WASP-19 were taken using the 128 × 128 sub-array mode with five non-destructive reads per exposure, sampled with the SPARS10 sequence. This resulted in an integration time of 21.657 s and 70 exposures per orbit, with a total of 274 exposures taken over four orbits. The WASP-12 data utilized the 256 × 256 sub-array mode with three non-destructive reads per exposure, leading to an integration time of 7.624 s and 99 exposures per orbit, with 484 exposures taken over five *HST* orbits. We discuss the implications of each sub-array size with respect to systematic trends in Section 4.1.

3. DATA REDUCTION

3.1. Image Files: .flt *versus* .ima

The WFC3 calwf3 calibration pipeline processes the raw detector output into two calibrated files per exposure: a file comprising the individual, non-destructive reads (called the .ima file) and a single final image produced by determining the flux
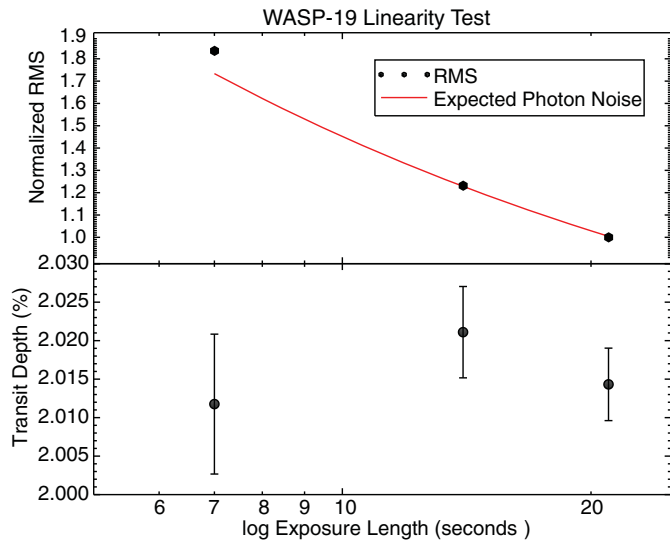
Figure 1. Top: the normalized rms compared to expected photon noise for a band-integrated light curve for WASP-19 created using different individual reads. Bottom: fitted transit depth for each read. The rms follows the photon-limited trend except for the first point, which most likely reflects read noise; the best-fit values are the same within uncertainties.

(A color version of this figure is available in the online journal.)

rate by fitting a line to the individual read-out values for each pixel (called the `.flt` file). The calibration steps implemented for the `.ima` files include reference pixel subtraction, zero-read, and dark current subtraction, and a nonlinearity correction; additional corrections applied using SUTR fitting for the `.flt` files include cosmic-ray and bad-pixel flagging and gain calibration. While it would seem that the `.flt` files would be the best choice for analysis, an analysis of the noise characteristics for each data type revealed that time series extracted from the `.flt` files have an rms that is on average 1.3× greater than time series created from the `.ima` files. It is unclear where this difference originates, though it is probably due to inaccurate cosmic-ray flagging for very bright sources (STScI WFC3 Support 2012, private communication). We therefore decided to determine our own flux values for each pixel directly from the `.ima` files and essentially re-create our own `.flt` files as a starting point for our analysis (this method was also advocated by Swain et al. (2013) for similar reasons).

Though the `.ima` files include a linearity correction, the exposures for some our objects approached or exceeded the established linearity limit for WFC3 and we therefore examined our data for signs of any remaining nonlinearity. The WFC3 detector generally remains linear up to 78 K e^- (WFC3 Handbook). However, Swain et al. (2013) suggest that known WFC3 issues with systematic increases in counts between buffer downloads (see Section 4.1) may be present when count levels exceed 40 K DN, or the equivalent of 100 K e^-. Our peak counts reach a maximum of 73 K e^- for WASP-19, with lower values for our other targets (see Table 1); we therefore chose WASP-19 to examine linearity. WASP-19 only has a total of four SUTR measurements, and in Figure 1 we show that the normalized rms of our band-integrated light curve follows the expected decrease for a photon-limited case. We also examined the linearity of each channel separately, in order to search for correlations with the final transit depth. Deviations from linearity were ~0.8% on average, but the channel-to-channel differences were only ~0.1% and would affect the transit depths for individual channels by only ~20 ppm, far below our uncertainty limits.

After binning up channels, this effect would be even less; we therefore did not use any additional linearity correction.

3.2. Spectral Extraction

The unique requirements of time-series photometry of bright sources necessitated the development of a custom-designed data reduction process for WFC3 exoplanet data. A data reduction package called aXe (Kümmel et al. 2009) exists for analyzing WFC3 data, but this software was designed with dithered observations in mind, and we used the package only for generating a wavelength solution and nominal extraction box sizes, since the package incorporates the most recent configuration files for the instrument. An object list was first generated by SExtractor (Bertin & Arnouts 1996), which uses the direct image to find the position of each source. aXe then calculates the trace and wavelength solution for each source, and produces FITS files with an extracted box from each grism image (with the extension `.stp`) and a one-dimensional spectrum (with the extension `.spc`) from which we extract the wavelength solution. For simplicity, we assumed that each pixel in a column has the same wavelength solution; measurements of the center of a Gaussian fit to the dispersion in the y direction showed that it changes by less than 0.02 pixels along the length of the spectra for all of our objects, so this assumption is valid. We also checked our wavelength solutions against the standard WFC3 sensitivity function to confirm accuracy for all sources.

We retrieved the coordinates for the extraction box from the headers of the `.stp` files, but we decided to expand the number of rows included in the extraction box from 15 pixels to 20 pixels to ensure that we included as much of the wings of the spatial point-spread function (PSF) as possible while avoiding any possible contamination from background sources. We also trimmed the extraction boxes to exclude regions of the spectrum with low signal-to-noise ratio (S/N), keeping the central 112 pixels of each spectrum.

3.3. Flat Field, Background Subtraction, Bad Pixel, and Cosmic-Ray Correction

The *calwf3* pipeline does not correct for pixel-to-pixel variations in grism images, but a flat-field cube is provided on the WFC3 Web site (Kuntschner et al. 2008). Each extension of the cube contains a coefficient, developed by ground tests, that can be fed into a polynomial function as follows, where $x = (\lambda - \lambda_{min})/(\lambda_{max} - \lambda_{min})$ and λ is the wavelength of pixel (i, j):

$$f(i, j, x) = a_0 + a1 * x + a_2 * x^2 + \ldots ai * x^n. \quad (1)$$

This polynomial gives the value of the flat field at each pixel in the extraction region, and we divided this flat field from our data. We also subtracted an average background flux from each spectral channel by using nearby uncontaminated regions of each image. These background regions cover the same wavelength space (extent in the x direction) as our science box, and are placed as far from the primary source as possible, leaving only a few pixels to guard against edge effects. We then averaged these background rows in the y direction, and subtracted this background spectra from each row of our science box. The average value of the background region was ~15–35 e^-, but for each source background, counts drop quickly at the beginning of each orbit and then continue to decrease slowly over the orbital duration (see Figure 2). The pattern is very similar in each
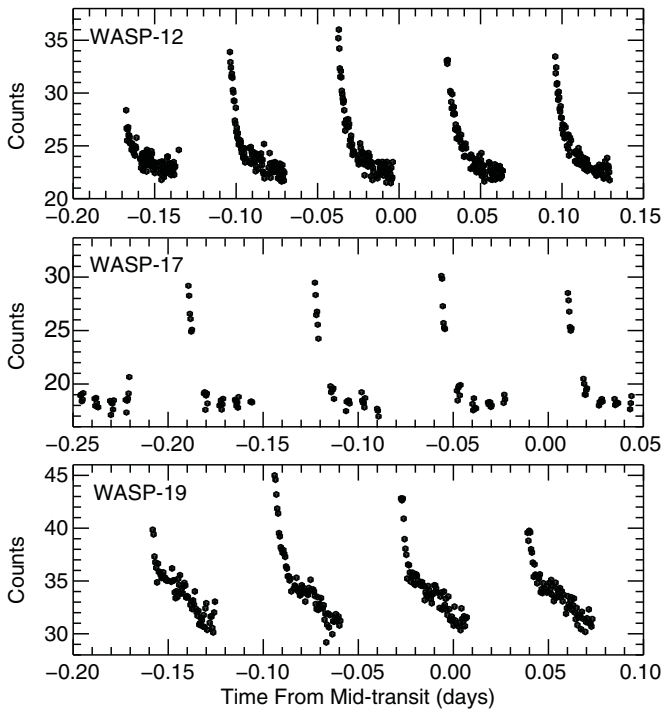
Figure 2. Background levels in counts for each target, given as a function of time from mid-transit. The drop in flux on a per-orbit basis is similar for each target, indicating that instrumental effects such as thermal variations during orbit are the likely cause.

channel, and is most likely due to thermal variations during the orbit.

To identify pixels that are either permanently bad or contaminated by cosmic rays, we employed several different bad-pixel identification strategies. First, to find individual pixels in individual images that were contaminated by cosmic rays or sensitivity variations, we created a three-dimensional image cube and examined each pixel in the two-dimensional images over time; any single-image pixels that were $>6\sigma$ higher than the median of their counterparts in time were flagged. We found 62 bad pixels for WASP-12, 30 for WASP-17, and 120 for WASP-19. We corrected most of these pixels through spatial interpolation in their individual frames. However, the linear interpolation that we used to correct bad pixels would clearly not be effective within the region covered by the stellar PSF due to the rapid change in flux across pixels in the spatial direction. Bad pixels within the PSF would also clearly have severe effects on the time series even if they were corrected, and we therefore left these pixels uncorrected.

We then summed over the spatial dimension of the corrected cube yielding a two-dimensional (wavelength, time) array and normalized this array in both the spectral and temporal dimensions, allowing us to remove the band-integrated transit signal and the stellar and instrumental spectral characteristics. This allowed us to identify both bad spectral channels in individual images as well as individual images and/or channels that showed increased noise or unusual characteristics. Through this analysis we found 20 individual bad data points for WASP-12, 6 for WASP-17, and 16 for WASP-19, which we corrected by linear interpolation in the spectral dimension. Additionally, we identified several spectral channels in each data set whose time series showed a significantly higher rms scatter compared with the rest of the channels; we removed two channels for WASP-12, four channels for WASP-17, and one channel for WASP-19 from further analysis as well.
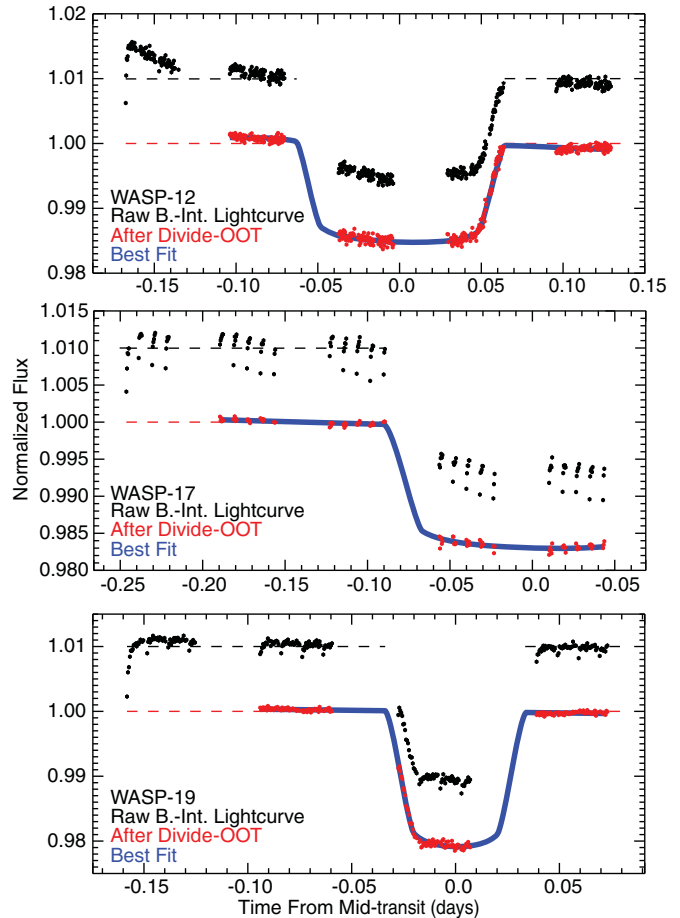

Figure 3. Combined-light time series for each source, before and after removing systematic trends. The presence of an intra-orbit pattern is easily identified for WASP-17 and WASP-19, repeating after every buffer read-out, but less obvious for WASP-12. After excluding the first orbit which is inconsistent with the others due to telescope settling, we removed the trends using the divide-oot method devised by Berta et al. (2012).

(A color version of this figure is available in the online journal.)

4. ANALYSIS AND RESULTS

4.1. Instrumental Systematics

Two out of our three data sets show strong systematic trends with time, which can be attributed to various instrumental effects, and have been seen in previous observations (Berta et al. 2012; Swain et al. 2013). The most obvious trend is the pattern of increasing counts after each data buffer download to the solid-state drive, possibly due to the use of charge-flush mode during the download (Swain et al. 2013). Depending on how quickly the count level stabilizes, this pattern can resemble a "ramp" (continually increasing until the next buffer download) or a "hook" (increasing for several exposures and then stabilizing). The effect may be associated with the well-known persistence effects inherent in HgCdTe detectors in general (Smith et al. 2008) and confirmed in WFC3 in particular (McCullough & Deustua 2008), but the relationship to the data buffer downloads suggests a connection to the data storage devices. Swain et al. (2013) performed an exhaustive analysis of the buffer-ramp effects in a number of different sources, and suggest that a smaller sub-array size, a fewer number of non-destructive reads, and a lower illumination level will decrease or eliminate the effect; for reference, we list the relevant attributes for each target in Table 1. The band-integrated light curves (Figure 3)
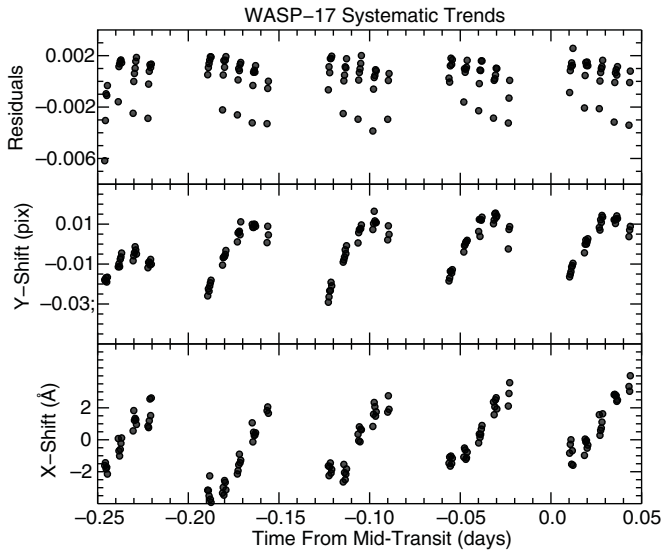
Figure 4. Top: the residuals of the combined-light fit for WASP-17, after subtracting our best-fit model. Middle: the shift in the position of the spatial profile of WASP-17 over time, in pixels. The vertical shift was calculated by fitting a Gaussian to sum of the spectral box in the spectral direction. Bottom: the shift in the position of the spectral profile over time, in pixels. The horizontal shift was calculated by measuring the change in flux over the edges of the spectrum and deriving the required shift of the spectral sensitivity function (see Section 4.1).

for the three objects we analyze here follow this general relationship—WASP-12 (intermediate array size, 3 reads, low peak pixel flux) has no buffer-ramp effect, while WASP-17 (large array size, 16 reads, high peak pixel flux) has a very steep ramp-up with no apparent stabilization before the next buffer dump. WASP-19 (small array size, 5 reads, high peak pixel flux) displays a shape intermediate between the two (a "hook"-like shape). We do not attempt a more detailed analysis of the cause of the buffer-ramp effects, and we find that the divide-oot method developed by Berta et al. (2012) is sufficient to remove the effect almost completely in the band-integrated light curve, provided sufficient out-of-transit data is available. We also see a visit-long decrease in flux; this effect has been noted in previous WFC3 analyses and may be due to a slow dissipation of persistence charge, and we correct for it using a linear trend component in our transit model fit. As noted in previous work, the first orbit for each target showed substantially higher scatter than all other orbits, and we do not use this orbit in our band-integrated divide-oot analysis; however, for our wavelength-dependent analysis we use a relative-depth analysis (see Sections 4.3 and 4.4 for a detailed description), and with this fitting strategy we are able to incorporate the first noisier orbit.

For each image, we also calculated the shift in the vertical (i.e., spatial) and horizontal (i.e., spectral) directions referenced to the first exposure in the time series. This allowed us to correct for any modulation in channel flux due to undersampling of the spatial PSF and/or spectral features. Since the FWHM of the PSF is ∼3 pixels, any vertical shifts can have a significant effect on the illumination of individual rows, and a similar effect can occur due to features in the stellar spectrum or the WFC3 sensitivity function that are several pixels wide. However, the shifts we measure are only a fraction of a pixel (see Figure 4) and the motion of a pixel across the spatial PSF or a spectral feature will be extremely small, creating a change in flux that is essentially linear. We can therefore decorrelate this effect against a scaled measurement of the image motion in each direction.

We measured the vertical shift by first summing our extraction box in the wavelength direction to get a one-dimensional array of the flux absorbed by each row of the detector for each exposure and then fitting a Gaussian to those arrays to determine the change in the location of the peak of the flux distribution from the first exposure. A precise measurement of the horizontal shift (i.e., the spectral shift) across all exposures was more difficult to calculate, since the sensitivity function of the grism does not allow for an analytical fit. We first attempted to cross-correlate the spectra against each other, but the scatter in the resulting measurements was too high to be useful. We then decided to utilize the edges of the spectrum where the sensitivity function of the detector rises and falls rapidly, and a small change in pixel position will have a strong effect on the illumination of each pixel. We fit a line to the slope for the same pixels at the edge of the spectrum for each exposure, and used the intercept of this fit to determine the shift of each spectrum in relation to the first exposure. The values from the fit to both the short-wavelength and long-wavelength edge of each spectrum were averaged to decrease the effective uncertainty of the measurement. In Figure 4 the vertical and horizontal shifts, as well as the final, band-integrated residuals after subtracting a light curve model, are plotted for WASP-17 as an example. All of the variables change relatively coherently within an orbit, and then reset at the beginning of the next orbit.

4.2. Background Source Correction

We also examined each object for contamination from background sources. Due to the slitless design of WFC3, spectra from background sources can be shifted both spatially and spectrally compared with the science target. In particular, a nearby background source or companion was discovered for WASP-12 (Bergfors et al. 2013) and more recently confirmed to be a double star (Bechter et al. 2013); the close companions have been shown to significantly affect the mid-IR photometry of this source with *Spitzer* (Crossfield et al. 2012). After averaging all of the images for each source, we examined each combined image by eye for evidence of background contamination, and then used a vertical profile cut to further constrain the amplitude and location of any identified sources. For WASP-19 there were no additional sources, and for WASP-17 the single background source identified nearby was very dim and significantly shifted in the spatial direction from the science target and therefore exterior to our extraction box.

For WASP-12 we identified a relatively bright contamination source very close to the science target; the peak of the spectral profile of the secondary source is located only ∼4 pixels away from the peak of the primary stellar PSF in the spatial direction (see Figure 5). This object is most likely the source identified by Bergfors et al. (2013) and subsequently referred to as Bergfors-6 by Crossfield et al. (2012) and WASP-12 BC by Bechter et al. (2013). After correcting for a shift of the secondary source in the spectral direction, the separation between the two sources matches up well with the previous measurements. As stated above, Bechter et al. (2013) resolved the source into two stars, but in the direct image from *HST* they are unresolved—the difference in the FWHM of the primary PSF compared to the secondary PSF is only 0.25 pixels. We therefore refer to the combined contamination from the two stars in our data as WASP-12 BC. Swain et al. (2013) also identified this contamination, and fit the profile of the PSF in the spatial direction by using the PSF shape from separate observations of a reference star. This method has the benefit of providing
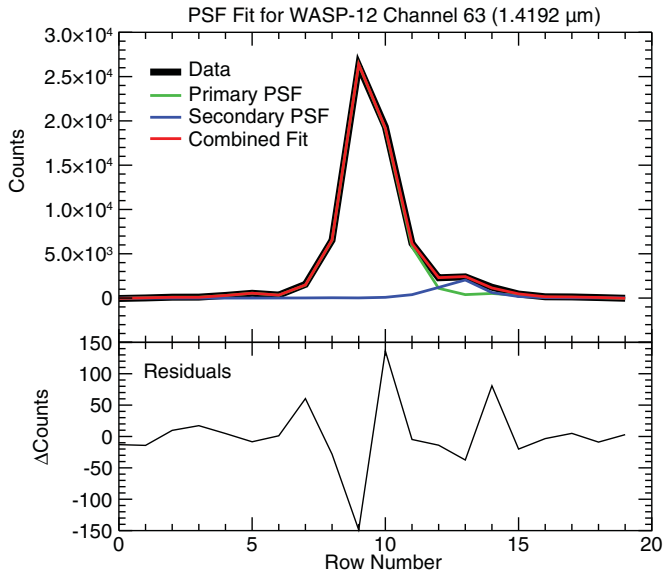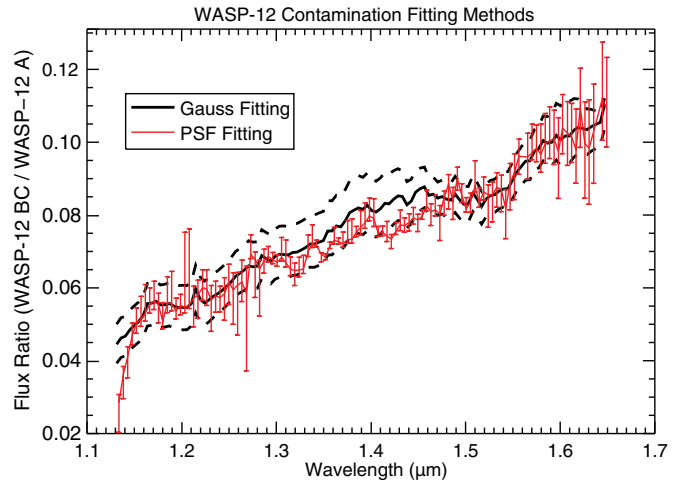
Figure 5. Top: data and the best-fit PSF model for a single channel for WASP-12, using the template-PSF method. The data are shown in black, the fit to the main peak is shown in green, and the fit to the contamination peak is shown in blue; the combined fit is shown in red. Bottom: remaining residuals after removing the model; the remaining flux under the region of contamination was used as the uncertainty in the contamination flux.

(A color version of this figure is available in the online journal.)

Figure 6. Flux ratio for WASP-12 BC compared to WASP-12 A for our two fitting methods. Gaussian fitting (black) subtracts one Gaussian centered on WASP-12's position, then fits another Gaussian to the residuals, centered on the contaminating source. Template PSF fitting (red) jointly scales two PSFs, using pre-determined columns from WASP-19 as a template. The uncertainties using the template PSF method are much smaller, even with the distortions of the secondary PSF due to the multiplicity of WASP-12 BC.

(A color version of this figure is available in the online journal.)

an empirical PSF shape that can be used for both the brighter primary star as well as the secondary star. This strategy is slightly complicated in this instance because of the multiplicity of the secondary source, but as stated above, the change in the width of the PSF is extremely small. The more difficult problem is that the angle of the spectrum on the detector is slightly offset from the horizontal pixel pitch. Therefore, the PSF changes shape with wavelength, and the primary and secondary PSFs are sampled differently.

Fortunately, our WASP-19 spectrum was also slightly angled on the detector, and since the flux levels remaining in the linear regime we were able to scale individual channels from our WASP-19 data as PSF "templates" for the WASP-12 channels (as suggested by Swain et al. 2013). The PSF of WASP-12 BC could also be fit in the same way, albeit with a different initial off-set for the starting template channel. We empirically determined the best-fit template channel off-set for both PSFs, and then performed a least-squares fit for the PSF amplitude of both stars at once. In Figure 5 we show an example of a fit to one of our WASP-12 channels; the remaining residuals in the region with the contaminating source will be impacted slightly by the distorted PSF of the double stars, so we summed them up to give uncertainties on the fit in the positive and negative directions. In addition to this PSF template strategy, we tested a straightforward sequential Gaussian fitting method, first fitting and subtracting the largest-amplitude signal (from the science target) and then fitting the additional contamination source. However, due to the under-sampling of the spatial PSFs and their overlap between the two sources, there was substantial uncertainty in the fundamental baseline of the individual PSF functions for each source, and considerable residual flux was left over after removing the contribution from both PSFs. In Figure 6 we plot our spectrum for WASP-12 BC derived from both methods. The results agree extremely well at short and long wavelengths except for an overall offset and some slight discrepancy between 1.35 and 1.45 μm. However, the

uncertainties are at least a factor of three smaller using the template-PSF method, even with the contributing error from the multiplicity of WASP-12 BC. We therefore adopted the results from the template-PSF fitting method, and corrected the data by subtracting the derived spectrum of the contaminating source from the one-dimensional spectrum at each time step.

For comparison, we calculated the expected ratio of the contaminating source to the primary star using stellar atmosphere models from Castelli & Kurucz (2004), assuming that the contaminating source is the combined light from WASP-12 B and WASP-12 C. Crossfield et al. (2012) determined a spectral type of M0V and an effective temperature between 3600 K and 3900 K for what they believed was a single star, depending on whether purely spectroscopic or a combination of spectroscopic and photometric data were used; for WASP-12, Hebb et al. (2009) determined an effective temperature of 6300^{200}_{100} K. Since Bechter et al. (2013) find that both companions have a similar spectral type and brightness, we can effectively treat them as one source. We assumed the same metallicity for all the stars, and used the direct image to derive a shift of 331 Å in the spectral direction for the contaminating source. We then scaled the ratio of two stellar models to match our results at 1.6 μm. In Figure 7 we plot our results from our PSF template method, with two analytic models spanning the range of effective temperatures for WASP-12 A and WASP-12 BC. A lower-temperature model for the combined flux from WASP-12 BC shows a significantly deeper water absorption feature from 1.4 to 1.6 μm compared with higher-temperature models, while a higher temperature for WASP-12 A makes a very small change in the overall slope. Our empirical fit to the data agrees very well with a model using a temperature of ~3900 K for WASP-12 BC, which matches well with the M0 spectral type derived by Bergfors et al. (2013) and Crossfield et al. (2012) but is inconsistent with the spectral type of M3V determined by Bechter et al. (2013) for both WASP-12 B and C.

A similar calculation of the contaminating flux was used by Stevenson et al. (2013); however, they assumed the lower effective temperature for WASP-12 BC from the spectroscopic
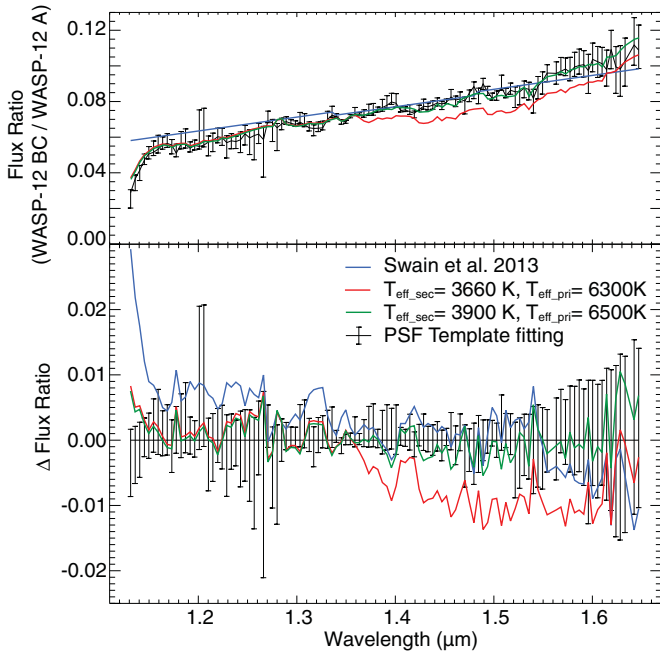
Figure 7. Top: flux ratio for the contaminating source (WASP-12 BC) from the template PSF fitting method (black), compared with analytical models for the flux ratio bracketing the range of values for the temperatures of WASP-12 and the contaminating source (red and green); an approximation to the same values from Swain et al. (2013) are also plotted (blue). Bottom: the same analyses as above, but both the analytical models and the Swain et al. (2013) results have the values from our fitting subtracted, in order to better show the discrepancies. The results from our PSF fitting match very closely with the high-temperature limit for the temperatures of both the primary source (WASP-12 A) and the contaminating source (WASP-12 BC); the low-temperature model shows a much larger signature of absorption from water vapor between 1.35 and 1.6 μm. The Swain et al. (2013) results are similar at most wavelengths, but there is a very large discrepancy at the shortest wavelengths.

(A color version of this figure is available in the online journal.)

analysis by Crossfield et al. (2012) a priori, without attempting to empirically determine the contaminating flux. Alternately, Swain et al. (2013) performed a similar fit to ours, but their results appear to lack the sharp downturn shortward of 1.15 μm and the upturn longward of 1.55 μm that are evident in our results. The slope of their results is also slightly shallower (a linear approximation to their results is plotted in Figure 7, and they did not publish their fitted values, but they are close to a single linear trend with a slight decrease between 1.34 and 1.48 μm). Given the close similarity between the high-temperature analytical model and our empirical fit to the data, we remain confident that our results are robust. However, it is clear that the choice of the spectral dependence for the dilution by WASP-12 BC has a significant impact on the final results for the spectrum of WASP-12 b; uncertainties of 1% for the dilution factor for WASP-12 BC will result in a difference of 150 ppm in the final transit depth, which is similar in magnitude to the uncertainties for the transit depths of our individual bins. We discuss this impact further in Section 5.

4.3. Band-integrated Transit Curve Fitting

Our analysis strategy relies on the assumption that almost all of the time-dependent trends present in the band-integrated time series are consistent across wavelength (even if the amplitudes of these trends change), since the systematics are related to either the general exposure parameters (array size, number

of read-outs, etc.), and/or correlated with the illumination of each pixel. We therefore decided to determine the band-integrated transit curve parameters first, and then use the residuals from this band-integrated fit as a component in our transit model when fitting individual spectral channels (with the amplitude of this component allowed to vary). This method allows us to incorporate any common-mode systematic trends into our fit, providing a more robust measurement of the relative change in transit depth across spectral channels, which is the most important factor when measuring the depth of spectral absorption features. We are also able to include the first orbit for each target into the wavelength-dependent analysis since the higher scatter in this orbit (which has caused most observers to discard it) is common across wavelength and can be removed accurately. We describe the fitting strategy in more detail in Section 4.4.

To achieve the best possible fit to the band-integrated light curve prior to fitting individual spectral bins, we utilized the divide-oot method developed by Berta et al. (2012), which uses the systematics in the out-of-transit data to correct the in-transit data by simply dividing all orbits by an average of the out-of-transit orbits. This method works very well to remove the repeated intra-orbit slope and buffer-ramp effects, which represent the largest instrumental effect in our data. We then fit the corrected light curve with a Markov Chain Monte Carlo (MCMC) routine with a Metropolis–Hastings algorithm within the Gibbs sampler (Ford 2005), using the light curve model from Mandel & Agol (2002), with an additional linear slope term to account for the gradual decrease in flux seen in all WFC3 exoplanet transit data to date.

All of the orbital parameters in our transit light curve model were locked to the literature values (see Table 2), since we are only analyzing single transits and lack full-transit coverage. The only exceptions are the mid-transit time and the two parameters for a quadratic limb darkening law, which we allow to vary under Gaussian priors since we are only analyzing a single transit with incomplete coverage of ingress and egress. For mid-transit times, we calculate the predicted mid-transit time from recent transit observations of our targets in the literature, and propagate the uncertainty on period in time to use as the width of our prior. For limb darkening, we use values calculated by Claret & Bloemen (2011) from analysis of ATLAS models. After selecting for the appropriate stellar parameters, Claret & Bloemen (2011) provide values at the centers of the J and H bands, with a choice between a least-square and flux conservation method. We interpolated between the J and H band points to find the central wavelength of our spectra, and took the average between the two methods as our starting limb-darkening parameter value. We used the standard deviation between the two methods, multiplied by two, as the width of our priors.

For each light curve we ran three MCMC chains with 100,000 links for analysis, with an additional initial burn period of 25,000 links. Our band-integrated time series for each of our targets are shown in Figure 3, with the best-fit transit curve overlaid; we tabulate our best-fit orbital parameters in Table 3. Our best-fit limb darkening parameters compare well with the expected values from Claret & Bloemen (2011), and best-fit mid-transit times are within the uncertainties based on prior measurements (see Figure 8).

4.3.1. Fitting for a Possible Thermal Contribution and Starspots

After fitting the integrated-light time series using the standard transit model, we determined that there appeared to be

Table 2
Stellar and Orbital Parameters Used for Model Fitting and Comparison

Parameters	WASP-12 b[a]	WASP-17 b[b]	WASP-19 b[c]
Period (days)	1.09	3.73	0.789
i (°)	82.5 ± 0.8	86.7 ± 0.500	79.5 ± 0.500
R_p/R^*	0.117 ± 0.00068	0.123 ± 0.037	0.139 ± 0.0457
T_c	55663.199	55750.285	55743.532
μ_1[d]	0.127 ± 0.0487	0.0901 ± 0.0487	0.153 ± 0.0487
μ_2	0.271 ± 0.0620	0.273 ± 0.0620	0.293 ± 0.0620
a/R^*	3.03 ± 0.0220	6.96 ± 0.0220	3.57 ± 0.0460
e	0.0447 ± 0.00430	0.00	0.00770 ± 0.00680
ω (°)	94.4 ± 0.0300	0.00	43.0 ± 67.0
Semi-major axis (AU)	0.02309 ± 0.00096	0.05105 ± 0.00128	0.01616 ± 0.00024
M_* ($M_\odot$)	1.38 ± 0.18	1.286 ± 0.079	0.904 ± 0.040
$M_p \times \sin i$ (M_J)	1.378 ± 0.181	0.477 ± 0.033	1.114 ± 0.04
Spectral type	G0	F4	G8V
H-band magnitude	10.228	10.319	10.602
[Fe/H]	0.3 ± 0.1	-0.25 ± 0.09	0.02 ± 0.09

Notes.
[a] Values from Southworth et al. (2012).
[b] Values from Maciejewski et al. (2013).
[c] Values from Lendl et al. (2013).
[d] Values for limb darkening derived from Claret & Bloemen (2011) quadratic limb darkening tables.

Table 3
Fitted Parameters From Band-integrated Time Series

| Parameters	WASP-12 b	WASP-17 b	WASP-19 b
R_p/R^*	0.11895 ± 0.0013	0.12316 ± 0.00058	0.14140 ± 0.00093
μ_1	0.085 ± 0.024	0.083 ± 0.031	0.092 ± 0.025
μ_2	0.281 ± 0.034	0.256 ± 0.046	0.305 ± 0.027
Mid-transit (MJD)	$55663.199736 \pm 0.000065$	55750.294793 ± 0.00088	$55743.532268 \pm 0.000040$
Slope[a]	-0.00793 ± 0.00034	-0.00578 ± 0.0010	-0.00407 ± 0.00039

Note. [a] Linear slope has units of normalized flux per day.
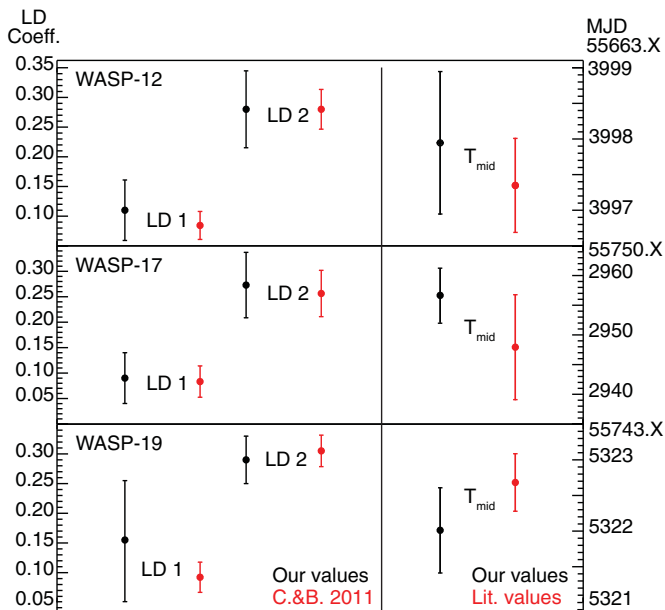
Figure 8. Limb darkening parameters for a quadratic limb darkening law shown as calculated using models from Claret & Bloemen (2011), and as found by our MCMC routine, using the Claret & Bloemen (2011) models and uncertainties as priors. Our final values match the expected values within uncertainties for all targets.

(A color version of this figure is available in the online journal.)

systematic deviations in the residuals of the out-of-transit orbits for both WASP-12 and WASP-19 as well as the in-transit orbit for WASP-19. The out-of-transit orbits appear to have trends in flux that are not perfectly fit by a single linear slope, with the first orbit having a steeper slope while the last orbit has a shallower slope (see Figure 9). It is difficult to determine the source of these trends due to the limited sampling in the orbital phase and the necessity of using the divide-oot correction method, which combines the data from all out-of-transit orbits (and therefore mixes underlying trends and/or red noise together). The current data can be fit using a second-order polynomial, or fit using a more physically motivated model including a sinusoidal component with a period equal to the planetary orbital period, representing the thermal phase variation due to the day–night temperature difference (Knutson et al. 2007). Either model results in a better fit to the data than the linear slope for WASP-12 and WASP-19, and we decided to use the Bayesian Information Criterion (BIC; Schwarz 1978; Liddle 2004) to determine whether the improvements from either of the more complex baseline models was sufficiently significant. The BIC includes a strong penalty for including additional parameters, and therefore provides a robust technique to distinguish between models; ΔBIC $\geqslant 2$ is considered to be positive evidence against the null hypothesis. The BIC was not increased using the nonlinear baseline models for either target (ΔBIC ~ -0.5). However, the best-fit peak-to-trough amplitude of 0.0018 ± 0.0006 for a possible sinusoidal component in the WASP-12 data is within the range predicted for the thermal phase variations of very hot planets (Cowan &
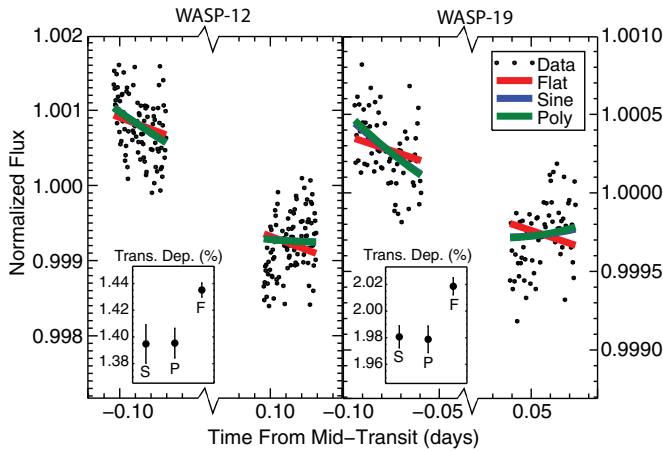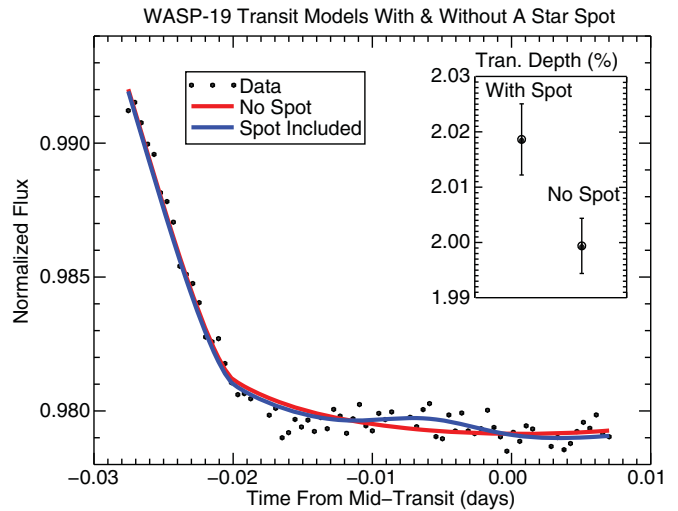
Figure 9. Out-of-transit portions of the band-integrated light curves for WASP-12 (left) and WASP-19 (right), with models including only a linear trend (red) and an additional sinusoidal component (blue) or 2nd-order polynomial function (green) overplotted. The best-fit transit depths for each model are also plotted (inset). The addition of sinusoidal or polynomial components produce a marginally better fit, but the improvements are not sufficient to yield a lower BIC.

(A color version of this figure is available in the online journal.)

Figure 10. Trough of the transit for the band-integrated light curve for WASP-19, with models including standard transit model (red) and a model with a star spot (blue) overplotted. The best-fit transit depths for each model are also plotted (inset). The value derived incorporating the spot model has a larger uncertainty from MCMC due to the additional free parameters, but the effects of red noise are not included and therefore the uncertainty on the spot-free fit is underestimated.

(A color version of this figure is available in the online journal.)

Agol 2011), though it is smaller than the value measured using *Spitzer* (Cowan et al. 2012). The best-fit amplitude for WASP-19 is similar to WASP-12 (0.0016 ± 0.0007). We conclude that due to the low significance of the fit, the limited time sampling and ambiguities introduced by the divide-oot method, the nature of the curvature is highly uncertain and therefore must be investigated with more complete observations before conclusions as to its validity or physical nature can be made. The light curve for WASP-17 does not include a post-egress portion, so we cannot evaluate the presence of a curved baseline.

The in-transit orbit of WASP-19 also has a region just after second contact (after the end of ingress) which deviates slightly from a standard transit curve (see Figure 10). The amplitude and duration of the deviation is similar to the amplitude and duration of starspots detected in optical transit data by Tregloan-Reed et al. (2013), so we experimented with including a Gaussian-shaped spot in our transit model. The spot model leads to a statistically better fit with ΔBIC = 7.8 (see Figure 10), leading us to adopt a model including a sunspot modeled as a Gaussian with a position centered at MJD 55743.526, a relative amplitude of 0.06%, and a width of 0.0036 days. We locked the amplitude of the spot when fitting each of the bins, since our data quality is insufficient to determine variations with wavelength. Neither of our other data sets showed evidence for star spots, which is expected since both WASP-17 and WASP-12 are significantly hotter than WASP-19.

Considering the ambiguity regarding the presence of additional visit-long components and star spots, we decided to use the average of all the model fits with and without a sinusoidal component or a spot for the band-integrated transit depth listed in Table 3, and augment the uncertainty values to encompass the full range of values. This increases the uncertainty by a factor of ~4 for WASP-19 and a factor of ~5 for WASP-12. To remove these ambiguities in the band-integrated transit depth, we would need a fully-sampled light curve and multiple visits to settle the question of spots. However, since we lock the values for any nonlinear or spot components when fitting the bins, the final choice of the best-fit band-integrated model makes no difference in the relative depths for our wavelength bins.

4.4. Fitting the Spectrally Binned Light Curves

Once we determined an adequate fit to the band-integrated light curves, we used the residuals of the fit to remove systematics common to all spectral channels (or bin of channels). Our transit models for each individual channel include a constant scaling of these residuals, with the scale factor varying as a free parameter. This strategy is similar to methods developed independently by Deming et al. (2012) and Stevenson et al. (2013) (though without a scaling term for modulating the amplitude of the band-integrated residuals), and it obviates the need for using the divide-oot method. Additionally, we introduced two more components into the light curve model (each with a scaling factor as a free parameter) based on our measurements of the horizontal and vertical shifts of the spectrum on the detector over time. The scaling factors for these components are insignificant for most bins, since a small shift for most points on the spectrum will not change the flux significantly. However, near spectral features or near the edges of the spectrum, these shifts can cause the flux within a single bin to drift up or down. Our final model light curve for comparison with the data takes the form

$$LC_{final} = LC_{transit} * (a + bt + C_1 * Res_{BI} + C_2 * Shift_y + C_3 * Shift_x), \qquad (2)$$

where $LC_{transit}$ is the light curve model calculated using the Mandel & Agol (2002) prescription, a and b are coefficients for a linear trend with time, Res_{BI} are the residuals from the band-integrated light curve, and the C coefficients are scaling parameters determined through our MCMC fitting.

For the light curve for each spectral bin we followed the above methods for bad pixel and bad channel correction and then fit for the best model using MCMC. We locked the same parameters as with the band-integrated light curve, and additionally locked the limb darkening and mid transit time to the best-fit values from the band-integrated light curve analysis. This allows us to measure the relative change in transit depth while maintaining
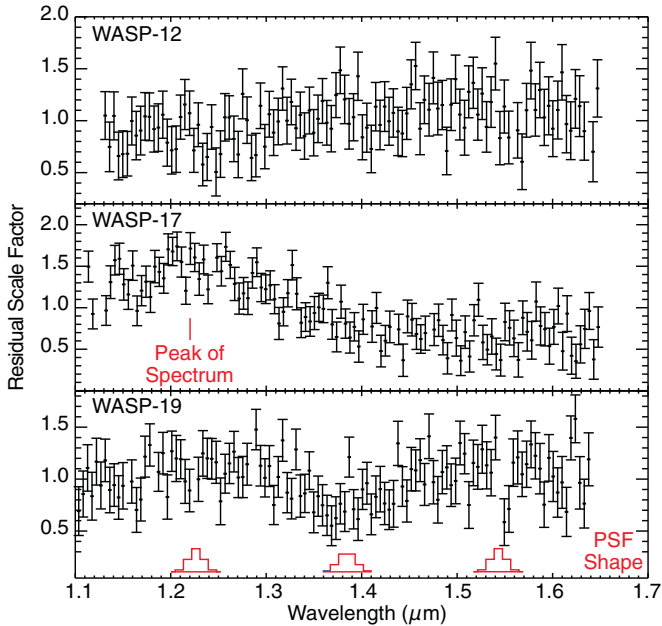
Figure 11. Best-fit scaling factors for the band-integrated light curve residuals derived for each channel (see Section 4.4). The relative amplitude of the scaled-residuals component of the model changes with wavelength based on the peak illumination in each channel, and varies between targets based on the sub-array size and sampling mode (see Table 1). For WASP-17, the scale factor peaks at the location of the peak flux in the spectrum, while for WASP-19, the scale factor varies based on the sampling of the spatial PSF. WASP-12 has very little structure in the band-integrated residuals, and therefore shows no clear correlation with flux.

(A color version of this figure is available in the online journal.)

the same transit shape. We experimented with fitting for the limb darkening parameters using priors based on a linear interpolation between the J and H-band values from Claret & Bloemen (2011), but we determined that there was no change in the final transit depths compared with exclusively using the band-integrated values.

In each bin, the importance of the different systematic trends varies. The amplitude of the common-mode residuals is related to (but not directly correlated with) the peak intensity in each channel (see Figure 11), and the x shift is only important near spectral features or other steep gradients in the spectral direction. To avoid including unnecessary components in our light curve model, we examined the importance and validity of including each model parameter using a nested model selection analysis. We began by assuming that the values determined for the band-integrated light curve except for R_p/R_* and the mean value of the out-of-transit flux would be valid for all the bins. We then calculated ΔBIC for models with the inclusion of free parameters for the slope of the linear trend, the scale factor for the band-integrated residuals, and scale factors for components based on the x and y shifts. We only included the parameters that provided an improvement in the BIC (ΔBIC ≥ 2) over the model that locked that parameter. The ΔBIC values for each of our 0.027 μm wide bins for each of our targets are shown in Figure 12. To further confirm that we are not over-fitting our data, we searched for correlations between different free parameters in our light curve model and the final transit depths. Most of the parameters in most of the bins remain locked to the band-integrated values (the slope of the linear trend remained locked for every bin for all targets), and we see no evidence of correlations between parameters for the fitted parameter values in any of our targets (see Figure 13).
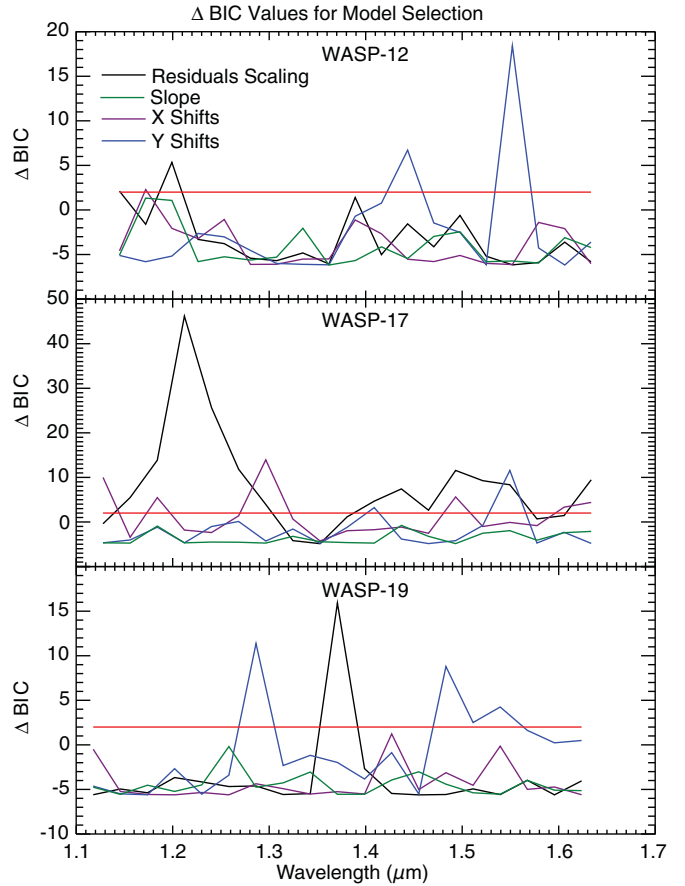

Figure 12. We calculate the change in BIC values for a model that fits for additional systematic trends (band-integrated residuals, the visit-long linear slope, x shift, y shift) compared with the default model (see Section 4.4). ΔBIC is shown for each of the 19 bins, for all targets (top: WASP-12, middle: WASP-17, bottom: WASP-19). The horizontal red line at zero indicates the level above which parameters are said to be significant—parameters are only allowed to vary from the best-fit band-integrated values if they have ΔBIC ≥ 2.

(A color version of this figure is available in the online journal.)

In Figure 14 we show final light curves for all of our 0.027 μm wide bins for each target after the various best-fit systematic trend components have been removed; they are overplotted with the best-fit transit light curve model. The light curves show no sign of correlated noise, and the posterior distributions (shown in Figure 15) are all fit well by a Gaussian distribution. Our final spectra for each of our science targets are shown in Figure 16. We plot the best-fit transit depth values for each individual channel, and two bin sizes (0.027 μm and 0.1 μm). The individual channels clearly show a high point-to-point scatter which appears to be largely due to photon noise, so we experimented with binning the channels using sequential bin sizes (two channels, three channels, etc.). The rms of the resulting spectra drops off quickly, but then stays elevated above the photon-noise limit for all stars beyond a five-channel bin width, suggesting structure in the spectrum on scales larger than 5 pixels (see Figure 17). We therefore chose to use the six-channel bins (0.027 μm) for our final spectrum, since they will largely conserve the overall structure of the individual-channel spectrum while decreasing the photon noise considerably and allowing for improved removal of systematic trends. Larger bin sizes, as used by Stevenson et al. (2013) and Huitson et al. (2013), do not fully encapsulate the structure in the smaller-bin spectrum. This smoothing is not incorporated into the uncertainty limits for the wider bins since the uncertainty is

Table 4
Derived Transit Depths for Binned Data

WASP-12 b		WASP-17 b		WASP-19 b	
λ (μm)	Transit Depth (%)	λ (μm)	Transit Depth (%)	λ (μm)	Transit Depth (%)
1.145	1.4131 ± 0.0235	1.128	1.5087 ± 0.0257	1.118	2.0159 ± 0.0175
1.172	1.4211 ± 0.0232	1.156	1.4867 ± 0.0250	1.146	2.0241 ± 0.0206
1.199	1.4302 ± 0.0224	1.184	1.5044 ± 0.0259	1.174	1.9905 ± 0.0172
1.226	1.4417 ± 0.0226	1.212	1.4957 ± 0.0216	1.202	2.0071 ± 0.0180
1.253	1.4376 ± 0.0224	1.240	1.4998 ± 0.0222	1.230	1.9269 ± 0.0189
1.281	1.4103 ± 0.0230	1.268	1.5166 ± 0.0226	1.258	1.9880 ± 0.0180
1.308	1.4143 ± 0.0207	1.296	1.4822 ± 0.0237	1.286	1.9941 ± 0.0187
1.335	1.4387 ± 0.0190	1.325	1.5362 ± 0.0197	1.314	2.0176 ± 0.0151
1.362	1.4338 ± 0.0186	1.353	1.5545 ± 0.0223	1.343	1.9943 ± 0.0174
1.389	1.4419 ± 0.0225	1.381	1.5686 ± 0.0239	1.371	2.0318 ± 0.0168
1.416	1.4414 ± 0.0207	1.409	1.5050 ± 0.0261	1.399	2.0317 ± 0.0157
1.443	1.4322 ± 0.0217	1.437	1.5578 ± 0.0250	1.427	2.0546 ± 0.0176
1.471	1.4505 ± 0.0237	1.465	1.5446 ± 0.0267	1.455	2.0363 ± 0.0171
1.498	1.4719 ± 0.0231	1.493	1.5300 ± 0.0247	1.483	1.9923 ± 0.0196
1.524	1.4645 ± 0.0229	1.521	1.5086 ± 0.0229	1.511	2.0470 ± 0.0187
1.552	1.4707 ± 0.0286	1.549	1.5410 ± 0.0316	1.539	2.0053 ± 0.0205
1.579	1.4170 ± 0.0296	1.577	1.5534 ± 0.0282	1.568	2.0350 ± 0.0196
1.606	1.4264 ± 0.0329	1.606	1.4875 ± 0.0278	1.597	2.0578 ± 0.0197
1.633	1.4073 ± 0.0400	1.634	1.4530 ± 0.0303	1.624	2.0142 ± 0.0188
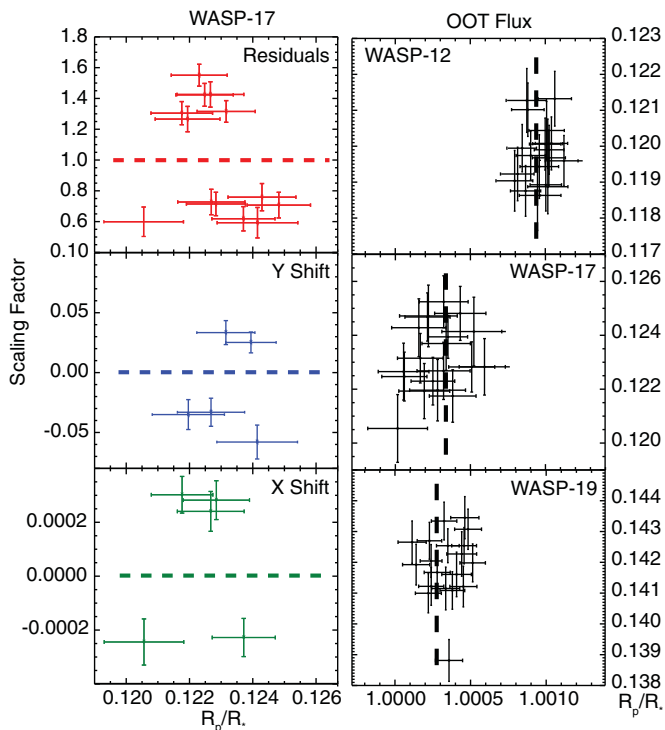
Figure 13. Left: correlation plots for the three model components vs. R_p/R_*, for WASP-17 (see Section 4.4). Parameters were only allowed to vary for those bins in which doing so resulted in ΔBIC $\geqslant 2$, and only the bins in which the parameters varied are plotted; the dotted lines represent the default value from the band-integrated results. Results for WASP-12 and WASP-19 are not plotted because the number of bins with open parameters for each component was small (1–4). Right: best-fit out-of-transit flux vs. R_p/R_* for all the targets. No correlation is seen between R_p/R_* and any of the parameters.

(A color version of this figure is available in the online journal.)

purely based on the goodness-of-fit of the transit model. We therefore believe the use of bin sizes <0.03 μm is necessary to avoid misinterpretation of spectral characteristics. The best-fit transit depths for the 0.027 μm wide bins for all of our targets are listed in Table 4.

4.5. Uncertainty Analysis

The uncertainty limits for our light curve parameters were derived from the widths of our MCMC posterior probability distributions. However, the uneven sampling before and after a transit, as well as across a transit event due to the gaps in the *HST* orbit, make the calculation of the expected noise limit difficult. We therefore decided to construct synthetic data sets for each of our targets in order to identify the different contributing sources of uncertainty in the final results, with each synthetic data set for an exoplanet constructed using the best-fit parameters from the fit to our band-integrated light curve and the timing array of our real data. Stochastic Gaussian noise was injected at the level of the final rms determined for our data, and the synthetic data was fit using MCMC in the same method described above for the real light curves. Since each data set has a relatively small number of data points (131 for WASP-17, 274 for WASP-19, and 484 for WASP-12), the impact of outliers due to purely stochastic noise can have a considerable effect, so we repeated this process 100 times with different randomly generated noise distributions in order to determine the range of uncertainties produced by MCMC. We can then compare the predicted noise based on the number of points in transit to the predicted uncertainty from MCMC fits to the synthetic data to estimate the increase in uncertainty due to the uneven sampling of the light curves. Also, by comparing the uncertainty derived for our real data to the range of uncertainties for the simulated data sets, we can estimate the amount of additional (red) noise in our data compared with a purely (white) stochastic noise distribution.

We also explored the use of residual-permutation analysis (RP) to estimate the effects of red noise. We fit the light curves using Levenberg–Marquardt least-squares fitting, subtracted the best fit model from the light curve, shifted the residuals by one position, and then added the model back in and re-fit the data, cycling through all the data points in each light curve. However, we found that with such a small number of data points in our light curves and the uneven sampling of the *HST* orbits, the RP method is not sufficiently robust; the final distributions for the fitted values of R_p/R_* showed a large scatter without any
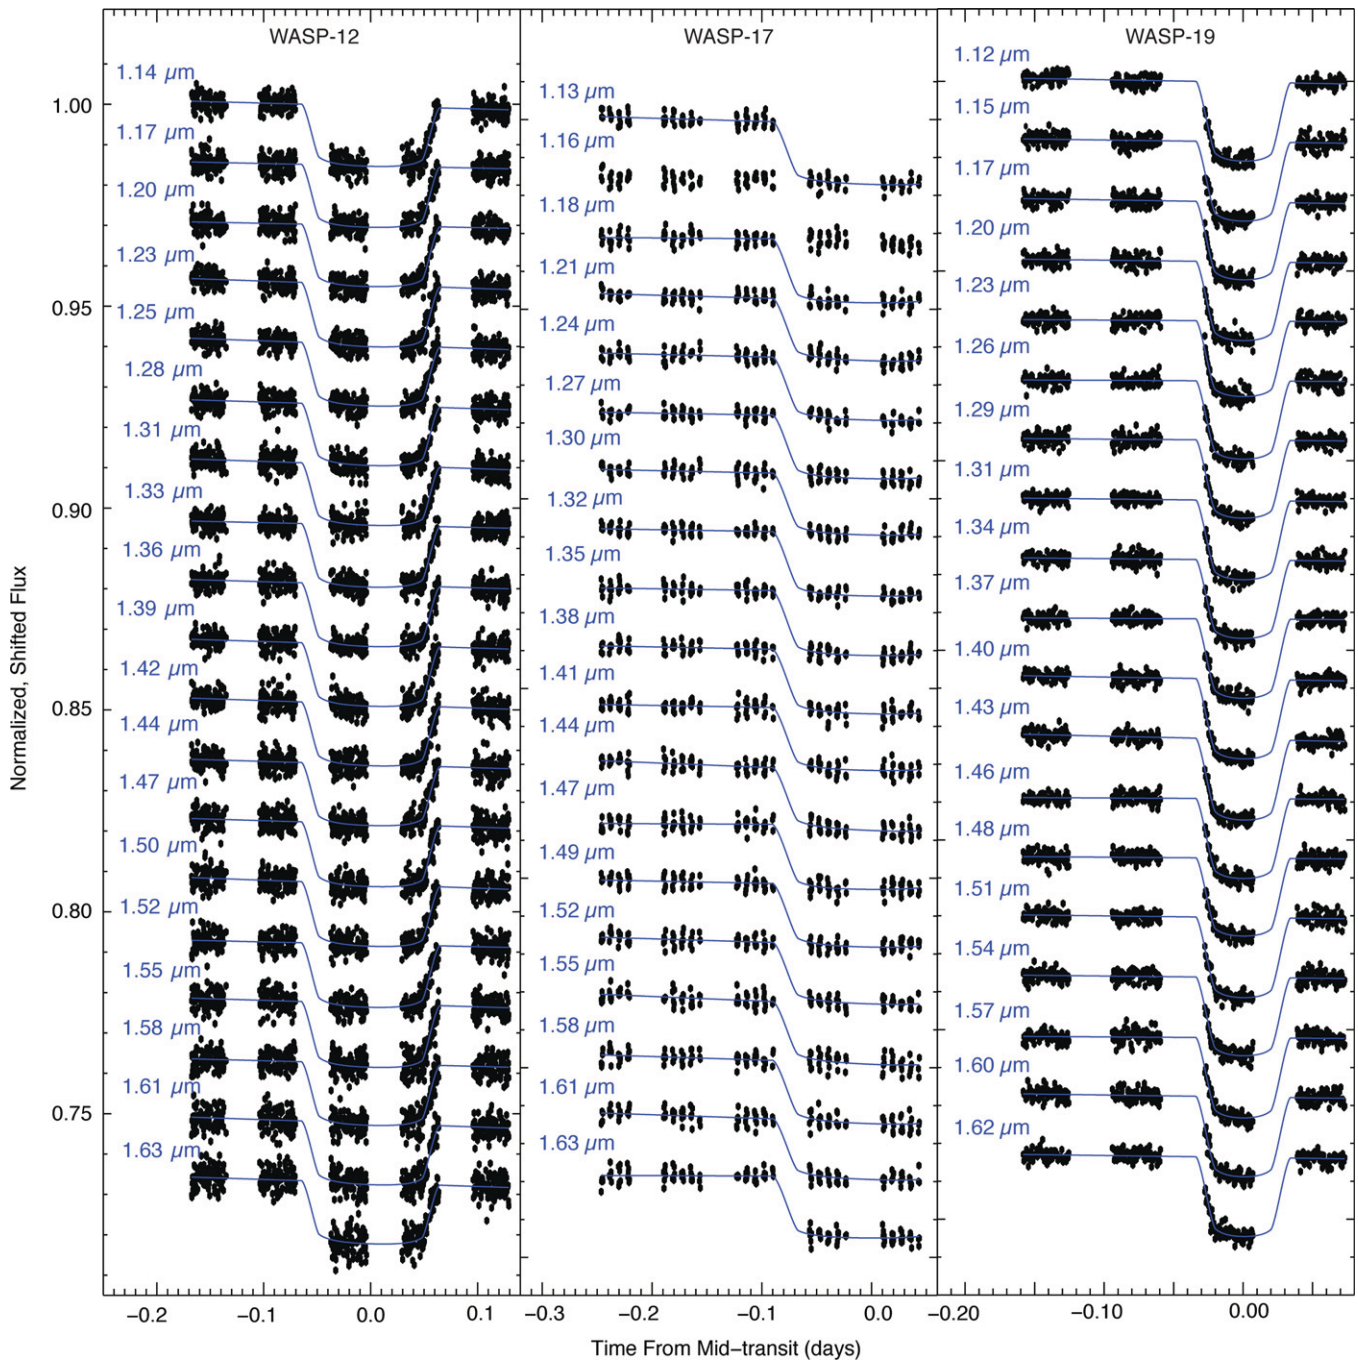
Figure 14. Final results for all the bins for each target are shown in black, after removing time series components based on the scaled residuals from the band-integrated light curve, as well as any scaled components based on the spectral shift in the *x* and *y* directions that were deemed statistically significant (see Section 4.4). The best-fit transit model from our MCMC analysis is shown in blue. The light curves all show essentially white noise, with no evidence of correlated noise or remaining systematic trends.

(A color version of this figure is available in the online journal.)

clear pattern. We therefore relied on our simulated data tests to determine how close we came to the expected photon noise.

The band-integrated photon noise statistics, rms uncertainty, and uncertainties in transit depth determined from MCMC fitting for the real and synthetic data sets are shown in Table 5. We find that the rms of the data is 1.2–1.44× the expected photon noise for band-integrated time series, but only 1.11–1.22× the photon noise limit for the binned data. For WASP-12 the MCMC results for the synthetic data match within a few percent to the predicted uncertainties based on the rms, suggesting

that the impact of light curve sampling is minimal. The real band-integrated data for WASP-12 are slightly noisier than the synthetic data suggesting some correlated noise, most likely due to trends in the out-of-transit portion of the data discussed previously Section 4.3.1. The WASP-19 results are similar, though the MCMC uncertainties and the dispersion in the range of value for the synthetic data are larger than predicted due to the impact of fitting for the presence of a spot (Section 4.3.1). For WASP-17, the uncertainty for the synthetic data is more than 2× larger than the predicted uncertainty,
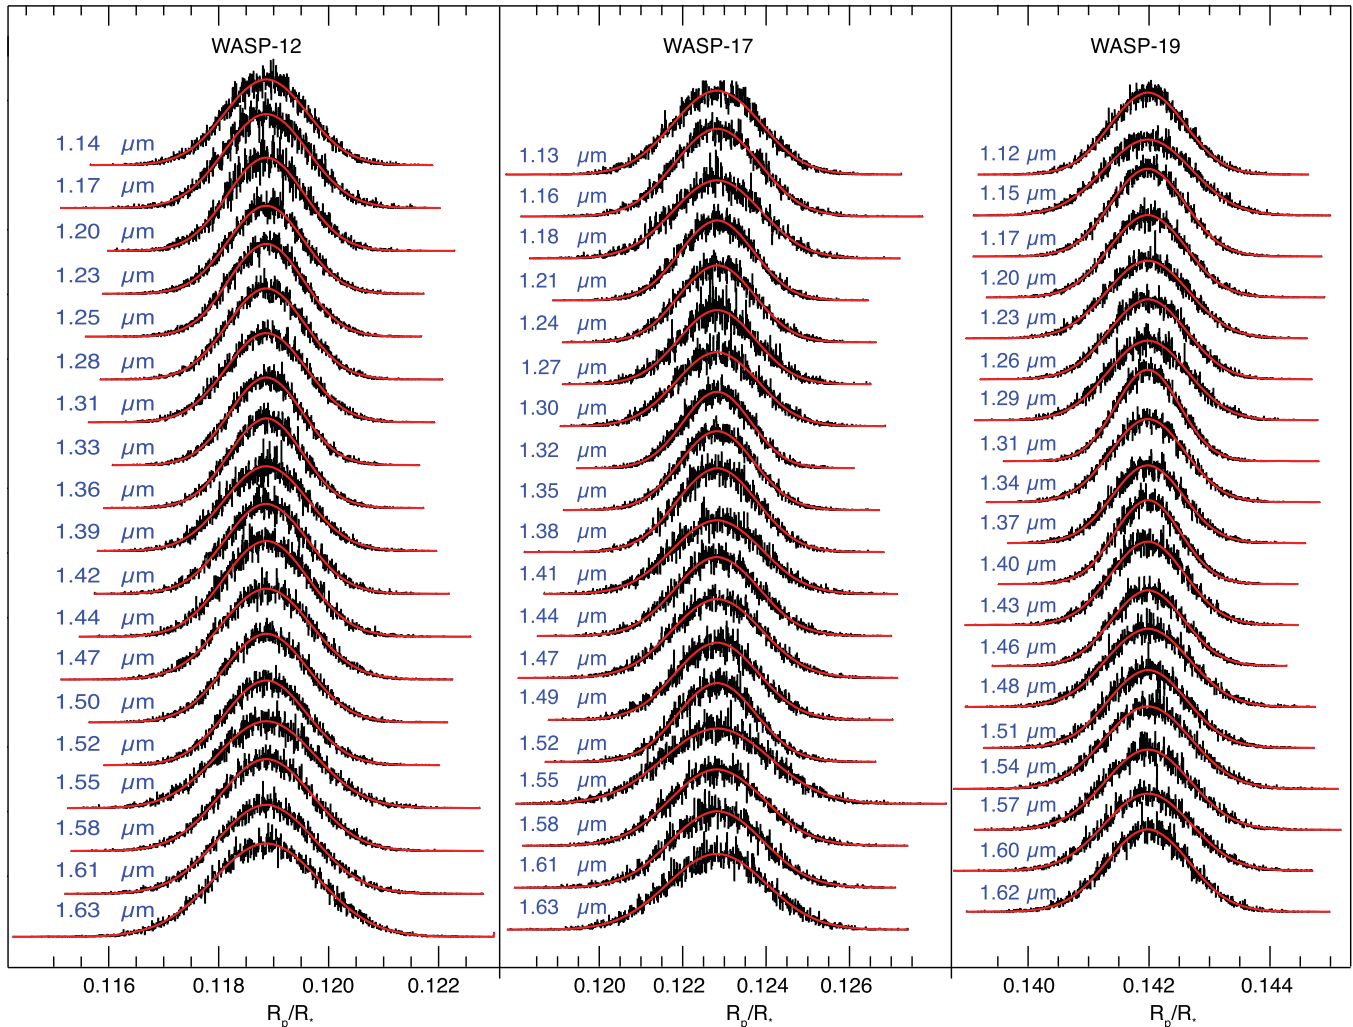
Figure 15. Posterior distributions from MCMC for R_p/R_* for every bin, for each of the three targets. All of the final distributions are symmetric and well-approximated by a Gaussian fit (red).

(A color version of this figure is available in the online journal.)

due to the lack of data covering ingress/egress or post-transit. However, we note that the effects of sampling and correlated noise are almost completely neutralized in the binned data by our residual subtraction—the ratio of the uncertainty for the simulated data to the analytical prediction for all the targets drops to essentially unity, demonstrating the effectiveness of our component removal method.

5. DISCUSSION

The observations analyzed in this study represent a preliminary sample of hot exoplanets observed with the WFC3 instrument on *HST*. The three planets include two extremely hot planets with temperature structures constrained by *Spitzer* occultation data (WASP-12 b and WASP-19 b) as well as a cooler planet with a highly inflated planetary radius (WASP-17 b), allowing us to investigate two classes of planets that pose significant challenges for current theories of exoplanet structure and evolution.

5.1. Comparison with Atmospheric Models

Absorption band depths in transit spectra probe the line of sight through the terminator of the planet, and are primarily sensitive to a combination of the atmospheric composition and the scale heights over which each species is absorbing. These factors can be significantly degenerate and it is difficult to place strong constraints on the overall abundances of different species with observations in only a single wavelength band. We therefore reserve a detailed examination of constraints on atmospheric composition and structure to a later study, and restrict our current analysis to a discussion of the general implications of qualitative comparison with several different sets of models.

In Figure 18 we plot the data for each planet and overplot two different sets of models, which utilize different strategies for constraining the atmospheric structure and composition. One set (top in Figure 18) is based on the framework of Burrows et al. (2000) and more recently Burrows et al. (2006), Burrows et al. (2008) and Howe & Burrows (2012). The Burrows models calculate the chemical and radiative equilibrium state of each planet based on the mass, size, and incident radiation, assuming solar abundances. The spectra were then calculated by combining day- and night-side model atmospheres joined at the terminator. Adjustments were made to the abundance of important molecular absorbers such as H_2O, CH_4, and CO and/or the inclusion of additional absorbers that affect
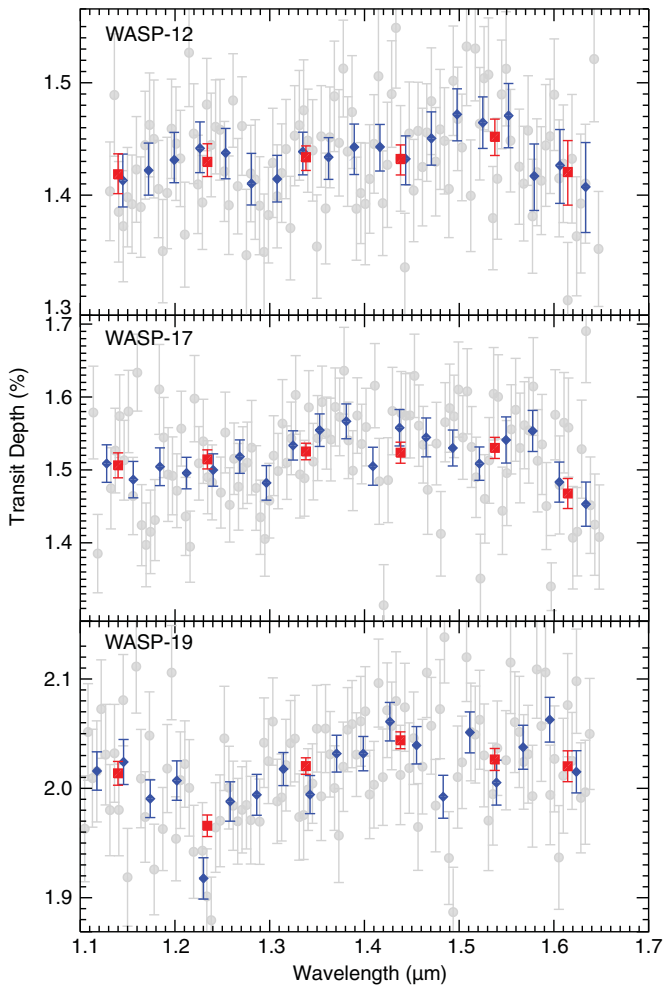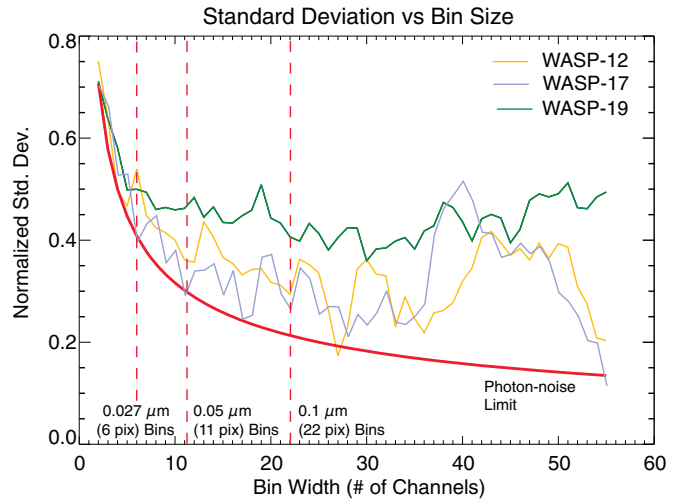
Figure 16. Final spectra for each of our targets. The individual channel depths are shown in gray, with the results for 0.027 μm wide (blue) and 0.1 μm wide (red) bins overplotted. The differences between the channels and the 0.027 μm wide bins are consistent with photon-noise variations, but the 0.1 μm wide bins appear to remove structure in the spectra that could be significant. We therefore chose to use the 0.027 μm wide bins in our analysis.

(A color version of this figure is available in the online journal.)

Figure 17. After fitting for the transit depths using individual channels, we binned the spectra using bin sizes between 2 and 55 points and then calculated the standard deviation of each binned spectrum; the results for each star are plotted as well as the expected relationship based on photon-noise statistics alone. The standard deviation for all the targets is approximately photon-limited up to five-channel bins, but then levels off. We use a six-channel bin size for our final results, and spectra using both the six-channel bins and 22-channel (0.1 μm) bins are shown in Figure 16 for comparison.

(A color version of this figure is available in the online journal.)

Table 5
Uncertainty Analysis

| Parameters	WASP-12	WASP-17	WASP-19
Data points during transit	196	54	70
Data points out of transit	288	77	204
Band-integrated time series			
Photon noise (ppm)	357	279	255
RMS of residuals (ppm)	515	350	305
Predicted[a] σ_{td} (ppm)	52	67	45
σ_{td} from MCMC, data	53	144	65
σ_{td} from MCMC, sm.[b]	53 ± 2	145 ± 13	63 ± 11
RMS/photon noise	1.44	1.26	1.20
Data/pred.	1.02	2.15	1.44
Sim./pred.	1.03 ± 0.04	2.16 ± 0.19	1.40 ± 0.24
0.027 μm bin width (19 total)			
Photon noise (ppm)	1560	1220	1110
rms of residuals (ppm)	1880	1400	1230
Predicted[a] σ_{td} (ppm)	174	249	170
σ_{td} from MCMC, data	180	257	180
σ_{td} from MCMC, Sim.[b]	181 ± 6	242 ± 14	187 ± 11
rms/photon noise	1.22	1.15	1.11
Data/pred.	1.02	1.03	1.06
Sim./pred.	1.03 ± 0.03	0.97 ± 0.06	1.1 ± 0.06

Notes.
[a] Calculated from the residual rms and the number of points during transit and out of transit.
[b] Simulated data was created with a sampling equivalent to that of the real data, and an rms equivalent to the rms of the final residuals.

the temperature structure and/or broadband optical depth of the atmosphere with the goal of improving fits to multi-wavelength observations. For example, additional opacity at optical wavelengths is required to produce a thermal inversion postulated to explain *Spitzer*/IRAC photometric measurements during occultation for a number of planets including WASP-12 b (Cowan et al. 2012; Crossfield et al. 2012) and possibly WASP-19 b (Anderson et al. 2013). TiO has been considered the most likely candidate (Hubeny et al. 2003; Fortney et al. 2008), but the lifetime for TiO in the upper atmosphere may be problematic for this hypothesis (Spiegel et al. 2009), and recent searches for spectral features of TiO have been unsuccessful (Huitson et al. 2013). On the other hand, a haze or dust with opacity through the optical and NIR is required to fit measurements of molecular absorption features for several hot Jupiters (Charbonneau et al. 2002; Pont et al. 2013; Deming et al. 2013). While the physical nature of these absorbers is currently unclear, we can test how different opacities for these parameters affect the model spectra in our wavelength region.

The Burrows models, which are characterized by broad H_2O absorption at 1.4 μm that slopes consistently downward toward longer wavelengths, fit the data for WASP-17 b reasonably

well—both a standard model and an isothermal model with haze yield a lower BIC (assuming 3 degrees of freedom) than simply fitting a line to the data (2 degrees of freedom), with the best-fitting model (the hazy model) giving a $\Delta\chi^2 \sim 10$. A model with haze is required to reproduce the flat region shortward of 1.3 μm, and a haze hypothesis may gain additional support from the fact that the best fits to the models are
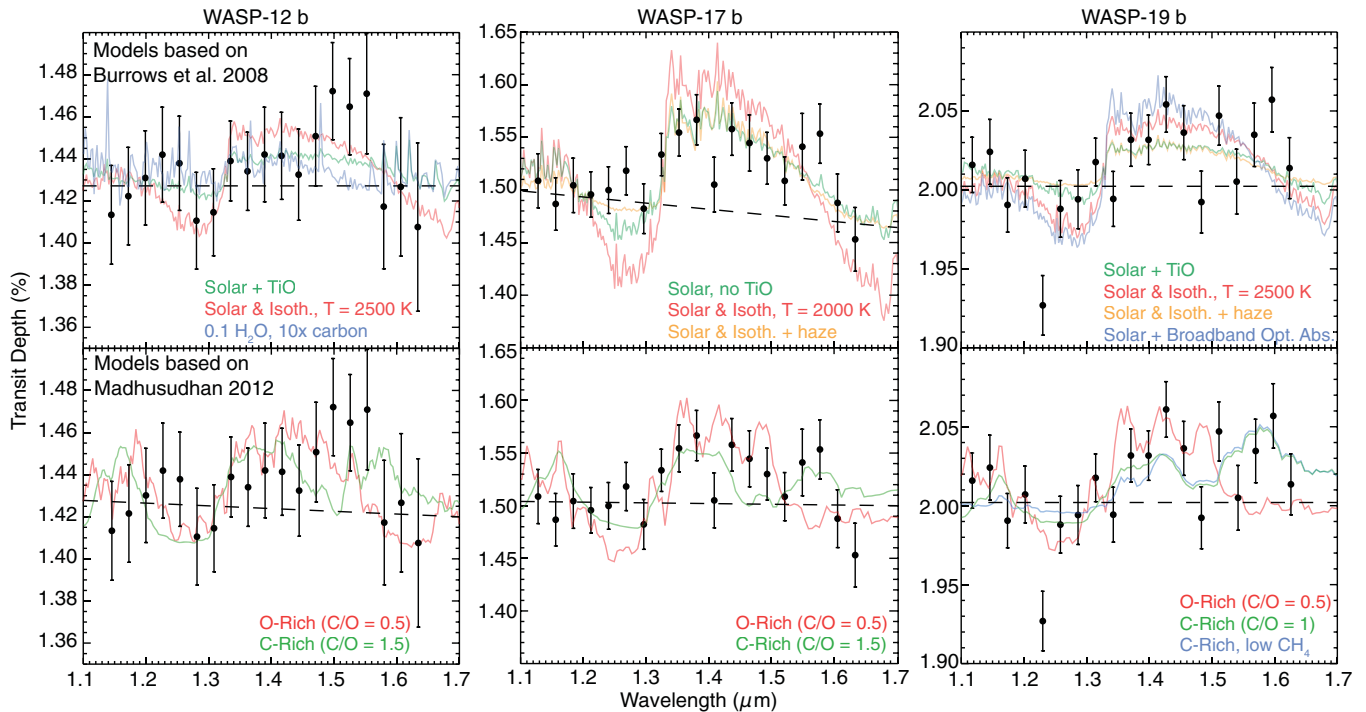
Figure 18. Transit depths for each of the 19 bins for each target, with models based on the framework of Burrows et al. (top) and Madhusudhan et al. (bottom). Standard models from Burrows et al. provide a good fit for WASP-17 b and a reasonable fit for WASP-12 b, but for WASP-19 b the models do not fit well beyond 1.45 μm. Models with a deep water absorption feature can also be adjusted to fit the data by adding an absorbing haze layer with an opacity of 0.01 cm^2 g^{-1}. The hazy model for WASP-17 b is further supported by the linear slope that is needed to match the models to the data. The oxygen-rich and carbon-rich models by Madhusudhan et al. fit equally well for WASP-12 band WASP-17 b, but for WASP-19 b the carbon-rich models provide a statistically better fit than the oxygen-rich models. However, except for WASP-17 b, the data is fit almost equally well by a flat spectrum, though WASP-19 b would require a very large scatter between the data points.

(A color version of this figure is available in the online journal.)

improved ($\Delta\chi^2 < 0$) in every case by including a linear trend to the models; we discuss the implications of these results in Section 5.2. However, the results for the two hotter planets are more ambiguous. The majority of the spectrum for WASP-12 b is consistent with a flat spectrum within the uncertainties ($\chi^2_{\rm red} = 0.57$), and the amplitude of the expected features do not allow us to discriminate between standard models with either an equilibrium temperature structure, an isothermal temperature structure suggested by Crossfield et al. (2012), or a model with a deficit of water and enhanced carbon abundance that best fits the analysis of *Spitzer*/IRAC occultation results by Cowan et al. (2012). WASP-12 b and possibly WASP-17 b also appear to have additional absorption in the region from 1.5–1.6 μm; these features are several bins wide, and do not appear to be the result of random noise. For WASP-19 b the results are even less consistent with the models—none of the models yield an improvement in BIC or χ^2 over a linear fit. The spectrum shows an increase in absorption beyond 1.35 μm, suggestive of H$_2$O, but does not include the consistent drop at longer wavelengths expected from the models and apparent in the WASP-17 b spectrum. Additionally, several bins in this region show a steep drop in absorption compared with the smooth downward trend expected from the Burrows models.

The second set of models we compare to our data (bottom in Figure 18) are based on the framework of Madhusudhan & Seager (2009) and Madhusudhan (2012), which relax the stringent requirements for radiative and chemical equilibrium in favor of flexibility when exploring the constraints on parameter space from available observations. In particular, the Madhusudhan models explore a range of C/O ratios for the overall composition of the atmosphere, and include a number of less abundant,

carbon-bearing species that may produce additional absorption features in NIR spectra at C/O $\geqslant$ 1. The models plotted roughly correspond to either an oxygen-rich chemistry (C/O $\sim$ 0.5, i.e., essentially the solar value) or a carbon-rich chemistry (C/O $\gtrsim$ 1) for specific temperature profiles (see Madhusudhan 2012 for details). It is clear that there are a number of overlapping spectral features that lead to degeneracies—the H$_2$O feature at 1.4 μm overlaps with CH$_4$ at 1.36 μm and HCN at 1.42–1.51 μm, while the H$_2$O feature at 1.15 μm overlaps with CH$_4$. The oxygen-rich and carbon-rich models primarily diverge between 1.45 and 1.65 μm, where the carbon-rich models include features from HCN and C$_2$H$_2$. While the additional absorption in WASP-17 b and WASP-19 b appears to line up well with these features, and produces an improvement in χ^2, the uncertainties in both our data and the range of potential model parameter values are large enough that we cannot discriminate between oxygen-rich and carbon-rich compositions based on these data alone.

We conclude that the data for all our targets are consistent for the most part with standard atmospheric models, but further improvements in S/N and a more comprehensive modeling strategy incorporating additional constraints on the molecular abundances and temperatures from other data sets are necessary to discriminate between them. In particular, the origin of significant deviations from the standard solar composition model predictions at wavelengths beyond 1.5 μm is unclear; these features could either be indicative of unexpected atmospheric absorption features, or they could be unexplained artifacts in the data. We have examined all of our data analysis routines in detail and we have found no obvious problems with the analysis of these bins, but repeated observations are necessary to confirm that the results are robust. We also point out the
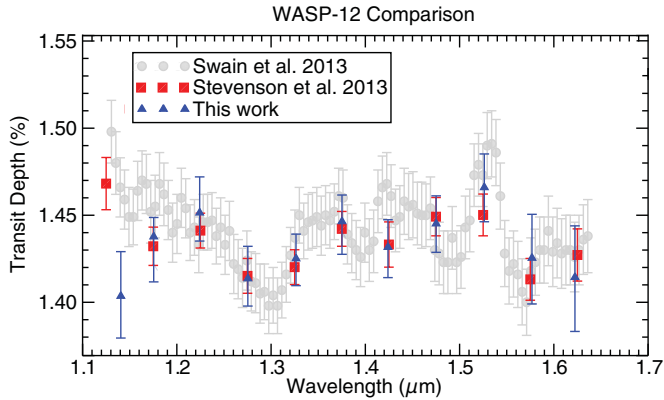
Figure 19. Results from Swain et al. (2013) shown in gray, results from Stevenson et al. (2013) shown in red, and from this work in blue. Results from this work have been binned to the same size and number of bins as those used by Stevenson et al. (2013), with edge bins offset due to different choices of spectral trimming. Results from Stevenson et al. (2013) have been shifted up slightly for comparison. The spectra are largely consistent, with the most noticeable offsets visible at the short edge of the spectrum.

(A color version of this figure is available in the online journal.)

importance of using appropriately sized bins to be sensitive to the possibility of narrower spectral features in the data. Figure 16 demonstrates that using bin sizes larger than ~0.03 μm significantly smoothes the data and has the potential to erase the signatures of small-scale fluctuations in the data.

5.2. Comparison to Previous Results

5.2.1. WASP-12

As mentioned previously, the data set that we analyzed for WASP-12 was originally observed and analyzed by Swain et al. (2013), and the data set has also recently been analyzed as part of a multi-wavelength study by Stevenson et al. (2013). Figure 19 shows our final spectrum for WASP-12 binned to match Stevenson et al. (2013) and plotted with the results from these two studies. While it is always difficult to pinpoint differences between independent analyses, there are two possible sources of significant variations between the results of the three different studies: the technique for fitting or modeling the flux from the nearby contaminating source, and the details of fitting the transit light curve model. Stevenson et al. (2013) demonstrated that by using two different transit modeling methods, small differences could be introduced in the spectrum; similarly, we have shown in Section 4.2 that the choice of the spectrum for the contaminating flux from WASP-12 BC can change the fitted transit depths by a factor comparable to the fitting uncertainty.

Remarkably, all the spectra show similar trends at wavelengths longer than 1.2 μm, with a high point at 1.225 μm and a broad peak from 1.325–1.575 μm. There are slight differences (at the 1–2σ level) for the bins at 1.425 and 1.525 μm, but the major disagreement is at the short-wavelength edge of the spectrum—the Swain et al. results show a steady rise at short wavelengths while the Stevenson et al. results show a upward spike in the shortest-wavelength bin (1.125 μm). In contrast, our spectrum shows a drop shortward of 1.2 μm. This region of the spectrum is particularly susceptible to the choice of the dilution factor for the contaminating star due to the wavelength shift of the spectrum (see Figure 7), and the edges of the spectrum also exhibit a steep gradient in flux due to the grism sensitivity which can lead to systematic trends if the spectrum drifts
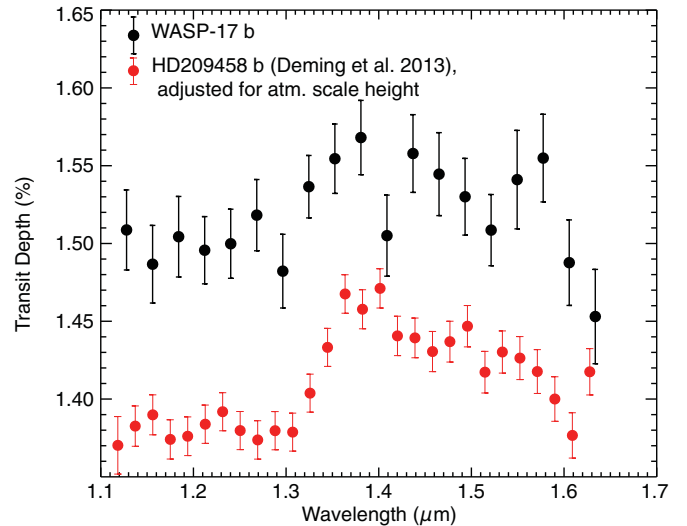
Figure 20. Results for WASP-17 b (black) compared with results for HD209458 b from Deming et al. (2013) in red. The spectrum for HD209458 b has been scaled to compensate for the difference in scale height for the two planets. The spectra match very well, suggesting commonality between the spectra for cooler, smaller planets.

(A color version of this figure is available in the online journal.)

over time (see Section 4.1). We therefore believe that a careful treatment of this spectral region is imperative. The downward slope of our final spectrum does not require any additional absorption from species such as TiH or CrH, as suggested by Swain et al. (2013). Our uncertainties are larger than those of Stevenson et al. (2013), but we believe the larger uncertainties are warranted based on the uncertainty in the contribution from WASP-12 BC.

5.2.2. WASP-17

There are no prior spectroscopic analyses of WASP-17 at H-band wavelengths, but we can compare our results with the recent WFC3 observations of HD209458 b by Deming et al. (2013). HD209458 b is similar in mass and temperature to WASP-17 b, but with a much smaller scale height—WASP-17 b has a scale height that is 3.4× larger than HD209458 b. In Figure 20 we plot our spectrum of WASP-17 b with the spectrum of HD209458 b from Deming et al. (2013), scaled up to compensate for the differences in scale height between the two planets; the spectra match very closely, though there is no evidence for the outlying peak at 1.575 μm in the spectrum of HD209458 b. The similarity between two cooler, lower-mass planets is especially notable considering that dissimilarity between the spectrum for WASP-17 b and the spectra for our other two targets, which are much hotter and more massive.

As stated earlier, we find that the models for WASP-17 b fit best when we include an additional linear slope in the models, which could be indicative of scattering due to haze or clouds. To determine the difference in the spectral slope compared with a cloudless model, we compared a linear fit to the two continuum regions (1.28 and 1.65 microns) for our standard haze-free model to the same fit for our best-fit hazy model including an additional linear spectral slope correction. We calculate a difference of −0.067% in the baseline transit depth across our band pass between the two models; this equates to a change in the radius of ~3000 km. If we assume that this spectral slope is due to a change in effective radius with wavelengths due to Rayleigh scattering, we can use Equation (4) from Lecavelier Des Etangs et al. (2008) to compare our spectral slope to similar
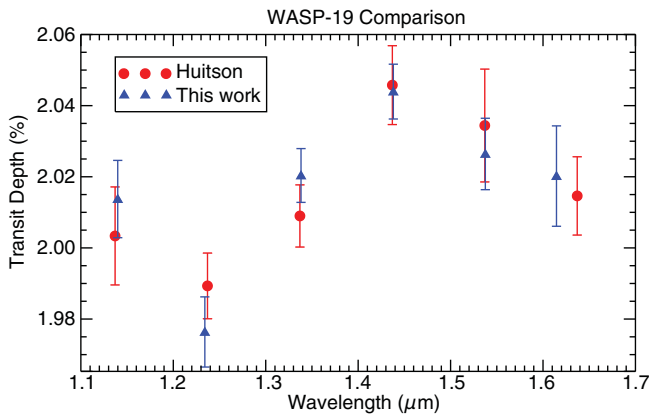
Figure 21. Results from Huitson et al. (2013) in red, with results from this work over plotted in blue, binned to the same size, with edges offset due to different choices of spectral binning. The spectra are largely consistent, but comparison with our smaller bin size suggests that the Huitson et al. (2013) may be missing statistically significant features in the spectrum.

(A color version of this figure is available in the online journal.)

results for the spectral slope of HD 189733 b across optical and IR wavelengths (Pont et al. 2008, 2013; Lecavelier Des Etangs et al. 2008). WASP-17 b is hotter than HD 189733 b by ~400 K, and the gravity is lower by a factor of ~7. Combining these factors leads to a change in altitude across our bandpass of ~4650 km −10× larger than for HD 189733 b, but only a factor of 1.5× larger than our best-fit value. Considering the lack of a detectable slope in the data for HD 209458 b, and the size of the uncertainty bars on our data, we consider this result highly speculative at this point; improved constraints through additional WFC3 observations and/or coincident radius measurements at other wavelengths will be necessary to examine this question in detail.

5.2.3. WASP-19

The current data set for WASP-19 was also recently analyzed by Huitson et al. (2013). Their published results utilized a bin size that is larger than ours by a factor of three (0.1 μm). They also subtracted the band-integrated residuals from each bin, but then used the divide-oot method on each bin separately and fit for transit depth and a linear trend. In Figure 21 we plot our results using a bin size matched to those of Huitson et al. (2013). The transit depths using larger bins are well matched to the Huitson et al. (2013) results, but as noted above, with smaller bins we see deviations from the smooth trend that appears to match the lower-resolution results. Huitson et al. (2013) state that they do not see any major differences beyond increased photon noise when using smaller bin sizes. However, the changes in our spectrum seem to be robust beyond a simple increase in photon noise. Bean et al. (2013) also presented a recent analysis of ground-based transit and occultation observations of WASP-19 at H-band wavelengths. Their results covered the region from 1.25–2.4 μm, with gaps near the peaks of the water features at 1.37 and 1.9 μm. The analysis of the transit observations yields only four broad bins in our wavelength region, similar in width and position to several of the wavelength bins used by Huitson et al. (2013) and generally consistent with both the Huitson et al. results and our own results for wide bins.

6. CONCLUSION

In this paper, we present our analysis of WFC3 observations of single transits for three exoplanets (WASP-12 b, WASP-

17 b, and WASP-19 b). We perform a careful analysis of the band-integrated time series for each target, revealing possible evidence of curvature in the out-of-transit data for WASP-12 and WASP-19 and evidence for a star spot in the light curve for WASP-19. We confirm that the repeating ramp-like or hook-like artifacts seen in a number of observations of exoplanets with WFC3 (which we call the "buffer-ramp") can be removed in the band-integrated light curve using the divide-oot method from Berta et al. (2012), but we develop an alternate method for removing the various systematic trends in the individual channels, or bins of multiple channels, that utilizes the residuals of the fit to the band-integrated light curve as well as measurements of the vertical and horizontal shift of the spectrum on the detector over time. We utilize a model selection strategy that relies on the BIC to determine the significance of fitting for individual systematic components, allowing us to identify trends due to changes in the amplitude of the buffer-ramp and the impact of spectral shifts on the flux in individual spectral bins. We present final transit spectra for each exoplanet using 0.027 μm channel bins, and argue that this is the optimal bin size for increasing S/N while avoiding any loss of spectral information that exceeds the photon-noise limit. When we use similar binning sizes to those used in previous analyses of the data for WASP-12 (Swain et al. 2013; Stevenson et al. 2013) and WASP-19 (Huitson et al. 2013), we can reproduce the earlier results to within uncertainties, except for the shortest-wavelength bin for WASP-12. This discrepancy may be due to treatment of data that falls on the steep spectral slope of the WFC3 sensitivity curve.

Our analysis demonstrates that precisions close to the photon-noise limit are possible for measurements of wavelength-dependent transit depths with WFC3 with the observation of only a single transit event even for relatively dim targets ($H > 10.2$). Measurements of the absolute transit depth are fundamentally limited by our ability to constrain parameters such as limb darkening and mid-transit time, and the phasing of HST orbits across the light curve has a significant impact on our final uncertainties in R_p/R_* for our band-integrated light curves. However, using our transit model including systematic trends, we show that the uncertainties for individual bins are not strongly affected by the light curve sampling and depend only on the number of photons acquired in transit and out-of-transit. Future observations of these targets that utilize the newly implemented spatial scan mode will allow for increased efficiency and improved sensitivity.

Comparison with theoretical models by Burrows et al. (2008) and Madhusudhan (2012) strongly suggest the presence of water absorption between 1.4 μm and 1.55 μm in WASP-17 b, and models with the inclusion of haze fit the data better than models without haze. For WASP-12 b and WASP-19 b the agreement with standard models including water absorption is not as clear. In particular, the spectral region beyond 1.45 μm shows increased absorption for all our targets beyond what is predicted from water-rich models; carbon-rich models provide a better match in this region, but significant discrepancies remain. We therefore believe that firm conclusions on atmospheric composition are impossible without more sensitive observations and/or a full analysis of multi-wavelength data at both optical and NIR wavelengths.

REFERENCES

Albrecht, S., Winn, J. N., Johnson, J. A., et al. 2012, ApJ, 757, 18
Anderson, D. R., Hellier, C., Gillon, M., et al. 2010, ApJ, 709, 159
Anderson, D. R., Smith, A. M. S., Lanotte, A. A., et al. 2011, MNRAS, 416, 2108
Anderson, D. R., Smith, A. M. S., Madhusudhan, N., et al. 2013, MNRAS, 430, 3422
Batygin, K., & Stevenson, D. J. 2010, ApJL, 714, L238
Bayliss, D. D. R., Winn, J. N., Mardling, R. A., & Sackett, P. D. 2010, ApJL, 722, L224
Bean, J. L., Desert, J.-M., Seifahrt, A., et al. 2013, ApJ, 771, 108
Bechter, E. B., Crepp, J. R., Ngo, H., et al. 2013, arXiv:1307.6857
Bergfors, C., Brandner, W., Daemgen, S., et al. 2013, MNRAS, 428, 182
Berta, Z. K., Charbonneau, D., Desert, J.-M., et al. 2012, ApJ, 747, 35
Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393
Bodenheimer, P., Lin, D. N. C., & Mardling, R. A. 2001, ApJ, 548, 466
Burrows, A., Budaj, J., & Hubeny, I. 2008, ApJ, 678, 1436
Burrows, A., Guillot, T., Hubbard, W. B., et al. 2000, ApJL, 534, L97
Burrows, A., Sudarsky, D., & Hubeny, I. 2006, ApJ, 650, 1140
Campo, C. J., Harrington, J., Hardy, R. A., et al. 2011, ApJ, 727, 125
Castelli, F., & Kurucz, R. L. 2004, arXiv:0405087
Charbonneau, D., Allen, L. E., Megeath, S. T., et al. 2005, ApJ, 626, 523
Charbonneau, D., Brown, T. M., Noyes, R. W., & Gilliland, R. L. 2002, ApJ, 568, 377
Claret, A., & Bloemen, S. 2011, A&A, 529, 75
Cowan, N. B., & Agol, E. 2011, ApJ, 726, 82
Cowan, N. B., Machalek, P., Croll, B., et al. 2012, ApJ, 747, 82
Croll, B., Lafreniere, D., Albert, L., et al. 2011, AJ, 141, 30
Crossfield, I. J. M., Barman, T., Hansen, B. M. S., Tanaka, I., & Kodama, T. 2012, ApJ, 760, 140
Deming, D., Seager, S., Richardson, L. J., & Harrington, J. 2005, Natur, 434, 740
Deming, D., Wilkins, A., McCullough, P., et al. 2012, BAAS, 219, 241.14
Deming, D., Wilkins, A., McCullough, P., et al. 2013, ApJ, 774, 95
Dressel, L. 2012, Wide Field Camera 3 Instrument Handbook, Version 5.0 (Baltimore, MD: STScI)
Fabrycky, D., & Tremaine, S. 2007, ApJ, 669, 1298
Ford, E. B. 2005, AJ, 129, 1706
Fortney, J. J., Lodders, K., Marley, M. S., & Freedman, R. S. 2008, ApJ, 678, 1419
Grillmair, C. J., Burrows, A., Charbonneau, D., et al. 2008, Natur, 456, 767
Grillmair, C. J., Charbonneau, D., Burrows, A., et al. 2007, ApJL, 658, L115
Guillot, T., & Showman, A. P. 2002, A&A, 385, 156
Hebb, L., Collier-Cameron, A., Loeillet, B., et al. 2009, ApJ, 693, 1920
Hellier, C., Anderson, D. R., Collier-Cameron, A., et al. 2011, ApJL, 730, L31
Howe, A. R., & Burrows, A. S. 2012, ApJ, 756, 176
Hubeny, I., Burrows, A., & Sudarsky, D. 2003, ApJ, 594, 1011
Huitson, C. M., Sing, D. K., Pont, F., et al. 2013, MNRAS, 434, 3252
Ibgui, L., & Burrows, A. 2009, ApJ, 700, 1921
Jackson, B., Greenberg, R., & Barnes, R. 2008, ApJ, 678, 1396
Knutson, H. A., Charbonneau, D., Allen, L. E., et al. 2007, Natur, 447, 183
Kümmel, M., Walsh, J. R., Pirzkal, N., Kuntschner, H., & Pasquali, A. 2009, PASP, 121, 59
Kuntschner, H., Bushouse, H., Walsh, J. R., & Kümmel, M. 2008, ST-ECF Instrument Science Report WFC3-2008-16 (Garching, Germany: Space Telescope European Coordinating Facility)
Lecavelier Des Etangs, A., Pont, F., Vidal-Madjar, A., & Sing, D. 2008, A&A, 481, L83
Lendl, M., Gillon, M., Queloz, D., et al. 2013, A&A, 552, 2
Liddle, A. R. 2004, MNRAS, 351, L49
Maciejewski, G., Dimitrov, D., Seeliger, M., et al. 2013, A&A, 551, 108
Madhusudhan, N. 2012, ApJ, 758, 36
Madhusudhan, N., Harrington, J., Stevenson, K. B., et al. 2011, Natur, 469, 64
Madhusudhan, N., & Seager, S. 2009, ApJ, 707, 24
Mandel, K., & Agol, E. 2002, ApJL, 580, L171
McCullough, P., & Deustua, S. 2008, Instrument Science Report WFC3 2008-33 (Baltimore, MD: STScI)
McCullough, P., & MacKenty, J. 2012, Instrument Science Report WFC3 2012-08 (Baltimore, MD: STScI)
Pont, F., Knutson, H., Gilliland, R. L., Moutou, C., & Charbonneau, D. 2008, MNRAS, 385, 109
Pont, F., Sing, D. K., Gibson, N. P., et al. 2013, MNRAS, 432, 2917
Rasio, F. A., & Ford, E. B. 1996, Sci, 274, 954
Richardson, L. J., Deming, D., Horning, K., Seager, S., & Harrington, J. 2007, Natur, 445, 892
Rogers, T. M., Lin, D. N. C., & Lau, H. H. B. 2012, ApJL, 758, L6
Schwarz, G. 1978, AnSta, 6, 461
Smith, R. M., Zavodny, M., Rahmer, G., & Bonati, M. 2008, Proc. SPIE, 7021, 70210J-1
Southworth, J., Hinse, T. C., Dominik, M., et al. 2012, MNRAS, 426, 1338
Spiegel, D. S., Silverio, K., & Burrows, A. 2009, ApJ, 699, 1487
Stevenson, K. B., Bean, J. L., Seifahrt, A., et al. 2013, arXiv:1305.1670
Swain, M., Deroo, P., Tinetti, G., et al. 2013, Icar, 225, 432
Swain, M. R., Tinetti, G., Vasisht, G., et al. 2009a, ApJ, 704, 1616
Swain, M. R., Vasisht, G., & Tinetti, G. 2008, Natur, 452, 329
Swain, M. R., Vasisht, G., Tinetti, G., et al. 2009b, ApJL, 690, L114
Tregloan-Reed, J., Southworth, J., & Tappert, C. 2013, MNRAS, 428, 3671
Triaud, A. H. M. J., Collier-Cameron, A., Queloz, D., et al. 2010, A&A, 524, 25
Weidenschilling, S. J., & Marzari, F. 1996, Natur, 384, 619