TRANSITS AND OCCULTATIONS OF HOT JUPITERS

by

Korey Haynes
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
In Partial fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Physics

Committee:

Dr. Michael Summers, Dissertation Director

Dr. Avram Mandell, Committee Member

Dr. Joseph Weingartner, Committee Member

Dr. Chi Yang, Interim Director, School of Physics, Astronomy, and Computational Sciences

Dr. Donna M. Fox, Associate Dean, Office of Student Affairs & Special Programs, College of Science

Dr. Peggy Agouris, Dean, College of Science

Date: 11-20-2014

Fall Semester 2014
George Mason University
Fairfax, VA

Transits and Occultations of Hot Jupiters

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

Korey Haynes
Master of Science
George Mason University, 2012
Bachelor of Arts
Macalester College, 2009

Director: Dr. Michael Summers, Professor
Department of School of Physics, Astronomy, and Computational Sciences

Fall Semester 2014
George Mason University
Fairfax, VA

Dedication

I dedicate this dissertation to Madeleine L'Engle, Diane Duane, Tamora Pierce, Bruce Coville, and Gene Roddenberry, who collaborated to raise someone who believed that heroes could be girls, stargazers, and scientists–maybe even all at the same time.

Acknowledgments

First and most importantly, I would never have made it to this final stage without Avi Mandell's guidance and support. He has been the adviser every grad student wishes for, helpful and patient and a wonderful teacher and collaborator. I've learned so much under his tutelage, and I can't fully express my gratitude not only for his support of my research, but equally for supporting my goals of anything other than research for a career.

I thank Mike Summers as well, for his guidance and generous gifts of time, advice, and even funding throughout my time at George Mason. Joe Weingartner has been an invaluable resource and coach since I arrived at grad school. He gave me the push I needed to go find a thesis project I would love, he's been a fantastic teacher and mentor, and I don't know what the department or I would do without him. And even though I didn't stick with her, Jessica Rosenberg was my first great graduate adviser, and I'm so glad she brought me in. Jess O'Connor, Derek Fertig, and Joey Rodriguez were the department for me my first year, and I miss you guys. Tanu Luke, I wish we'd had more time together; your pep talks were a lifeline, and they made more of a difference than you know.

I want to thank John Cannon, who set me on this wonderful path and provided the means by which I could transition from student to researcher. He made me believe I could do this whole crazy grad school notion, and I am unbelievably grateful for his inspiration, kindness, humor, and limitless support. I most certainly would not have made it here without him. To my other Mac professors, who laid the groundwork for everything that came after, and to my cohort, but especially Jillian, Matt, and Stephen, whose friendship was–and remains–even more valuable than our study sessions, even the ones augmented by Shish's hummus platters: thank you.

All of my outreach buddies, from my fellow Astrobites writers to my KICP Outreach Summer School cohort to my friends and mentors at Project Astro and the Smithsonian's Air and Space Museum: you kept me going when I might otherwise have lost interest and lost faith. You made this crazy ride more rewarding than I could have dreamed, and I am immeasurably grateful for your company, teaching, and example on this journey. My research probably went slower thanks to my activities with all of you, but more crucially, I'd never be taking the next step in my career without the skills and experience you've helped me gain. Thank you so much.

Evan Sinukoff, my first NASA research buddy, who wrote the math so I didn't have to: thanks. I'm grateful for your mad IDL skills, but more importantly, you made this thing fun. I still miss sharing a desk with you.

I thank my family, who have been proud of me even when they haven't had a clue what I was talking about. But I have to highlight my grandfather, who never stops clipping out articles from the Baltimore Sun's science section to send me, who never stops asking if I've found any new planets, and who might still be concerned that I'm moving on "to sell magazines." I know he's proud of me anyway.

Tom, my love, my rock: your support in this endeavor is such a small drop in the ocean of your love I don't even know where to start. Thank you for everything.

My friends have been tireless cheerleaders, and while I thank all of them from the bottom of my heart, I'd be remiss if I didn't single out Erika Nesvold. My best friend and role model for well over a decade, through the years she has also been my roommate, officemate, lunch buddy, commiserator, human whiteboard for talking over problems, and first and last stop for, "Do you know how to get IDL to...?" Not many people are lucky enough to have their best friend so literally at their side for the horrors and joys of the PhD journey, and I can't imagine how I would have made it through without her.

Finally, thanks to SPACS and NASA for financial and institutional support over the past four and a half years.

Table of Contents

List of Tables

List of Figures

Abstract

TRANSITS AND OCCULTATIONS OF HOT JUPITERS

Korey Haynes, PhD

George Mason University, 2014

Dissertation Director: Dr. Michael Summers

Since the first discovery of an extrasolar planet less than two decades ago, astronomers have learned how to measure not only the masses, radii, and orbital elements of a wide range exoplanets (far exceeding the parameters of our own solar system), but also their atmospheric temperatures and chemical compositions. Even with plentiful observations, many questions remain unanswered. Measuring atmospheric abundances based on observed absorption features can answer questions about carbon-to-oxygen (C/O) ratios, but many of the literature results rely on broadband photometry, where multiple absorption features become blended, thus complicating interpretation. Combining measurements across a long spectral baseline using multiple different instruments can be a powerful lever for studying the spectral energy distributions (SEDs) of exoplanets, but there is often a lack of consensus between observing teams and instruments. Some differences may be due to genuine temporal variations in the exoplanet atmospheres, while others are more likely due to differences in instrument characterization and data analysis. Resolved spectra of exoplanets, particularly in the infrared, where strong features due to H_2O, CO, CO_2, and CH_4 are expected, could break model degeneracies and answer

many questions about C/O ratios and pressure-temperature atmospheric structures. While not the first, Wide Field Camera 3 (WFC3) on the *Hubble Space Telescope* is the only current space-based opportunity to study spectrally resolved exoplanet atmospheres in the infrared. We focus on hot Jupiter type exoplanets, and use WFC3 (as well as ancillary data from *Spitzer* and ground based facilities) to try to break degeneracies between models, resolve past observing conflicts, and unambiguously determine these planets' atmospheric composition and structure. We discover unambiguous detections of water in exoplanet atmospheres, and the first spectroscopic evidence for a temperature inversion due to TiO in an exoplanet atmosphere.

Chapter 1: Background

1.1 Exoplanets

Exoplanets, while long theorized, were only discovered around a main sequence star in 1995 with the discovery of 51 Pegasi b by Mayor & Queloz (1995). A decade later, Charbonneau et al. (2005) and Deming et al. (2005) observed the first light from an exoplanet, and set the stage for a shift from detection to characterization. Now, almost a decade after that breakthrough, over 1800 exoplanets have been detected, with over 100 exoplanets bright enough for characterization observation. Many of these planets have multi-band photometry from both space and ground-based observatories, and some even resolved spectra.

Exoplanets are generally found via the transit or radial velocity method. The transit method relies on the star-planet system having a favorable geometric alignment with the observer, such that once during its orbit, the planet passes across its star from the observer's viewpoint, thus dimming the light from its star. The fraction of light blocked tells the observer the radius of the planet (larger planets block more light), and the duration of the transit tells the observer about the planet's orbital radius. A planet with a wider orbit will take longer to cross its star than a closer-in planet. The *Kepler* mission has found thousands of exoplanets by monitoring tens of thousands of stars and recording this periodic dimming of the star's light.

The radial velocity method measures the motion of a star. If a star is orbited by a planet, the planet will cause the star to orbit around the center of mass of the star-planet system, instead of rotating "in place" (any net motion of the star is removed in this investigation). The magnitude of the red- and blue-shifting of the star as it orbits tells observers the mass of the planet(s) in orbit around the star.

Other methods are occasionally used to detect planets, such as direct imaging (taking an image of the planet) and gravitational lensing (a planet causes an object behind the planet to brighten as it passes in front), but these methods are more rare.

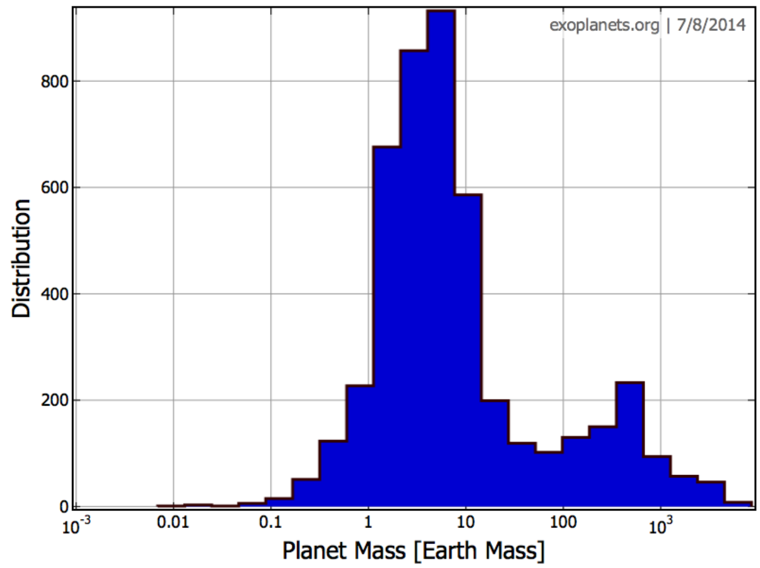

Figure 1.1 Histogram showing exoplanets detected as of July 9, 2014, arranged by mass. From exoplanets.org

Many different kinds of planets have been found through these methods. No known star system mirrors our own solar system, though our own observational limitations (we are insensitive to planets the mass of the solar system planets at their orbital separations) prevent us from finding such a system, should it exist. We are more sensitive to close-in planets for RV and transit searches, and very massive planets for direct imaging and gravitational lensing. The planets we have found include small rocky planets like Mercury, super-Earths and mini-Neptunes, gas giants that dwarf our own Jupiter, and even gas giants orbiting extremely close to their host stars (see Figures 1.1 and 1.2 for distributions of known exoplanets). These last types, usually called hot Jupiter exoplanets, are especially interesting because their large masses and radii, as well as their brightness compared to their host stars (because they are very hot), makes them prime targets for follow-up observations and characterization.

In this research, I will focus on the characterization of hot Jupiter atmospheres.

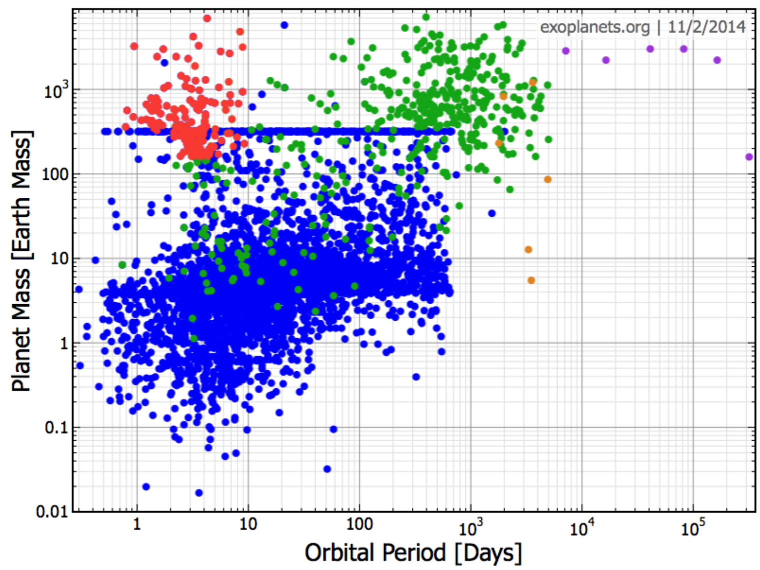

Figure 1.2 Exoplanets detected as of November 2, 2014. Planets are arranged by mass and orbital period. Planets detected by the transit method in blue, by radial velocity in green, by direct imaging in purple, and by gravitational microlensing in gold. Hot Jupiter planets have been over plotted in red. From exoplanets.org.

1.2 Methods for Observing Exoplanet Atmospheres

At present there are three proven methods of observing exoplanet atmospheres: direct imaging, radial velocity analysis, and transit analysis.

The first of these methods is to directly image the planet in question. Unfortunately, even the nearest stars have prohibitively short angular separations from most of their planets, making it impossible for even the most advanced telescopes to resolve both planet and star. Planets orbiting far from their stars are usually too dim for telescopes to detect at all, unless the planets are extremely hot and massive. Since planets far from their stars are heated mainly by formation processes and not stellar irradiation, this creates a selection bias for young planets that are $\gtrsim 1$ Jupiter masses. A few planets are both far enough from their star and hot enough to directly observe with modern instrumentation utilizing large collection areas and adaptive optics. Astronomers have observed several of these systems, including β Pictoris b (Lagrange et al. (2009), also see Figure 1.3), Fomalhaut b (Kalas et al. (2008)), and also the four-planet system around HR 8799 (Marois et al. (2008)). Observing HR 8799 b at H and K bands (roughly 1.5 - 2.4 μm) (Barman et al. (2011), Konopacky et al. (2013)) has yielded information about the H_2O, CO_2, and CH_4 abundances of the planet's atmosphere, implying a C/O ratio > 1 and thick clouds. Unfortunately, the sample of planets for which this approach is possible remains small, though it will grow with future observational platforms (see Madhusudhan et al. (2014) for an overview of upcoming ground- and space-based direct imaging opportunities).

The second method, that of using radial velocity (RV) measurements to study a planet's atmosphere, is a very recent development. Identifying changes in the RV of the star due to the motion induced by a planet in orbit around the star is a common method for discovering exoplanets. Signals from spectral features either emitted or transmitted by the planet's atmosphere will have a different radial velocity than the star and can be isolated to determine both the mass of the planet and the abundance of the atmospheric component producing the observed spectral feature. Astronomers have been using such methods to study double-line eclipsing binary stars for decades, and a few exoplanets exist that meet the requirements for such studies. Snellen et al. (2010) observed HD 209458 b using the cryogenic high-resolution infrared echelon spectrograph (CRIRES, Kaeufl et al. (2004)) on the Very Large Telescope (VLT) at the European Southern Observatory (ESO) to measure a spectral signal due to CO that suggested enhanced carbon abundances in the upper atmosphere. Brogi et al. (2012) observed the non-transiting planet τ Boötis b, also using CRIRES. They were able to detect a strong absorption feature due to CO (see Figure 1.4), and concluded that the atmosphere lacks a temperature inversion. These results were confirmed by Rodler et al. (2012) using the same instrument and methods. Rodler et al. (2013) used the near infrared spectrometer (NIRSPEC, McLean et al. (1998))

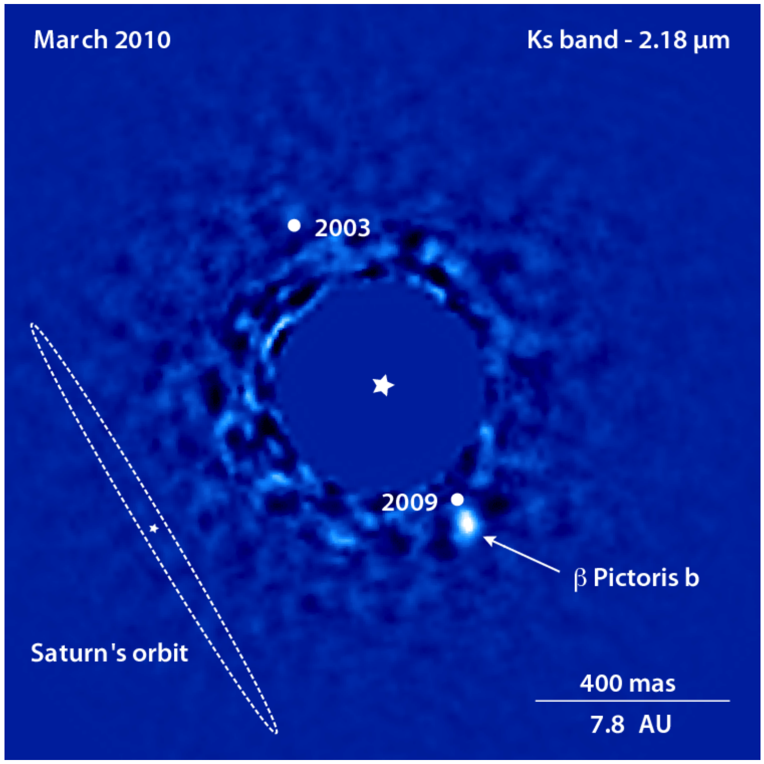

Figure 1.3 VLT/NaCo imaging of the β Pictoris system at 2.18 μm. The center of the system, containing the star, has been masked to allow for detection of the planetary companion. Superimposed on the image are the locations of the star, the planet in 2003, and the orbit of Saturn, for size reference. Adapted from Bonnefoy et al. (2011).

on Keck (R $\sim$ 20,000) to observe the presence of CO and indications for a lack of thermal inversion in HD 189733 b. These results were confirmed by de Kok et al. (2013), who used CRIRES to observe CO in the dayside of the planet and find upper limits for the presence of H_2O, CO_2, and CH_4. These are powerful observations, but they are limited by the fact that very few exoplanets are amenable to these types of high resolution spectroscopy.

For the foreseeable future, transiting exoplanets offer the best opportunity to observe exoplanet atmospheres. A transit is when an exoplanet passes in front of its host star, blocking some fraction of the star's light. This method does require the observer to be positioned such that the transit is visible. A planet that orbits perpendicular to the observer's line of sight will not appear to transit. If the planet is on a circular orbit, a

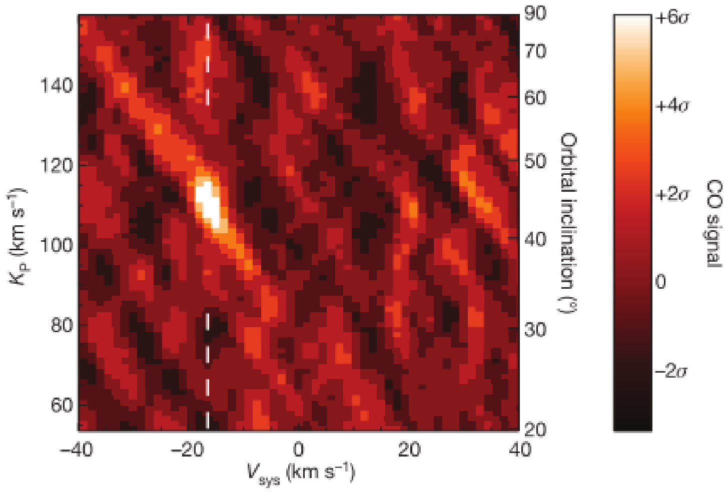

Figure 1.4 CO signal in the dayside spectrum of the exoplanet τ Boötis b. Shown is a color scale plot of the CO signal as a function of heliocentric (systemic) velocity (V_{sys}) on the x axis, and the maximum radial velocity of the planet, K_P, on the left-hand y axis. Lighter colors indicate CO in absorption. A clear signal at the 6.2 σ level is visible at the systemic velocity of τ Boötis (216.4 kms^{-1}), as indicated by the vertical dashed line, for a maximum orbital radial velocity of the planet of $K_P = 110.0 \pm 3.2$ km s^{-1}. The signal is obtained by cross-correlating a template spectrum of CO lines with the CRIRES/VLT spectra, which were each shifted in wavelength using the planet's ephemeris assuming a particular value of K_P. This is to compensate for the changing Doppler effect caused by the change in the planet's radial velocity over the large range in phase. Adapted from Brogi et al. (2012).

similar but smaller dip in flux will be seen when the planet passes behind its host star and the light from the planet is hidden from view (elliptical orbits may or may not offer both a transit and an eclipse to the observer). This is called occultation or secondary eclipse. See Figure 1.5 for an illustration of this effect.

Planets with larger radii will cause deeper transits than smaller planets, and planets on longer orbits will also take longer to transit their star.

In planets with extended atmospheres, some light from the star will pass through the limb of the planet as it transits. This light carries the imprint of the molecules that make up the planet's atmosphere. At wavelengths where these molecules absorb, the planet's atmosphere is opaque, causing a deeper transit than at wavelengths free from absorption

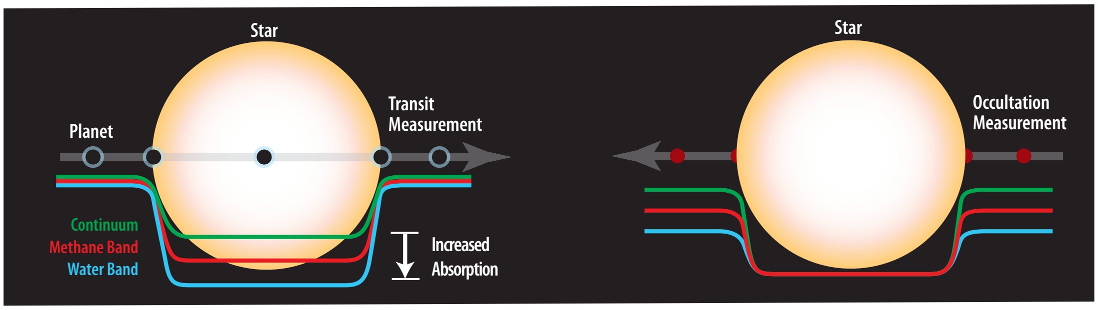

Figure 1.5 Left: Diagram depicting a planet transiting its host star. Transits viewed at wavelengths corresponding to features due to, for example, water or methane will show a deeper transit depth than over continuum regions. Light collected during transit carries an imprint of signatures from the limb of the planet's atmosphere it passes through on its way to the detector. Right: Diagram depicting an occultation of the same planet. In this case, observations are sensitive to the thermal emission from the planet itself, as we measure the light from star+planet on either side of the occultation, and the star alone during occultation.

features. This change in transit depth yields a transmission spectrum that can be used to determine the composition and temperature structure of the planet's atmosphere. This technique primarily probes regions of low pressure (i.e., high altitude) at the day-night terminator. Sodium absorption was one of the first spectral signals to be observed in an exoplanet atmosphere (Charbonneau et al. (2002)), but in the infrared the strongest signals are due to molecular absorbers such as H_2O, CH_4, CO, and CO_2.

The difference in transit depth across a spectral feature can be subtle, but this is exactly the change we wish to observe in order to determine the spectral composition of atmospheres via the transit or eclipse method. A few basic calculations can illustrate the amplitude of the change in depth one can expect. In the most basic case where no clouds or hazes affect the measurements, the amplitude of the absorption features are proportional to the atmospheric scale height, which is calculated as:

$$H = \frac{kT}{\mu g}. \tag{1.1}$$

Here, k is Boltzmann's constant, T is the temperature of the planet's atmosphere, μ is the mean molecular mass of the atmosphere, and g is the surface gravity (Madhusudhan et al. (2014)). Since the atmospheric composition is usually unknown, hot Jupiters are typically assumed to have roughly solar composition (C/O $\approx$ 0.4).

The expected change in transit depth is estimated to be 5-10 scale heights above the planet radius (R_p), is measured relative to the stellar radius (R_*), and is calculated as:

$$\delta_{depth} \simeq \left(\frac{R_p + 10H}{R_*}\right)^2 - \left(\frac{R_p}{R_*}\right)^2 \tag{1.2}$$

(Madhusudhan et al. (2014)).

WASP-12 b is a classic example of a hot Jupiter and one of the planets observed for this research study. As an example, WASP-12 b has an average temperature of 2525 K and a surface gravity of 975 cm/s^2. Assuming an H_2 composition atmosphere, this yields a scale height of 1058 km. The star WASP-12 has a radius of 1.63 $R_\odot$, and its planet has a radius of 1.79 R_{Jup}, so one could expect a nominal transit depth of approximately 1.2%, with additional absorption on the order of 0.2%, using Equation 2.

During occultation, the planet passes behind its host star, and the difference in brightness from star+planet to the star alone yields information about the thermal emission from the planet's dayside. The secondary eclipse depth can be estimated as:

$$depth = \left(\frac{R_p}{R_*}\right)^2 \left(\frac{T_p}{T_*}\right) \tag{1.3}$$

where T_p and T_* are planetary and stellar temperatures, respectively. Assuming the planet radiates as a blackbody, the equilibrium temperature can be estimated from the Stefan-Boltzman law as:

$$T_p = \left(\frac{(1-A)L_*}{16\pi\sigma d^2}\right)^{1/4} \simeq T_*\sqrt{\frac{R_*}{2d}}, \tag{1.4}$$

where A is the planet's Bond albedo (the fraction of incident light reflected at all wavelengths), σ is the Stefan-Boltzmann constant, and d is the distance between the planet and its star. In most cases, the amount of light reflected by the planet is negligible, and the equation simplifies to the right-hand side of the equation. However, if the planet lacks circulation and instead absorbs and emits only on its dayside, then T_p increases by a factor of $2^{1/4}$.

Spectral features can theoretically be observed from occultations. However, due to the lower contrast compared to transits, resolving spectral features during occultation is more difficult. For current observational limits, resolved spectral observations remain unattainable. Most occultation measurements are taken in the infrared, because the flux contrast between the planet and star is most favorable at those wavelengths, where the star is relatively faint and the planet relatively bright. However, occultation measurements taken in the visible can constrain the planetary albedo by measuring the amount of light reflected or scattered off the top of the planet's atmosphere. Regardless, even broadband measurements across a range of wavelengths can probe the blackbody curve of an exoplanet, yielding information about the energy budget, and perhaps still revealing information about broad molecular features in emission.

1.3 Theories of Hot Jupiter Atmospheres

The observational spectra that comprise the results of this research must be compared against spectra of model atmospheres in order to be interpreted. In the following section I summarize the crucial features of the models that can be constrained observationally.

1.3.1 Chemical Disequilibrium and C/O Ratios

In the simplest physical case, one should assume that a planetary atmosphere is in chemical equilibrium. This means that for a given temperature, pressure, and set of elemental abundances, the components of the atmosphere's composition are determined by minimizing the Gibbs free energy of the system. The temperature of the planet is an observable, and the pressure can be calculated at each level of the atmosphere. Solar composition is typically assumed, which means an atmosphere dominated by H and He, with a C/O ratio of 0.5.

In this case, water forms very favorably, and is the largest reservoir of oxygen (See Figure 1.6). Carbon is mostly stored in CO at high temperatures or CH_4 at $T \lesssim 1300$ K. Other molecules such as CO_2, NH_3, and other hydrocarbons may also play an important role, depending on the temperature (Madhusudhan et al. (2014), see also Figure 1.7). However, in the case of non-solar abundances, where the C/O ratio has been greatly enhanced (C/O $\gtrsim 1$), water is depleted, and carbon species dominate the spectrum. A planet with a C/O ratio that differs from its host star would have implications for formation process, particularly where in the disk it initially formed.

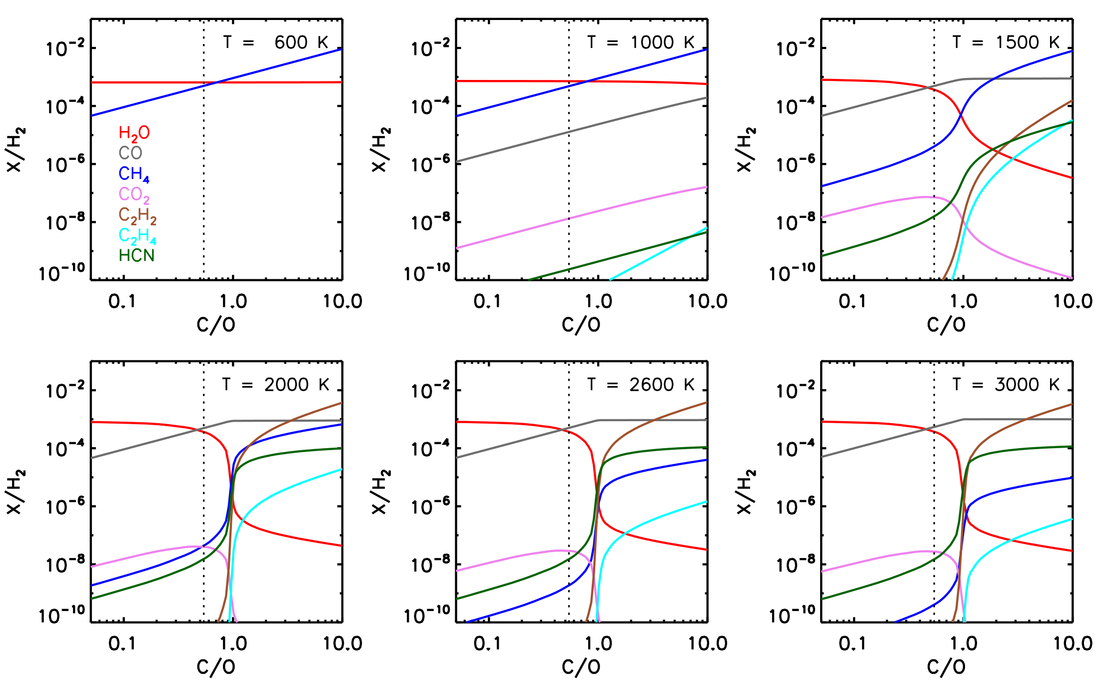

Figure 1.6 Molecular mixing ratios for common atmospheric constituents as a function of C/O for various temperatures. Figure from Madhusudhan & Seager (2011).

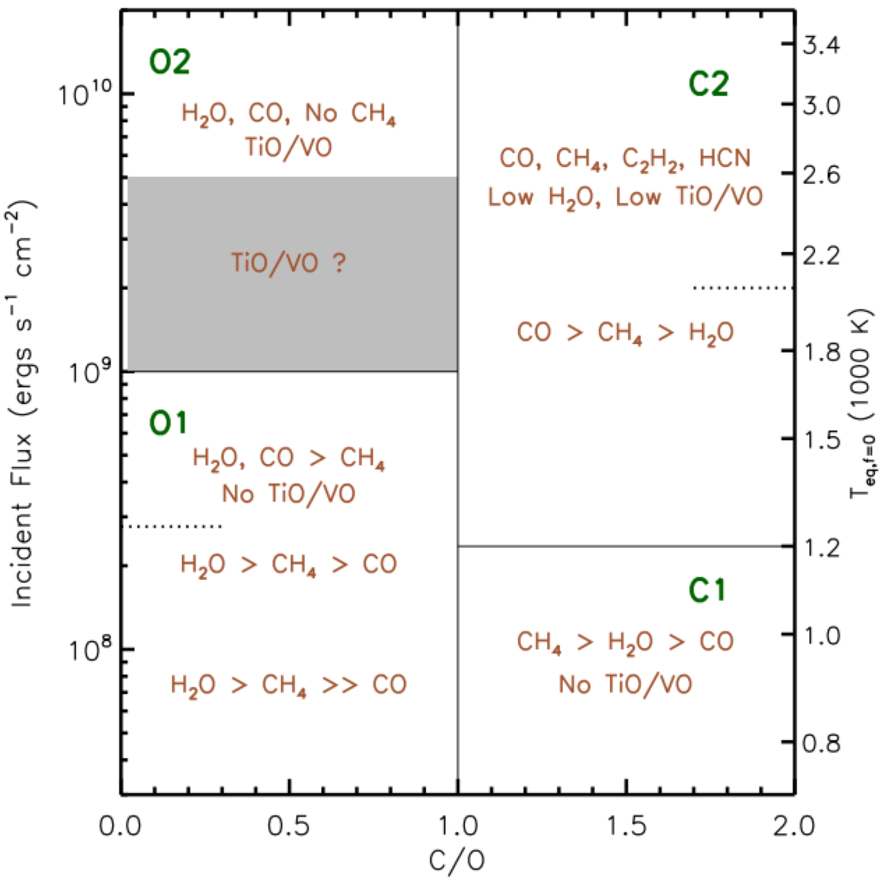

Figure 1.7 Overview of strongest features observed depending on C/O ratio and incident stellar flux. Figure from Madhusudhan & Seager (2011).

In cases of non-equilibrium chemistry, different reservoirs for carbon and oxygen may emerge. Non-equilibrium chemistry is common in solar system planets, and can be caused by photochemistry, turbulent mixing, or gravitational settling. Mixing in particular tends to be more commonly observed in lower temperature atmospheres ($T \lesssim 1300$ K). If it is common in higher temperature atmospheres, it must happen at high pressure regions below the level probed by observations. The temperatures where non-equilibrium chemistry might be expected is the same where CH_4 is expected to be the main carbon-bearing species. This means that CO might be seen at higher abundances than would be expected under equilibrium conditions, as it is dredged up from hotter regions of the atmosphere.

The conditions to cause turbulence and vertical mixing in the atmosphere are dependent on atmospheric dynamics. At the most basic level, circulation on hot Jupiters is driven by the strong temperature gradient between the day and night sides of the planet, particularly if the planet is tidally locked, a common occurrence for close-in planets. The planet's rotation speed dictates the formation of Hadley cells, which convect longitudinally. Slowly rotating planets develop almost global Hadley cells, while more rapidly rotating planets may have many bands of winds (Jupiter itself has $\sim$20).

1.3.2 Temperature Inversions

In the simplest of toy models, a planet's temperature should decrease with increasing altitude. However, with the presence of strongly absorbing molecules in the upper layers of a planet's atmosphere, a temperature inversion may form, where the temperature increases with altitude at certain layers. See Figure 1.8 for a model that supports a temperature inversion. Such an inversion layer is seen in the Earth's own stratosphere due to the ozone layer, as well as in the atmospheres of the giant solar system planets. Temperature inversions have been predicted in some hot Jupiters, commonly attributed to absorption by TiO or VO (Hubeny et al. (2003), Fortney et al. (2006), Fortney et al. (2008), Burrows et al. (2008)), as has been observed in brown dwarfs. However, many studies have explored the difficulty of retaining gas-phase TiO in the upper atmosphere of exoplanets. Cold layers lower in the atmosphere or on the planet's nightside would cause the Ti to condense out, leaving it absent from the upper layers when parcels return to warmer regions of the atmosphere. Perhaps the most straightforward problem is that TiO is a heavy molecule, and requires significant amounts of vertical mixing to keep it aloft in the upper atmosphere. Furthermore, Knutson et al. (2010) has explored the relationship between stellar activity and the presence of temperature inversions, theorizing that high UV flux can dissociate TiO. In high C/O ratio atmospheres, it is also possible that the oxygen reservoirs are depleted enough that TiO and VO will not form at all. While TiO and VO are the most commonly modeled sources of hot Jupiter temperature inversions, with the lack of any observational evidence for their existence, it is equally plausible to model some ad hoc absorber, as has been done by Burrows et al. (2008). However, there are at present no alternative molecules that absorb in the necessary wavelength regions and can also exist at the high temperatures of most hot Jupiters.

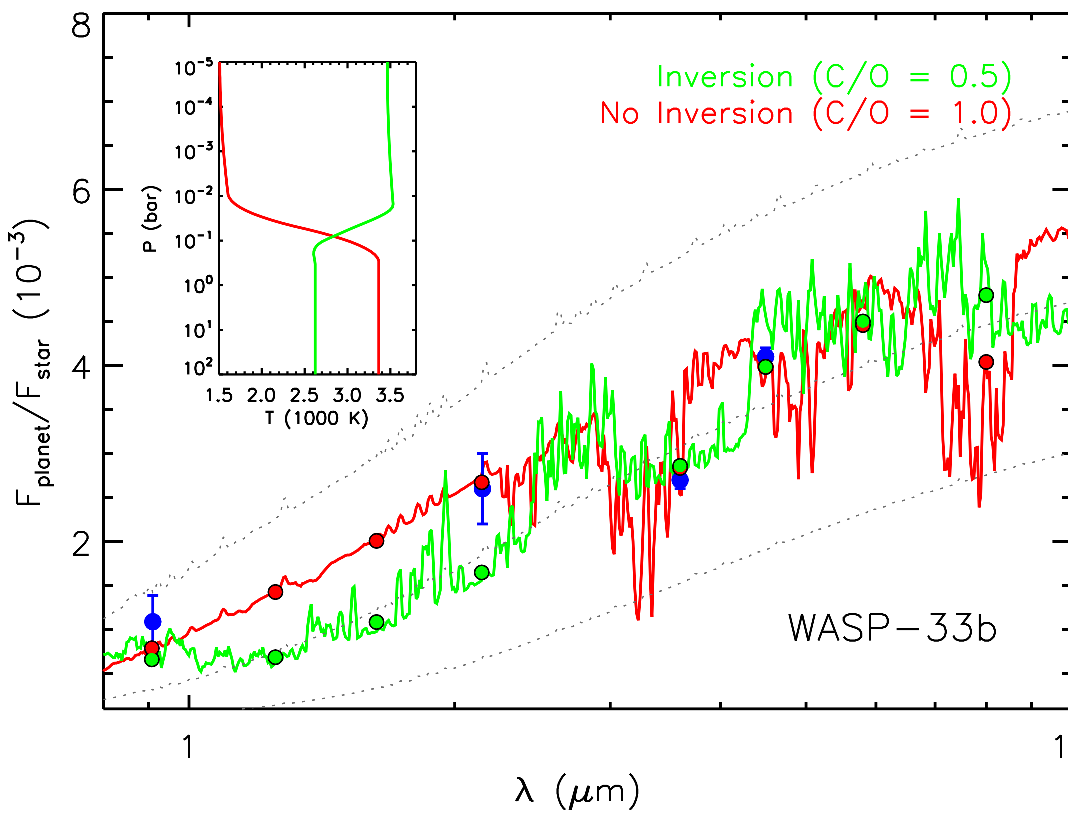

Figure 1.8 Two models for the hot Jupiter WASP-33 b, one of the planets studied for this research. This shows a model with and without a temperature inversion, depending on the chemical composition of the atmosphere. Figure from Madhusudhan & Seager (2011).

1.4 Previous Observational Studies of Hot Jupiter Atmospheres

In this section I review milestones and the most relevant observational studies in the field of exoplanet atmospheres, with a particular focus on hot Jupiters.

1.4.1 Early and Parallel Spectral Work

The first observations of exoplanet atmospheres were largely broadband, photometric points. The *Spitzer Space Telescope* (Werner et al. (2004)) and *Hubble Space Telescope* were heavily utilized in the search to observe and understand emission and transmission from these atmospheres. Once the techniques had been proven to work, studies accumulated quickly. An exhaustive review is beyond the capabilities of this work, but a sample of noteworthy observational results is presented here. Specifically, I point out the disagreements and revisions common to this field, since different instruments and observing teams often come to different conclusions about the same targets.

The first detections of light from an exoplanet atmosphere came from Deming et al. (2005), and Seager et al. (2005), who used the *Spitzer's* Multiband Imaging Photometer for *Spitzer* (MIPS, Rieke et al. (2004)) at 24 μm to observe the hot Jupiter HD 209458 b during occultation. This marked the first time light had been observed directly emitted from an exoplanet, and the target was fitting, since HD 209458 b was also the first exoplanet observed to transit (Charbonneau et al. (2000)). Both observations reported a non-detection of water or CO. Near the same time, Charbonneau et al. (2005) used *Spitzer's* Infrared Array Camera (IRAC, Fazio et al. (2004)) at 4.5 and 8.0 μm to observe the hot Jupiter TrES-1 b, also during occultation. They found a modest departure from a blackbody curve, implying that some additional opacity is present in the 4.5 μm band. Barman (2007) found evidence for water and CO in the atmosphere of HD 209458 b by observing the planet's transit with *HST* STIS. However Richardson et al. (2007) observed HD 209458 b in eclipse using *Spitzer* IRS and found evidence for silicate clouds, but a poor model fit for water. Possible reasons for this lack of consensus were offered as differences in the region of atmosphere probed (dayside during the occultation observations, and the day-night terminator for the transit), or even the large difference in time between observations. Knutson et al. (2008) observed the same planet in eclipse with all four IRAC bands, and found results inconsistent with a cloudless atmosphere. Instead, they suggest a temperature inversion that would cause water emission at the 4.5 and 5.8 μm bands.

A water detection was also claimed by Tinetti et al. (2007) for HD 189733 b using *Spitzer* IRAC 3.6, 5.8 and 8.0 μm bands to observe the planet during occultation. Désert et al. (2009) later observed HD 189733 b using updated IRAC 4.5 and 8.0 μm band observations, and re-analyzing the Tinetti et al. (2007) 3.6 and 5.8 μm data, but did not find evidence for water absorption. They did find evidence for absorption due to CO, perhaps implying a high C/O ratio. Grillmair et al. (2008) observed HD 189733 b during occultation with *Spitzer* IRS and found strong evidence for water. This result was confirmed by Barman (2008) by comparing a selection of available data and models. Machalek et al. (2008) observed XO-1 b during occultation using all four IRAC bands, and found evidence for a thermal inversion and possible water detection.

HST NICMOS provided a step forward, allowing for spectroscopic observations from a space-based facility. Its wavelength range from 1.8 - 2.3 μm covered bands including H_2O, CH_4 and CO_2. Using this instrument, Swain et al. (2008) claimed a detection of both water and methane in the atmosphere of HD 189733 b, observed during transit. The same group examined HD 189733 b's dayside emission spectrum in (Swain et al. (2009a)), to find H_2O, CO, and CO_2. Similar findings were reported of H_2O, CH_4, and CO_2 for HD 209458 b's emission spectrum (Swain et al. (2009b)), and of H_2O, CH_4, CO_2, and CO for the transmission spectrum of XO-1 b (Tinetti et al. (2010)). Pont et al. (2009) observed GJ 436 b, a hot Neptune, during transit, but report no significant molecular signals. However, the NICMOS detections as a whole were challenged by Gibson et al. (2011), who argued that the NICMOS results suffered from inadequately corrected instrument systematics and should have significantly higher uncertainties, which in most cases removed the evidence for molecular features in the data. These conclusions for HD189733 b were later revised in Gibson et al. (2012a), which proposed instead that the results were robust, although larger uncertainties were required to accurately represent the telescope's capabilities. These results were largely confirmed by Waldmann et al. (2013) and Swain et al. (2014). Burke et al. (2010) re-analyzed the XO-1 b NICMOS data and concluded that certain instrument systematics were unaccounted for in the Tinetti et al. (2010) study, but could be corrected for through additional correction procedures. They did not offer a wavelength dependent analysis, however. Crouzet et al. (2012) observed XO-2 b and also offered a re-analysis of XO-1 b, and similarly conclude that the instrument systematics are strong enough that any detection of water cannot be called significant. This controversy highlighted the need in the field not only for careful analysis, but also the importance of reporting uncertainties accurately, and the strong impact that instrument systematics can have on data interpretation.

Until recently, ground-based measurements were mostly constrained to photometric observations. Multiple reference stars are generally needed in order to correct for variations in the Earth's atmosphere during the period of observation, and until recently, multi-object spectrographs did not exist on the telescopes that had the capability to do exoplanet work. However the photometric work accomplished by Rogers et al. (2009), Gillon et al. (2009), de Mooij & Snellen (2009), Sing & López-Morales (2009), Croll et al. (2011), Smith et al. (2011), and Deming et al. (2012b), among many others, provides valuable information in bands often inaccessible by current space-based observatories. Furthermore, these bands can provide anchors at very different wavelengths, affording a long spectral baseline with which to study the exoplanet atmospheres.

Ground based spectral work (Mandell et al. (2011), Bean et al. (2013), among others) can offer complementary information to the work referenced so far, but instrument capabilities are as yet limited.

Much of the forward progress in characterization of exoplanet atmospheres has continued to rely upon *Spitzer* and in particular the four IRAC bands, which were reduced to two (the 3.6 and 4.5 μm bands) when *Spitzer's* cryogen reserves were depleted.

Combined with developments in modeling, particularly retrieval methods (Madhusudhan & Seager (2009), Madhusudhan & Seager (2010), Madhusudhan & Seager (2011), Lee et al. (2012), Line et al. (2012)), these four bands allowed constraints on temperature inversions and the H_2O, CH_4, CO_2, and CO contents on a variety of planetary atmospheres. These results were also bolstered by developments from the ground, allowing long spectral baselines to leverage information even when resolved spectra were unobtainable or questionable. Even with resolved spectra, these long baselines of information are still vital to understanding the atmospheres of exoplanets across ranges of pressure and temperature, and helping to break degeneracies between models. While the remainder of this work will focus on medium resolution spectra, we attempt to leverage observations taken at other wavelengths in order to have the most complete possible understanding of the exoplanets at hand.

1.4.2 Exoplanet Spectroscopy with WFC3

The Wide Field Camera 3 (WFC3, Dressel (2012)) on *HST* provides the bulk of the observations for this work. WFC3 can provide spectra of planets during transit and occultation, covering a wavelength range from 1.1-1.7 μm at a spectral resolution of 130, neatly bracketing a broad absorption feature due to water 1.4 μm. Another water feature exists at 1.15 μm, as well as features due to CO and CH_4. These water features probes the overall C/O ratio of the planet's atmosphere during transit. This spectrally resolved wavelength region also provides a near-infrared anchor when combined with longer wavelength photometric observations during occultation that give insight into the thermal emission from the planet.

The first exoplanet observations using WFC3 were accomplished by Gibson et al. (2012b), who observed HD 189733 b during transit. Due to saturation and non-linearity issues across the peak of the spectrum, only two spectral bins representing the blue and red ends of the spectrum were extracted from each of two visits. They reported tentative evidence

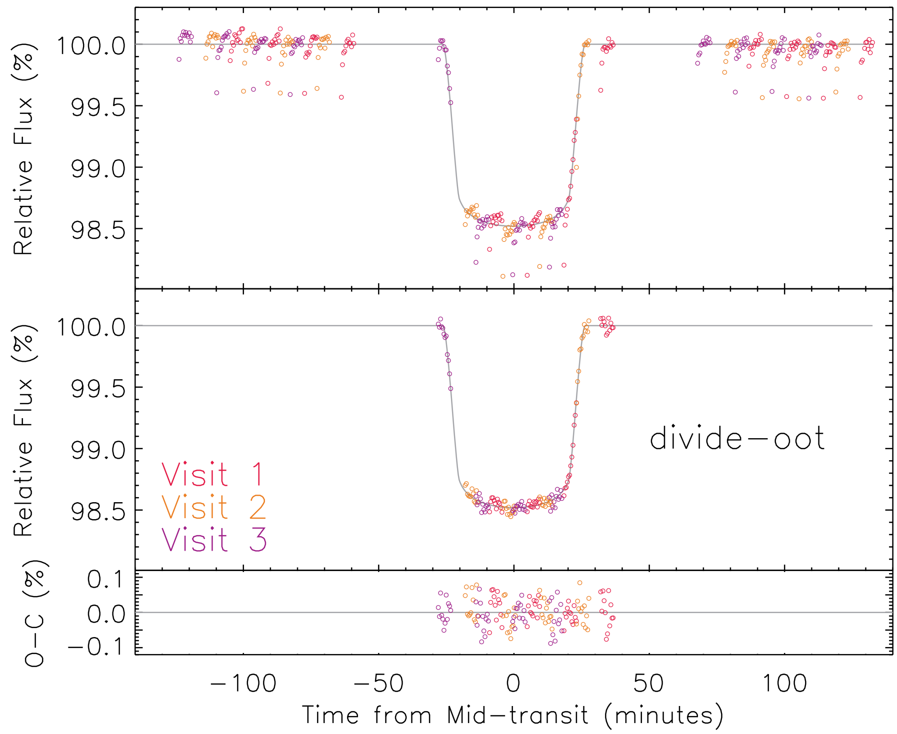

Figure 1.9 Lightcurves for three separate visits of GJ 1214, before (top) and after (middle) averaging the out-of-transit data and dividing it from all the data, in order to remove systematic trends. Bottom panel shows flat residuals. Figure from Berta et al. (2012).

for a hazy transmission spectrum. Soon after, Berta et al. (2012) published results of the super-Earth GJ 1214 b, using the full extent of the WFC3 spectral coverage. They characterized "ramp" features associated with the timing of detector buffer dumps, and further provided strategies for removing these features from the data (see Figure 1.9). These ramp features have been commonly observed in later campaigns, though it is possible to avoid them in certain observing modes (see Swain et al. (2013), Deming et al. (2013)). Berta et al. (2012) found a flat spectrum for GJ 1214 b, which at the time could be attributed to either a high mean molecular weight atmosphere or high-altitude clouds (see Figure 1.10). Subsequent observations published in Kreidberg et al. (2014) conclusively ruled out a cloud-free atmosphere (see Figure 1.11).

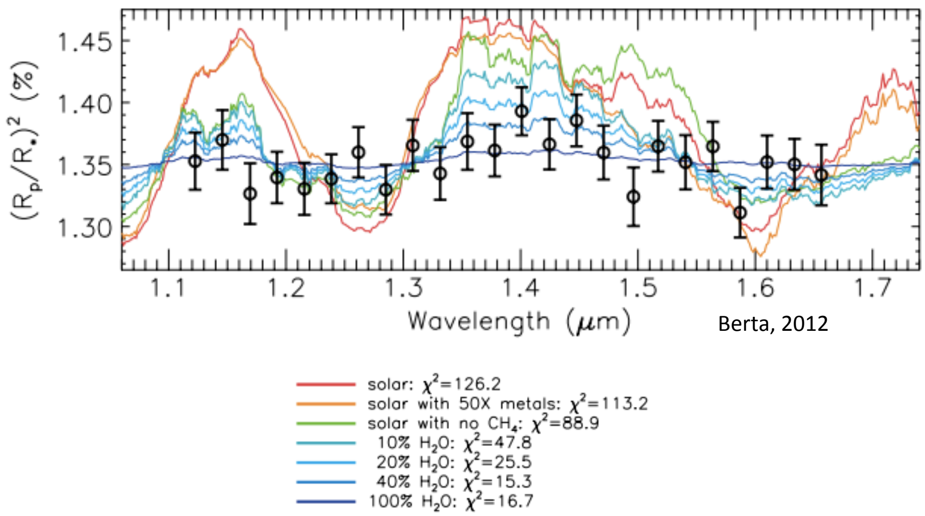

Figure 1.10 Various atmospheric models are shown in comparison to the transmission spectrum of GJ 1214. The closest match is a 100% water atmosphere, but results could be explained equally well by a cloud layer flattening spectral signatures. Figure from Berta et al. (2012).

There have been many successes already using WFC3 to observe hot Jupiter atmospheres. Swain et al. (2013) observed WASP-12 b, but did not find evidence for a C/O ratio >1, as had been identified by Madhusudhan et al. (2011b). They did undertake a study of WFC3 data available at the time in order to understand the ramp effect first observed by Berta et al. (2012) and other systematics of the WFC3 instrument, and offered guidance for future observers.

Huitson et al. (2013) analyzed observations of WASP-19 b (combining the WFC3 measurements with observations from STIS), and identified water in the atmosphere, as well as ruling out the presence of TiO features in the transmission spectrum. Deming et al. (2013) identified water in the atmospheres of XO-1 b and HD 209458 b by observing them during transit, though the signatures are weaker than expected, likely due to clouds or hazes (see Figures 1.12 and 1.13). Similarly, Wilkins et al. (2014) identified no significant absorption or emission features in the occultation measurements of CoRoT-2 b,

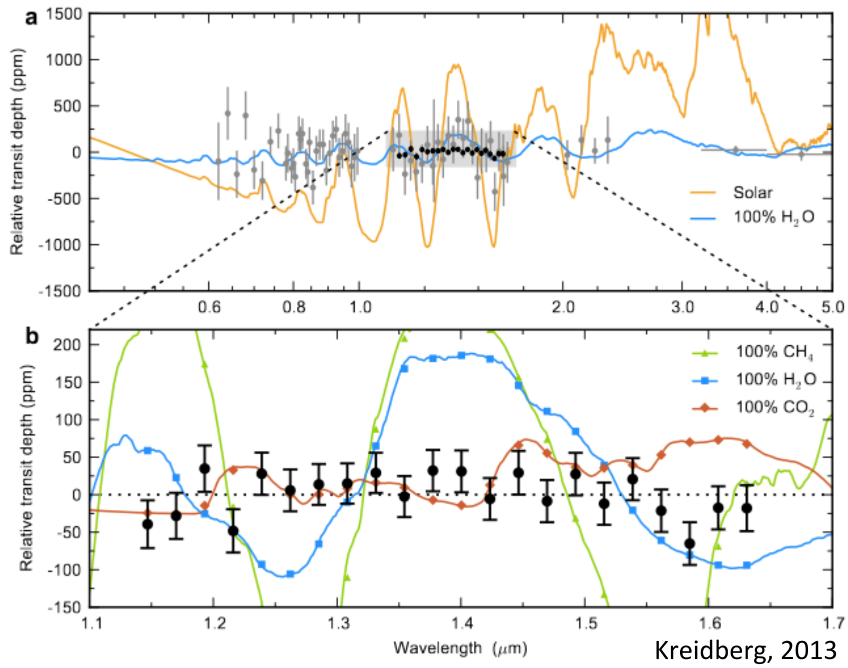

Figure 1.11 More recent spectrum of GJ 1214 b based on 15 transits of its host star. This spectrum clearly shows a featureless atmosphere explained only by a cloud or haze layer. Figure from Kreidberg et al. (2014).

and concluded that clouds or hazes must be present in the atmosphere of the planet. They also corrected wavelength solution coefficients from those given by STScI for the WFC3 instrument. Ranjan et al. (2014) observed a variety of transits and occultations, finding featureless spectra for all planets observed (TrES-2b, TrES-3b, TrES-4b, CoRoT-1b, and WASP-4).

WFC3 has therefore already been proven as an effective tool in the effort to characterize exoplanet atmospheres, and it is the primary instrument used in the research I propose here.

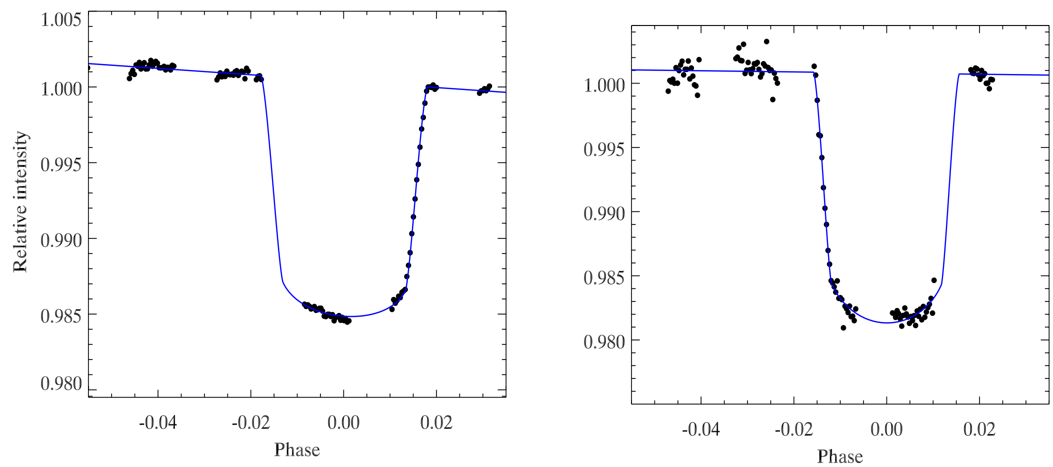

Figure 1.12 Lightcurves for HD 209458 b (left) and XO-1 b (right). These data were taken in spatial scan mode, and do not display the systematic ramp or hook pattern observed in stare mode data. Figure from Deming et al. (2013).

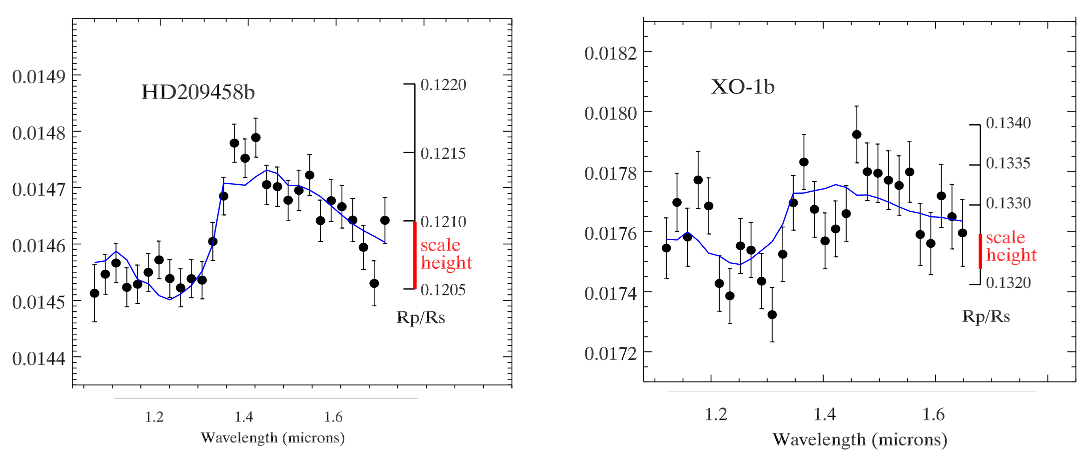

Figure 1.13 Transmission spectra for HD 209458 b (left) and XO-1 b (right). Models from Burrows et al. (2008), with amplitude of features scaled to match observations. The fitted amplitude of HD 209458 b's absorption is 0.57 of the modeled value, and for XO-1 b, the fitted amplitude is 0.84 of the model. Figure from Deming et al. (2013).

Chapter 2: Statement of Problem and Goals for This Work

Exoplanet atmospheres are a rapidly developing field in astronomy. With recent advances in observing strategies and platforms, there now exist sufficient observations to begin addressing questions of exoplanet composition, as well as atmospheric and interior processes and evolution. Observing exoplanet atmospheres yields spectral information inaccessible through any other currently available strategies. Specifically, comparing this spectral information with atmospheric models involving known physical processes can shed light on the pressure-temperature profiles, chemical compositions, energy budgets, non-equilibrium processes, and presence of clouds or hazes in exoplanet atmospheres.

To accomplish this, I use data mostly from Hubble's WFC3 instrument to search for spectral signatures of water (and some carbon species) in the infrared. I developed, wrote, and executed my own reduction and analysis programs for this work, which will also be applicable for future observations with WFC3. While the data I use for this research was obtained prior to my involvement with the project, I also contributed to several proposals for related research, and additionally conducted ancillary observations with the multi-object spectrograph MOSFIRE on the Keck Telescope, which were unfortunately unusable due to weather.

I will compare the final, observed spectra with model spectra of exoplanet atmospheres contributed by collaborators in order to interpret my findings and apply them to the current body of research.

Chapter 3: Methods

In order to produce meaningful science from raw observations, it is necessary to first reduce the data, then fit the data with model light curves in order to ascertain transit or eclipse depths, and then compare these wavelength-dependent transit or eclipse depths to spectra from models of exoplanet atmospheres. In the following section I summarize each of these stages.

3.1 Reducing the Raw Data

A data reduction package called aXe Kümmel et al. (2009) developed by the Space Telescope Science Institute (STScI) exists for analyzing WFC3 data, but this software was not designed for time-series observations of bright objects, and so I developed a custom pipeline for data reduction. I summarize the requirements for this pipeline below.

WFC3 data can be taken in either stare or spatial scan McCullough & MacKenty (2012) mode. In stare mode, the telescope remains pointed at the target for the duration of the observations, with the target staying fixed on one pixel (or collection of pixels). In spatial scan mode, the target is dragged vertically along a column of pixels, in order to decrease the flux accumulated on any one region of the detector (see Figure 3.1). This decreases the tendency for intrument systematics to appear in the data, though other factors than accumulated flux can affect the strength of the systematics.

In either mode, the target must be identified in the raw image (See Figure 3.2). An extraction box is drawn around the target that maximize the target flux while excluding background objects. The vertical direction corresponds to the dimension through which the target is moved during spatial scan mode. The horizontal direction corresponds to the spectral dimension. Due to the slitless design of WFC3, a wavelength solution must be

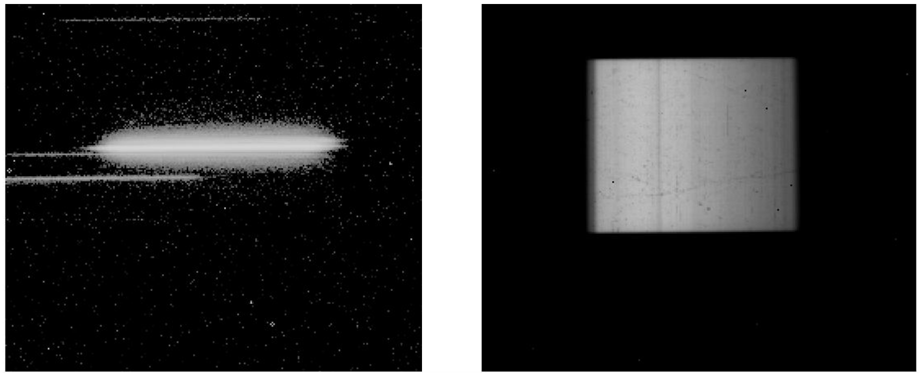

Figure 3.1 Raw WFC3 images for the two possible observing modes: stare mode (left) and spatial scan (right).

found for each source, calcuated using coefficients provided by STScI (amended by Wilkins et al. (2014)) and the image coordinates of a direct (photometric) image taken immediately prior to the start of observations.

I built further routines to identify bad pixels and perform flat-fielding and background subtraction. I also trim the edges of the data in wavelength space, as the sensitivity of the detector drops dramatically near the edges, and this data is unreliable (see Figure 3.3. After these steps have been completed, the images may be summed in the vertical direction, resulting in a time series of 1D spectra. Conversely, this data cube may be thought of as a spectral series of light curves. This is the initial data product.

3.2 Modeling Exoplanet Transits

The transit light curve can be modeled by geometric assumptions that depend on the ratios of the planet and star's radii and the distance between the planet and star. Mandel & Agol (2002) provide such a light curve model with the additional inclusion of a quadratic law for the effects of limb-darkening. Limb-darkening causes a star to be brighter near the center than around the edges, which changes the shape of a transit lightcurve.

In the case of no limb darkening, a planet's transit of a uniform source can be described as follows Mandel & Agol (2002). It is helpful to use normalized units, so we introduce

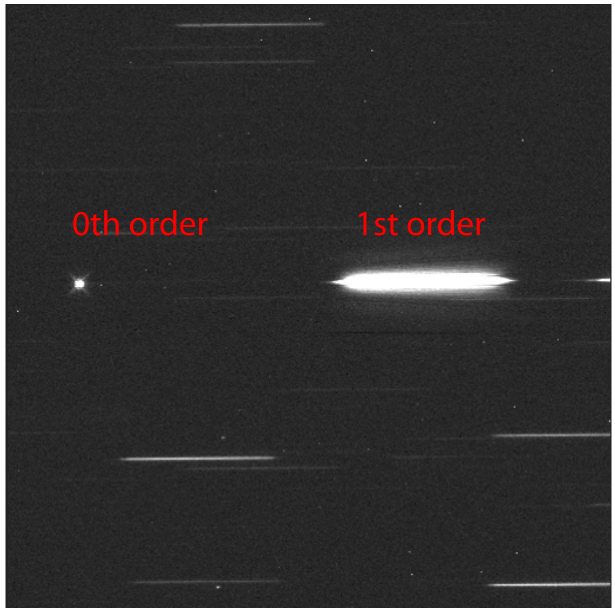

Figure 3.2 A raw WFC3 image, with the 0th and 1st order spectra labeled for a target. Background objects can be seen as additional spectra.

$z = d/R_*$, the normalized separation distance, and $p = R_p/R_*$, the ratio of the planet and star's radii. The flux relative to the unobscured flux is given as $F(p, z) = 1 - \lambda(p, z)$, where λ depends on the size of the planet relative to its star and its distance from the star. The following cases describe the possible situations:

$$\lambda(p,z) = \begin{cases} 0, & 1+p<z, \\ \frac{1}{\pi}\left[p^2\kappa_0 + \kappa_1 - \sqrt{\frac{4z^2-(1+z^2-p^2)^2}{4}}\right], & |1-p|<z\le 1+p, \\ p^2, & z \le 1-p, \\ 1, & z \le p-1, \end{cases} \tag{3.1}$$

Here, $\kappa_0 = cos^{-1}[(p^2 + z^2 - 1)/2pz]$ and $\kappa_1 = cos^{-1}[(1 - p^2 + z^2)/2z]$.

Limb-darkening is a necessary consideration during transit for cooler stars. While a generalized model is given in Claret (2000), a quadratic limb-darkening law is an order of

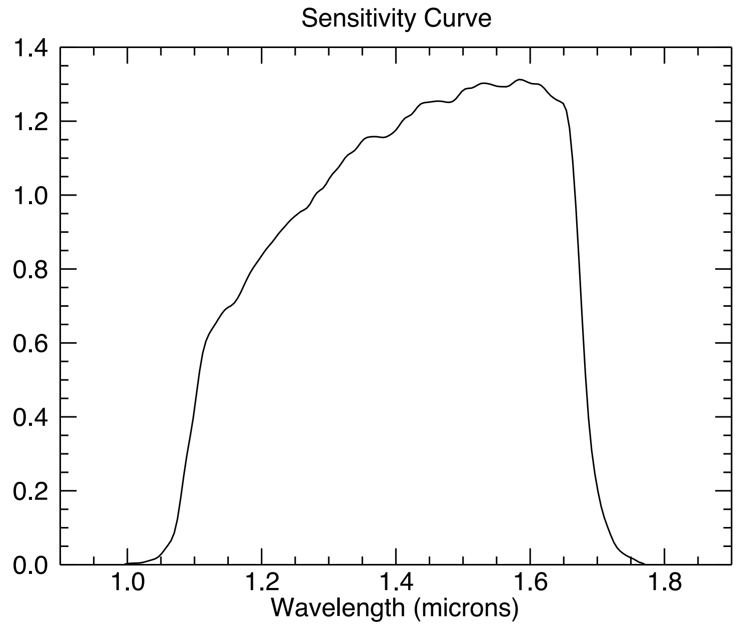

Figure 3.3 Sensitivity curve for the WFC3 G141 grism.

magnitude faster to compute and is a reasonable approximation in most cases. The quadratic model defines the specific intensity of the star, I, as a function of radius r, or more typically, the normalized radial coordinate $\mu = cos\theta = (1 - r^2)^{1/2}$ for $0 \leq r \leq 1$. This limb-darkening law is given as

$$I(r) = 1 - \gamma_1(1 - \mu) - \gamma_2(1 - \mu)^2. \tag{3.2}$$

Here, $\gamma_1 + \gamma_2 < 1$.

In the specific case of fiitting models to the lightcurves we observed with WFC3, we found it helpful to add additional parameters to the light curve model to account for instrument systematics and stellar activity. These parameters account for a decrease in flux throughout the course of an observational visit, stellar oscillations, and/or trends to account for motion of the source on the detector and instrument systematics related to high exposure levels and long read-out times on the detector. These are described in greater detail in the individual observation sections, as the nature of the additional parameters changes depending on the source and the observing mode.

The effective radius of the planet will change with wavelength depending on the presence of absorbers in the atmosphere. Since the transits at different wavelengths are obtained simultaneously by the WFC3 spectrometer, model parameters describing the orbital elements (as well as any secondary astrophysical or instrumental effects) will be unchanged in spectral channels compared to a spectrally integrated (i.e., white) light curve. By modeling the high S/N white light curve first, we can assess these secondary effects and determine the best orbital characteristics. Differences between models of the light curves of each channel or bin of channels will therefore be dependent only on transit depth and the scaling of secondary effects.

In order to determine the best model parameters for combined or channel light curves, we choose to utilize a Markov Chain Monte Carlo (MCMC) simulation using the Metroplois-Hastings algorithm and the Gibbs sampler (Ford (2005)). MCMC works by generating a chain or sequence of states (parameter sets), sampled from some probability distribution. A Markov Chain is one where the next state of the system depends only on the current state, and not any previous states. The Monte Carlo aspect ensures that the generation of each new state is random. The Gibbs sampler is responsible for selecting some random subset of the parameter set to vary for each trial state. The Metropolis-Hastings algorithm dictates generation of a trial state according to a candidate transition probability function, and randomly accepting or rejecting that trial state so as to obtain some desired acceptance probability.

If the trial state is worse (has a higher χ^2) than the current state, then it is rejected. If the trial state is a better fit, it is still rejected some percentage of the time, such that the acceptance rate remains $\sim$0.25 for multi-dimensional systems, or ~ 0.44 for parameter sets of one dimension. This ensures that the simulation is not trapped in local minima of the parameter space. The candidate transition probability function is given as follows:

$$q(x'_\mu|x_\mu) = \frac{1}{\sqrt{2\pi\beta_\mu^2}} \exp\left[-\frac{(x'_\mu - x_\mu)^2}{2\beta_\mu^2}\right]. \tag{3.3}$$

Here x is the current state, x' is the trial state, μ identifies the parameter currently being perturbed, and β is the scaling factor that determines how large a step to take when generating a new trial. The β factors are determined through trial and error. Sample chains are constructed with estimated scaling factors. If the β factors are too small, then the new parameters will be accepted too often; too large, and the acceptance rate will be too low. Once reasonable β factors are determined, the trial chains are discarded and the simulation begins in earnest. At this point, chains are run until they have converged to the stationary distribution. For our studies we find that generally on the order of 10^4 links in 3 independent chains is necessary to ensure thorough exploration of the parameter space. The width of the posterior distribution function defines the uncertainties as drawn from MCMC. If the chains have converged, the posteriors should be Gaussian in nature, centered at the "best-fit" value, with a width corresponding to the 1σ uncertainty.

3.3 Modeling Exoplanet Atmospheres

I do not develop any atmospheric models for this research, but interpretation of my results relies on comparison to models contributed by several collaborators, and I summarize their development here.

Burrows et al. (2000) used models of exoplanet atmospheres primarily controlled by the amount of stellar irradiation, as well as the planet's mass and size. This was expanded upon in Burrows et al. (2006) and Burrows et al. (2008), which additionally used calculations of the redistribution of heat from the planet's day to night side. These models assume solar abundances, but can be tuned to adjust the levels of the major infrared absorbers, as well as adding optical opacity to explain temperature inversions.

Madhusudhan & Seager (2009) developed a new method by which they parameterized the pressure temperature (P-T) profile of the planet, rather than the base-level physical processes that give rise to such profiles. They further only calculate an energy balance at the top of the atmosphere, rather than layer by layer. These improvements allow the running of millions of models across parameter space. They achieved good results with comparisons to observations not only of hot Jupiters (Madhusudhan & Seager (2010)), but also the hot Neptune GJ 436 b (Madhusudhan & Seager (2011)). Madhusudhan (2012) introduced the idea of varying the C/O ratio for exoplanet atmospheres. This was motivated by the inability of models to explain some exoplanet observations, and the understanding that changing the C/O ratio can dramatically affect the expected abundances and therefore observed opacities in exoplanet atmospheres. These models introduced the interpretation of WASP-12 b as a carbon-rich exoplanet (Madhusudhan et al. (2011b)), though these results have since been called into question (Swain et al. (2013)).

Chapter 4: Paper 1: Transit Spectroscopy of Three Hot Jupiters

4.1 Overview

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, 2009a,b; Madhusudhan & Seager, 2009; Madhusudhan et al., 2011b). 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 IRS on *Spitzer*. WFC3 is an optical/NIR camera capable of slitless grism spectroscopy, with wavelength coverage in the 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 μm to 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. (2011b) 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 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.

4.2 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 4.1. The observations were taken with the G141 grism on WFC3's infrared channel, providing slitless spectra covering the wavelength range 1.1 μm to 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 x 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

Table 4.1. Observation Parameters

	WASP-12	WASP-17	WASP-19
Date of Observation	2011-04-12	2011-07-08	2011-07-01
Integration Time	7.624	12.795	21.657
Subarray Mode	256	512	128
CALWF3 version	2.7	2.3	2.3
NSamp	3	16	5
Timing Sequence	SPARS10	RAPID	SPARS10
Peak Pixel Value[1]	38,000	64,000	73,000

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

sampling sequence has a minimum exposure time of ∼7 sec but no maximum. Observations of WASP-17 were taken using the 512 x 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 seconds 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 x 128 sub-array mode with 5 non-destructive reads per exposure, sampled with the SPARS10 sequence. This resulted in an integration time of 21.657 seconds and 70 exposures per orbit, with a total of 274 exposures taken over four orbits. The WASP-12 data utilized the 256 x 256 sub-array mode with 3 non-destructive reads per exposure, leading to an integration time of 7.624 seconds 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 §4.4.1.

4.3 Data Reduction

4.3.1 Image Files: `.flt` vs `.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 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 non-linearity 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, 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 non-linearity. The WFC3 detector generally remains linear up to 78K e^- (WFC3 Handbook); however, Swain et al. (2013) suggest that known WFC3 issues with systematic increases in counts between buffer downloads (see §4.4.1) may be present when count levels exceed 40K DN, or the equivalent of 100K e^-. Our peak counts reach a maximum of 73K $e^-$for WASP-19, with lower values for our other targets (see Table 5.1); we therefore chose WASP-19 to examine linearity. WASP-19 only has a total of 4 SUTR measurements; in Figure 4.1 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 $\sim$0.8% on average, but the channel-to-channel differences were only $\sim$0.1% and would affect the transit depths for individual channels by only $\sim$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.

4.3.2 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 1D 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 pixel 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 PSF as possible

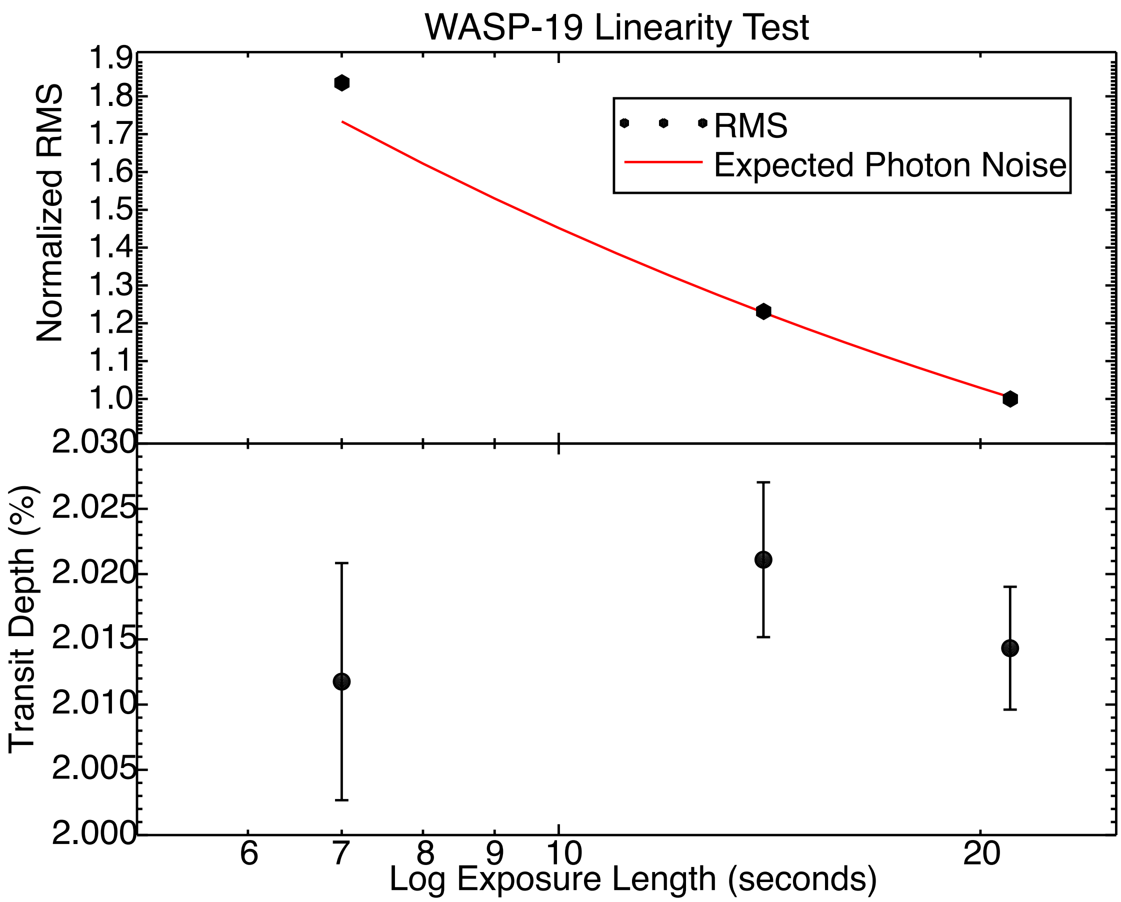

Figure 4.1 Top: The normalized rms compared to expected photon noise for a band-integrated light curve for WASP-19 created using different individual reads. Bottom: Fitted transit depth for each read. The rms follows the photon-limited trend except for the first point, which most likely reflects read noise; the best-fit values are the same within uncertainties.

while avoiding any possible contamination from background sources. We also trimmed the extraction boxes to exclude regions of the spectrum with low S/N, keeping the central 112 pixels of each spectrum.

4.3.3 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 website (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):

$$f(i,j,x) = a_0 + a1 * x + a_2 * x^2 + ...ai * x^n \tag{4.1}$$

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 4.2). The pattern is very similar in each 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

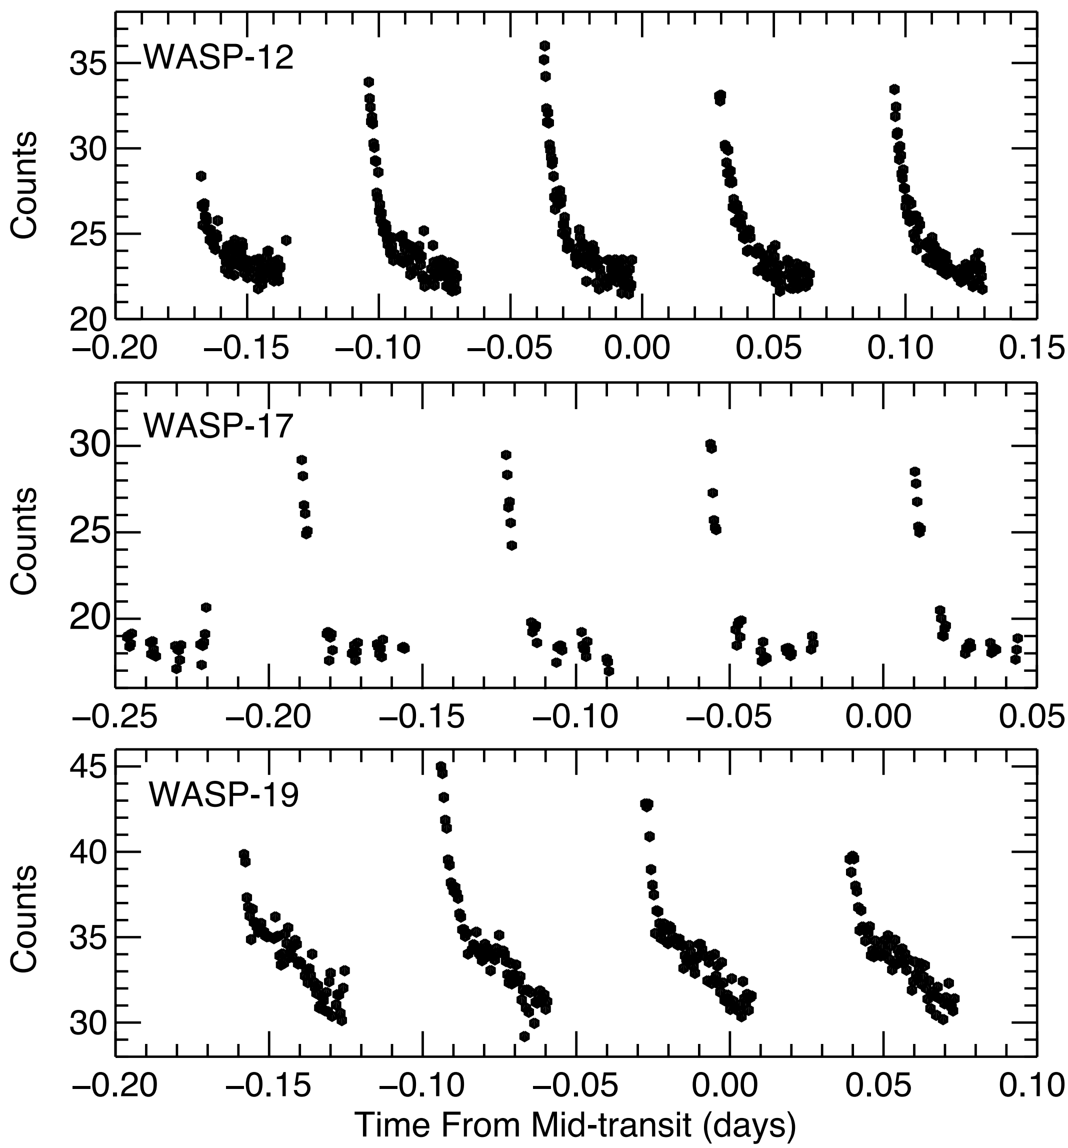

Figure 4.2 Background levels in counts for each target, given as a function of time from mid-transit. The drop in flux on a per-orbit basis is similar for each target, indicating that instrumental effects such as thermal variations during orbit are the likely cause.

variations, we created a 3D image cube and examined each pixel in the 2D 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 2D (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 2 channels for WASP-12, 4 channels for WASP-17, and 1 channel for WASP-19 from further analysis as well.

4.4 Analysis and Results

4.4.1 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 5.1. The band-integrated light curves (Figure 4.3) 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; 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 §4.4.3 and §4.4.4 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

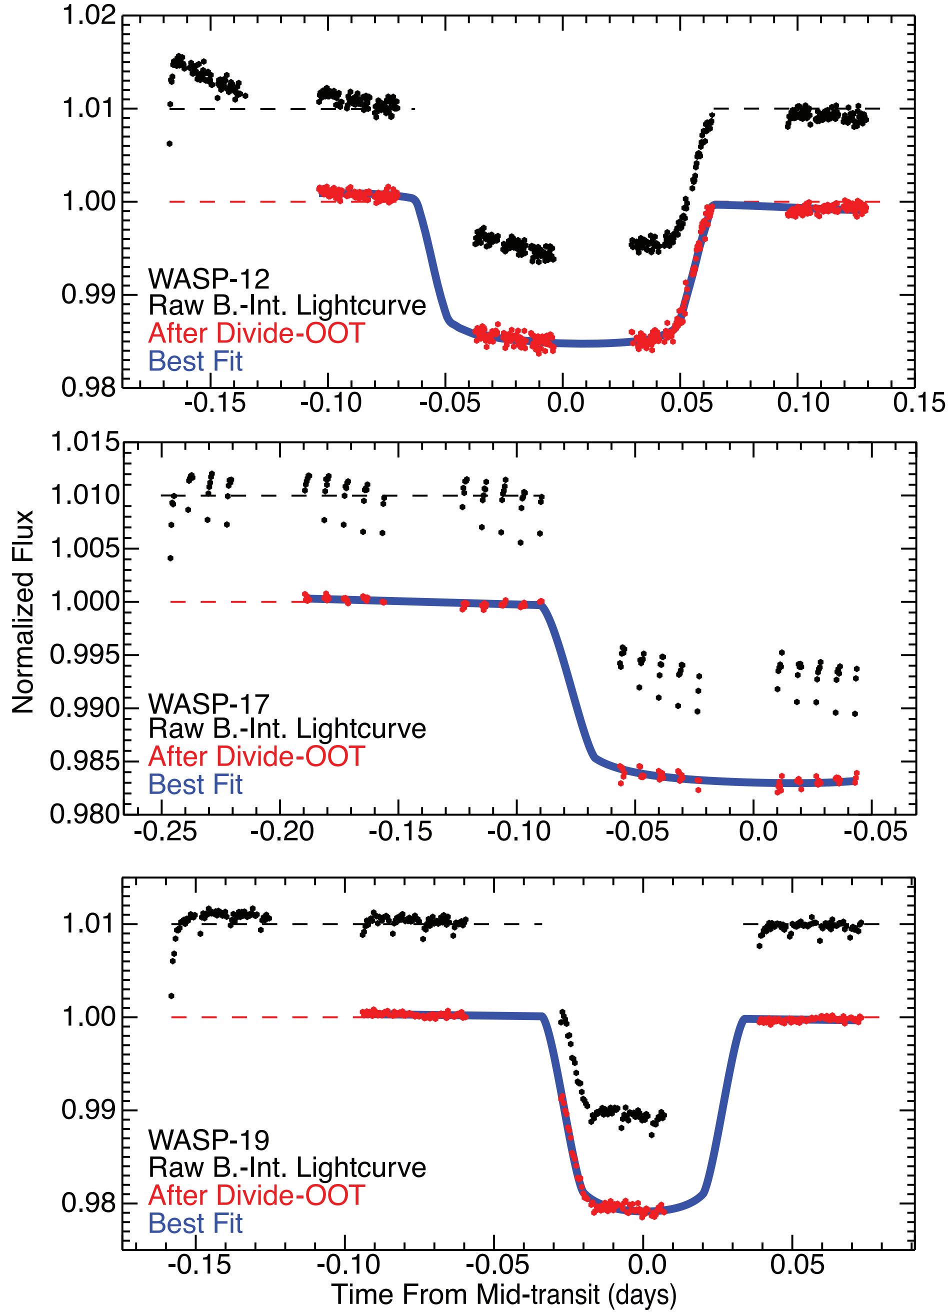

Figure 4.3 The combined-light time series for each source, before and after removing systematic trends. The presence of an intra-orbit pattern is easily identified for WASP-17 and WASP-19, repeating after every buffer read-out, but less obvious for WASP-12. After excluding the first orbit which is inconsistent with the others due to telescope settling, we removed the trends using the divide-oot method devised by Berta et al. (2012).

(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 $\sim$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 4.4) 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 1D 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 4.4 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.

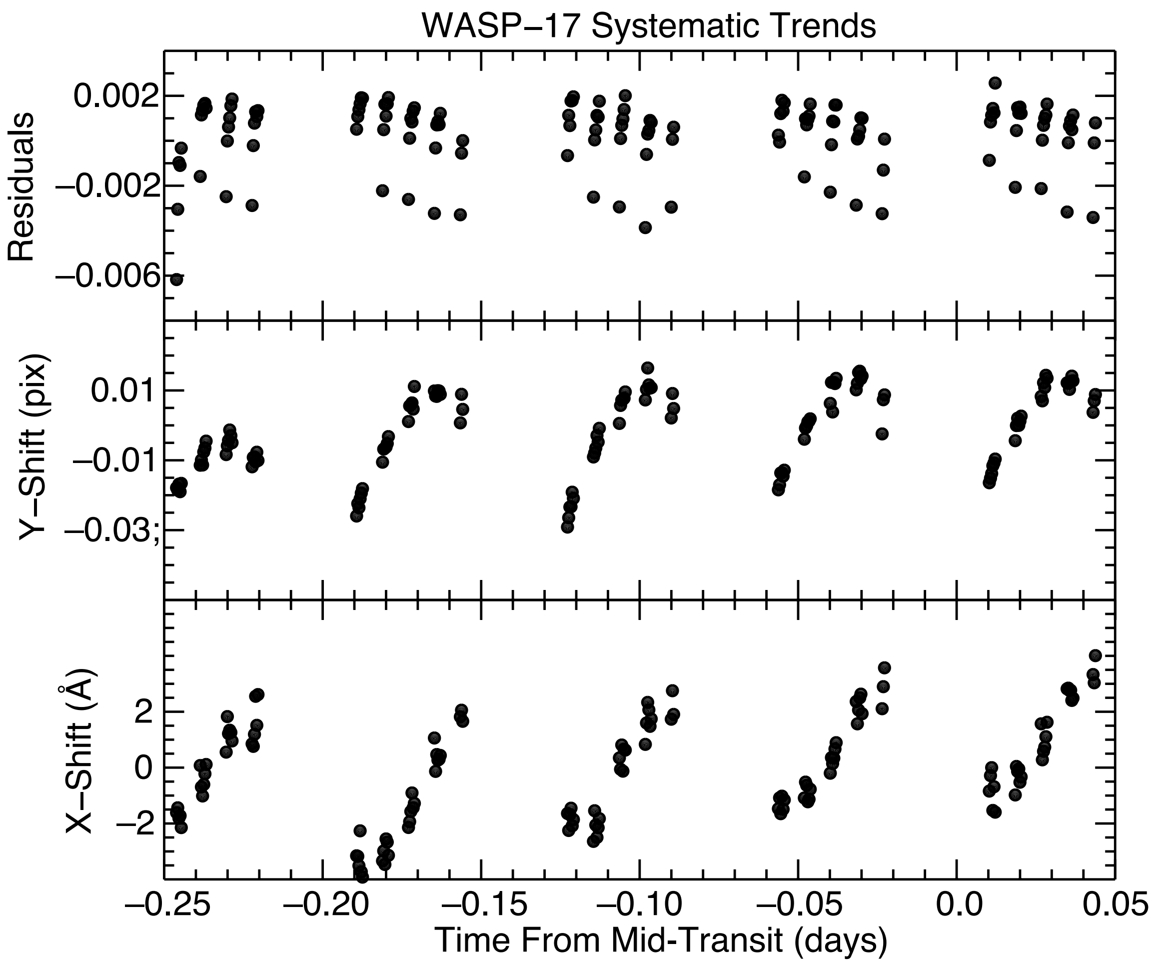

Figure 4.4 Top: The residuals of the combined-light fit for WASP-17, after subtracting our best-fit model. Middle: The shift in the position of the spatial profile of WASP-17 over time, in pixels. The vertical shift was calculated by fitting a Gaussian to sum of the spectral box in the spectral direction. 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 §4.4.1).

4.4.2 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., 2014); 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 $\sim$4 pixels away from the peak of the primary stellar PSF in the spatial direction (see Figure 4.5). This object is most likely the source identified by Bergfors et al. (2013) (referred to as Bergfors-6 by Crossfield et al. (2012) and WASP-12 BC by Bechter et al. (2014)); after correcting for a shift of the 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. (2014) 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 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 point spread functions 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 4.5 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 4.6 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 3 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 1D spectrum at each time step.

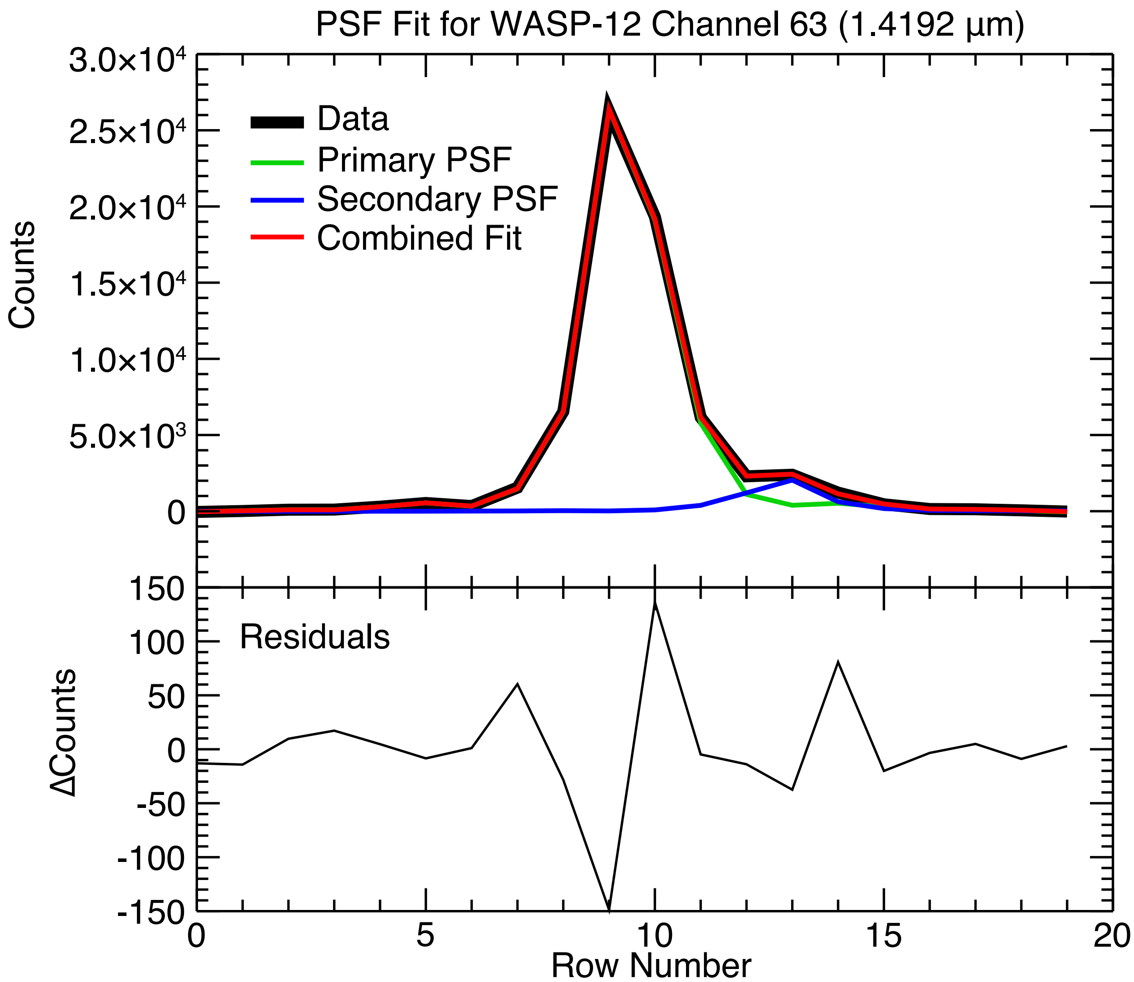

Figure 4.5 Top: Data and the best-fit PSF model for a single channel for WASP-12, using the template-PSF method. The data are shown in black, the fit to the main peak is shown in green, and the fit to the contamination peak is shown in blue; the combined fit is shown in red. Bottom: Remaining residuals after removing the model; the remaining flux under the region of contamination was used as the uncertainty in the contamination flux.

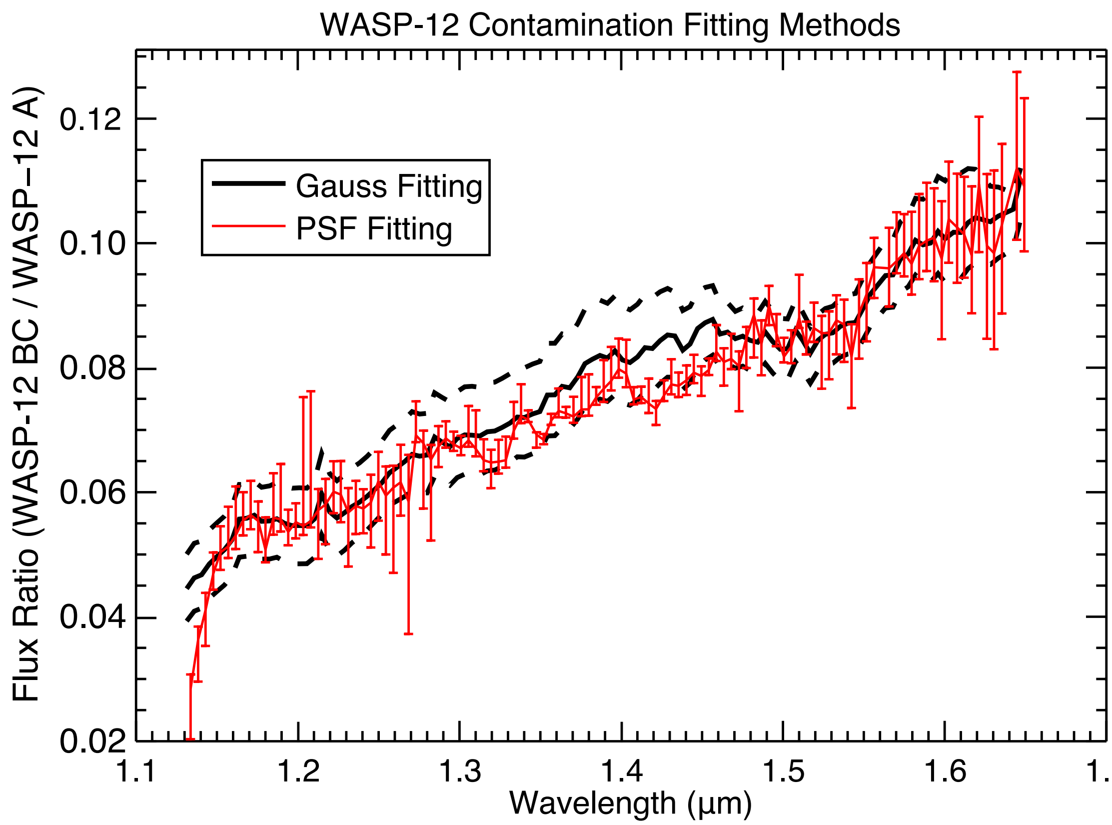

Figure 4.6 Flux ratio for WASP-12 BC compared to WASP-12 A for our two fitting methods. Gaussian fitting (black) subtracts one Gaussian centered on WASP-12's position, then fits another Gaussian to the residuals, centered on the contaminating source. Template PSF fitting (red) jointly scales two PSFs, 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.

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. (2014) 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 4.7 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 $\sim$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. (2014) for both WASP-12 B and C.

A similar calculation of the contaminating flux was used by Stevenson et al. (2014); however, they assumed the lower effective temperature for WASP-12 BC from the spectroscopic analysis by Crossfield et al. (2012) *a priori*, without attempting to determine the contaminating flux empirically. Alternately, Swain et al. (2013) performed a similar fit to ours, but their results appear to lack the sharp downturn shortwards of 1.15 μm and the upturn longwards 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

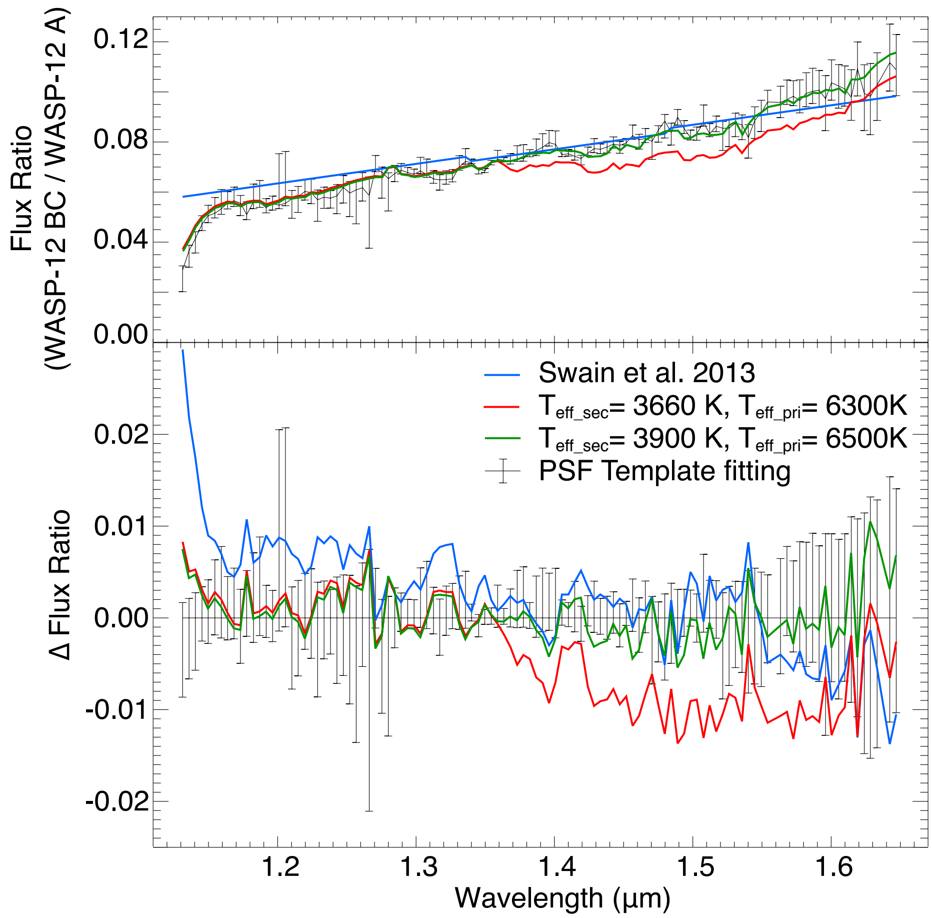

Figure 4.7 Top: Flux ratio for the contaminating source (WASP-12 BC) from the template PSF fitting method (black), compared with analytical models for the flux ratio bracketing the range of values for the temperatures of WASP-12 and the contaminating source (red and green); an approximation to the same 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.

Figure 4.7; 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 §4.5.

4.4.3 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 §4.4.4.

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 4.2), 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 4.3, with the best-fit transit curve overlaid; we tabulate our best-fit orbital parameters in Table 4.3. Our best-fit limb darkening parameters compare well with the expected values from Claret & Bloemen (2011), and best-fit

Table 4.2. 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

[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.

mid-transit times are within the uncertainties based on prior measurements (see Figure 4.8).

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 systematic deviations in the residuals of the

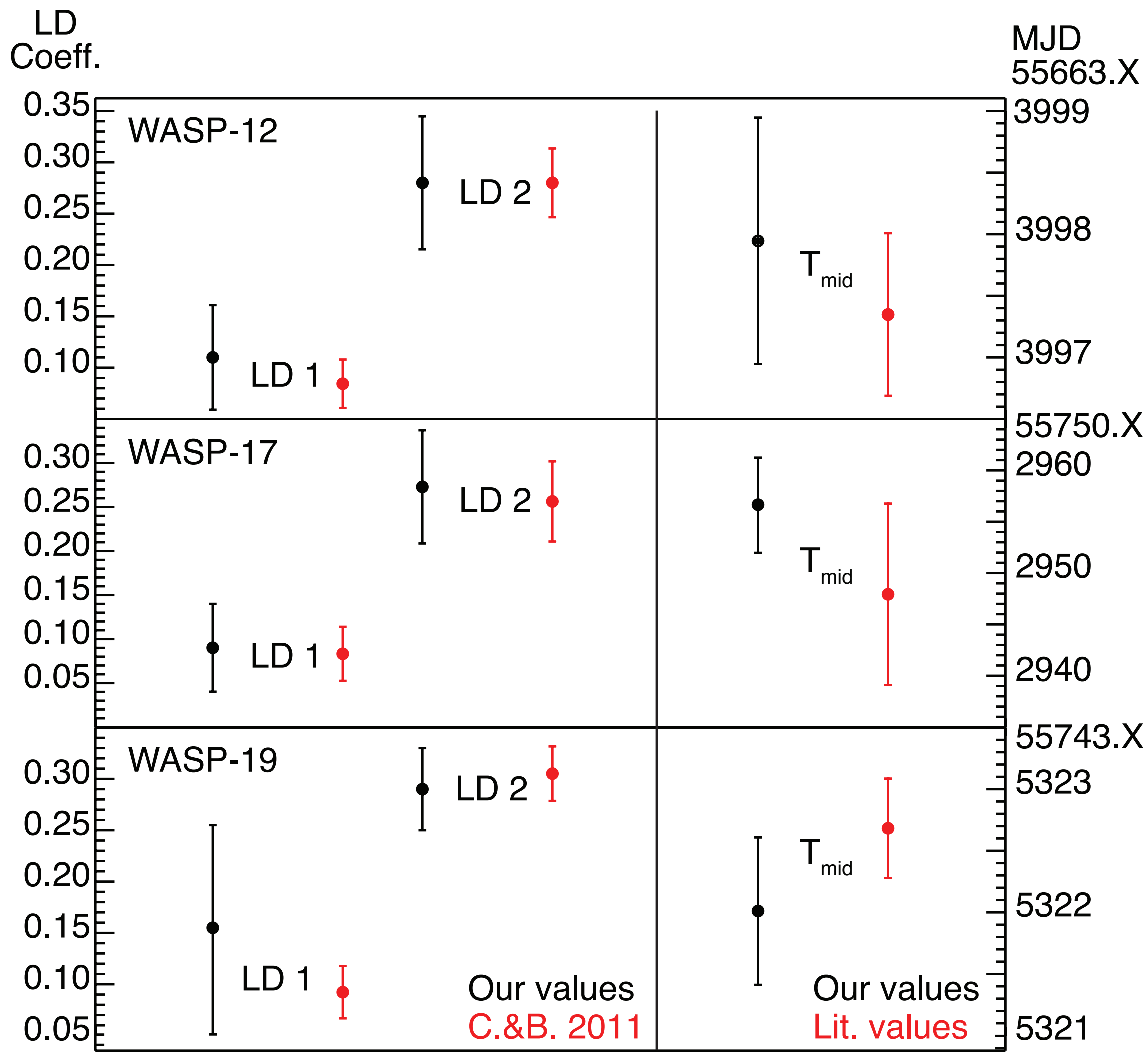

Figure 4.8 Limb darkening parameters for a quadratic limb darkening law shown as calculated using models from Claret & Bloemen (2011), and as found by our MCMC routine, using the Claret & Bloemen (2011) models and uncertainties as priors. Our final values match the expected values within uncertainties for all targets.

Table 4.3. Fitted Parameters From Band-Integrated Time Series

Parameters	WASP-12 b	WASP-17 b	WASP-19 b
R_p/R^*	0.11895 $\pm$ 0.0013	0.12316 $\pm$ 0.00058	0.14140 $\pm$ 0.00093
μ_1	0.085 $\pm$ 0.024	0.083 $\pm$ 0.031	0.092 $\pm$ 0.025
μ_2	0.281 $\pm$ 0.034	0.256 $\pm$ 0.046	0.305 $\pm$ 0.027
T_0 (MJD)	55663.19974 $\pm$ 0.00007	55750.29479 $\pm$ 0.0009	55743.53227 $\pm$ 0.00004
Slope[a]	-0.00793 $\pm$ 0.00034	-0.00578 $\pm$ 0.0010	-0.00407 $\pm$ 0.00039

[a]Linear slope has units of normalized flux per day.

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 4.9). It is difficult to determine the source of these trends due to the limited sampling in 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 2nd-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 ≥ 2 is considered to be positive evidence against the null hypothesis. The BIC was not increased using the non-linear baseline models for either target (ΔBIC ~ -0.5). However, the best-fit peak-to-trough amplitude of 0.0018$\pm$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 & 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 must therefore 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 4.10). 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 4.10), 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 4.3, 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

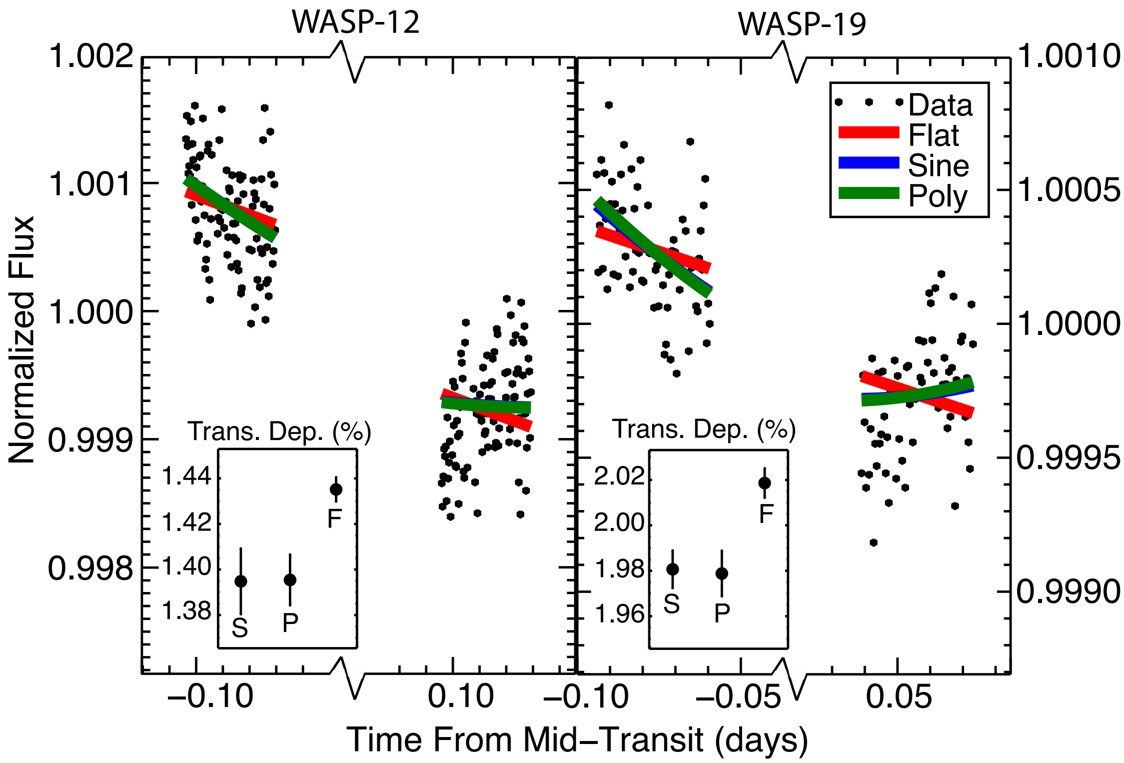

Figure 4.9 The out-of-transit portions of the band-integrated light curves for WASP-12 (left) and WASP-19 (right), with models including only a linear trend (red) and an additional sinusoidal component (blue) or 2nd-order polynomial function (green) over-plotted. The best-fit transit depths for each 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.

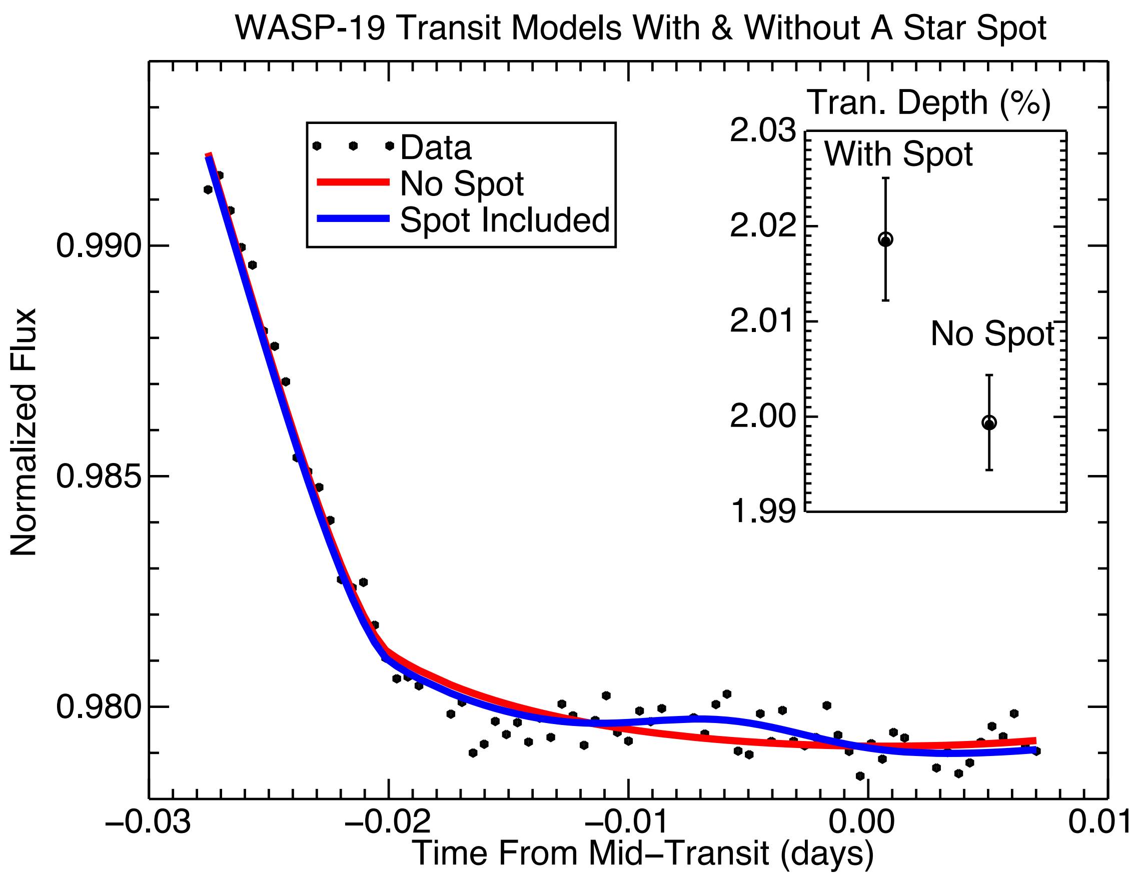

Figure 4.10 The trough of the transit for the band-integrated light curve for WASP-19, with models including standard transit model (red) and a model with a star spot (blue) over-plotted. The best-fit transit depths for each model are also plotted (inset). The value derived incorporating the spot model 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.

for any non-linear 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.

4.4.4 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. (2012a) and Stevenson et al. (2014) (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

$$LC_{final} = LC_{transit} * (a + bt + C_1 * \mathrm{Res}_{BI} + C_2 * \mathrm{Shift}_y + C_3 * \mathrm{Shift}_x) \quad (4.2)$$

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 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 4.11), 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 included only the parameters that provided an improvement in the BIC ($\Delta\text{BIC} \geq 2$) 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 4.12. 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 4.13).

In Figure 4.14 we show final light curves for all of our 0.027 μm-wide bins for each target

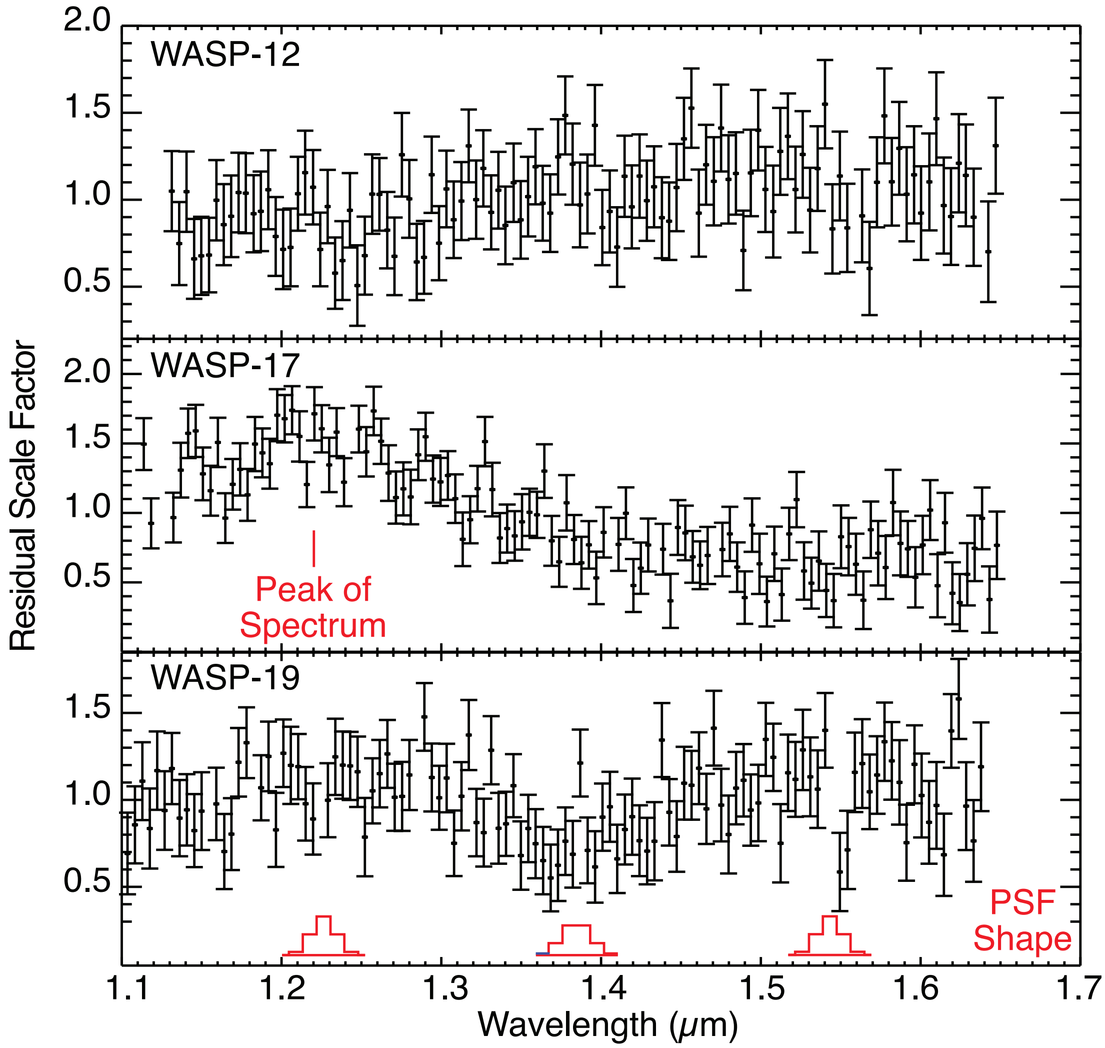

Figure 4.11 Best-fit scaling factors for the band-integrated light curve residuals derived for each channel (see §4.4.4). The relative amplitude of the scaled-residuals component of the model changes with wavelength based on the peak illumination in each channel, and varies between targets based on the sub-array size and sampling mode (see Table 5.1). 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.

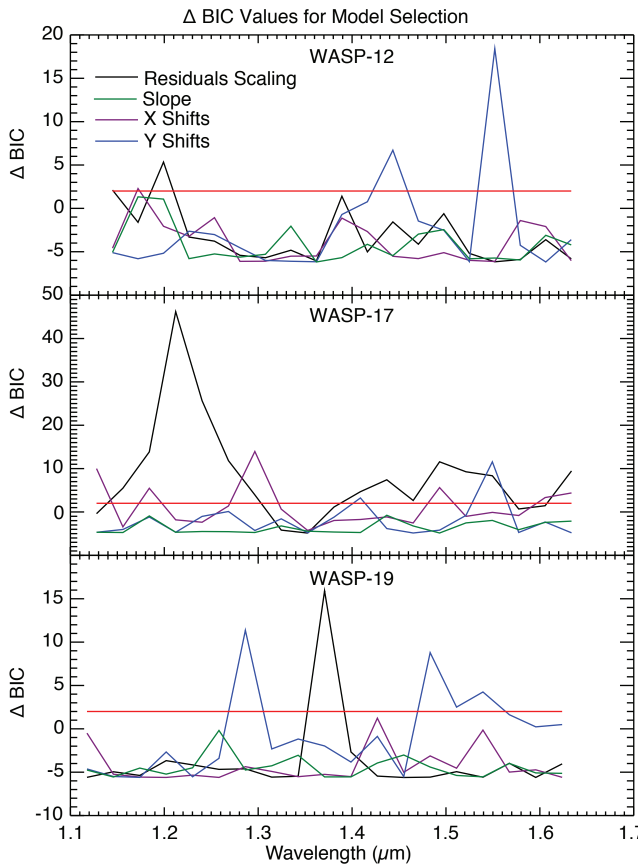

Figure 4.12 We calculate the change in BIC values for a model that fits for additional systematic trends (band-integrated residuals, the visit-long linear slope, x shift, y shift) compared with the default model (see §4.4.4). ΔBIC is shown for each of the 19 bins, for all targets (top: WASP-12, middle: 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 $\Delta\text{BIC} \geq 2$.

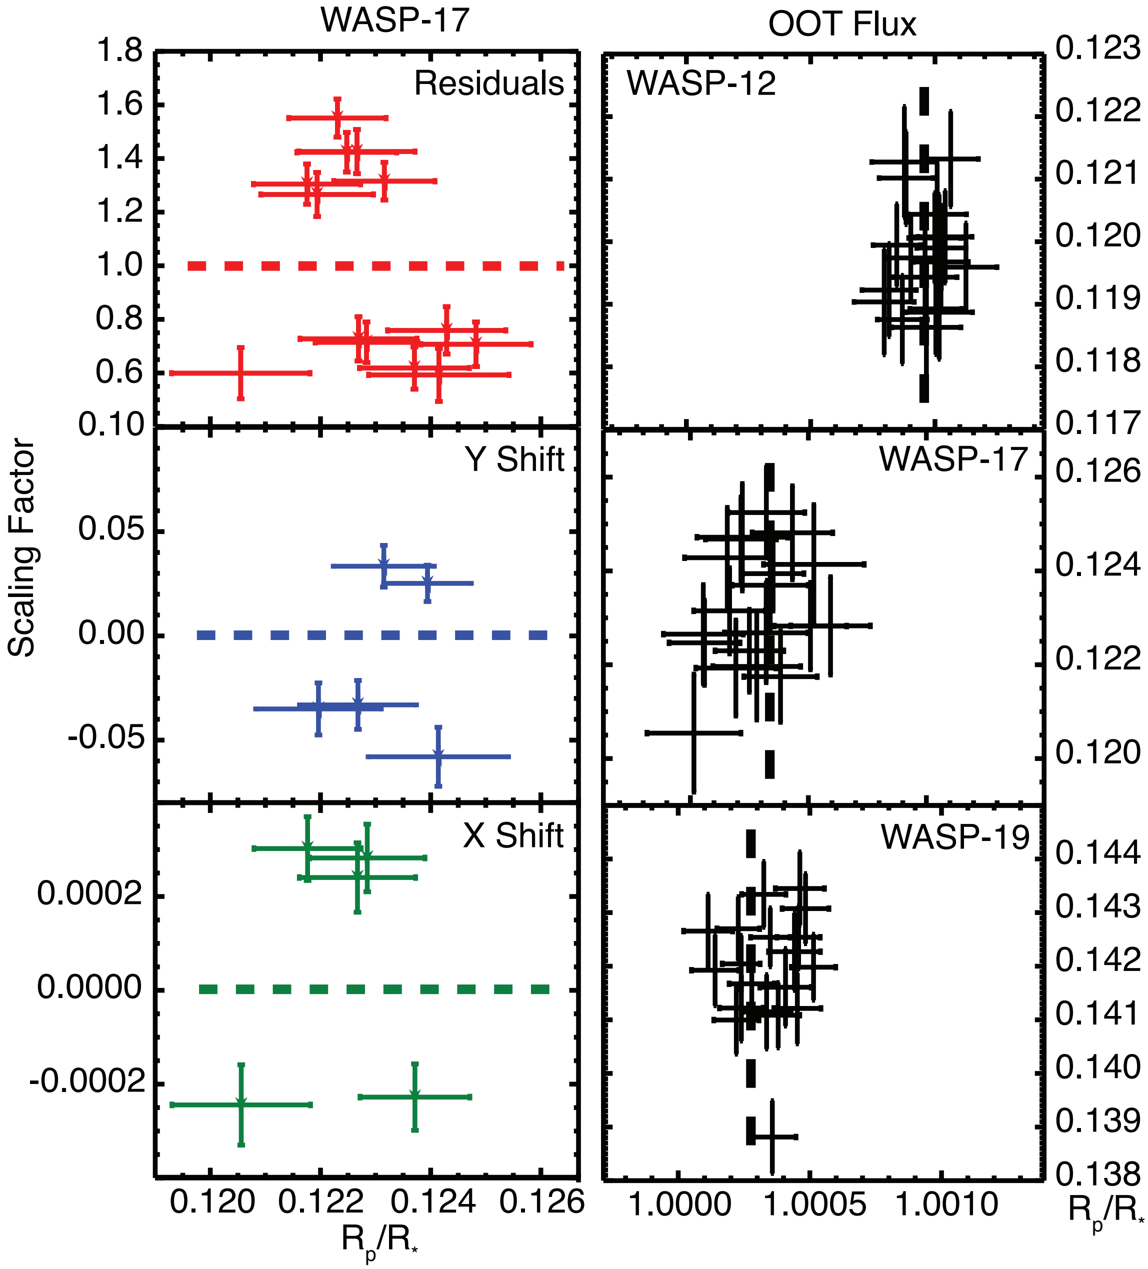

Figure 4.13 Left: Correlation plots for the three model components versus R_p/R_*, for WASP-17 (see §4.4.4). Parameters were only allowed to vary for those bins in which doing so resulted in ΔBIC ≥ 2, and only the bins in which the parameters varied are plotted; the dotted lines represent the default 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 versus R_p/R_* for all the targets. No correlation is seen between R_p/R_* and any of the parameters.

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 4.15) are all fit well by a Gaussian distribution. Our final spectra for each of our science targets are shown in Figure 4.16; 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 (2 channels, 3 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 5-channel bin width, suggesting structure in the spectrum on scales larger than 5 pixels (see Figure 4.17). We therefore chose to use the 6-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. (2014) 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 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.

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

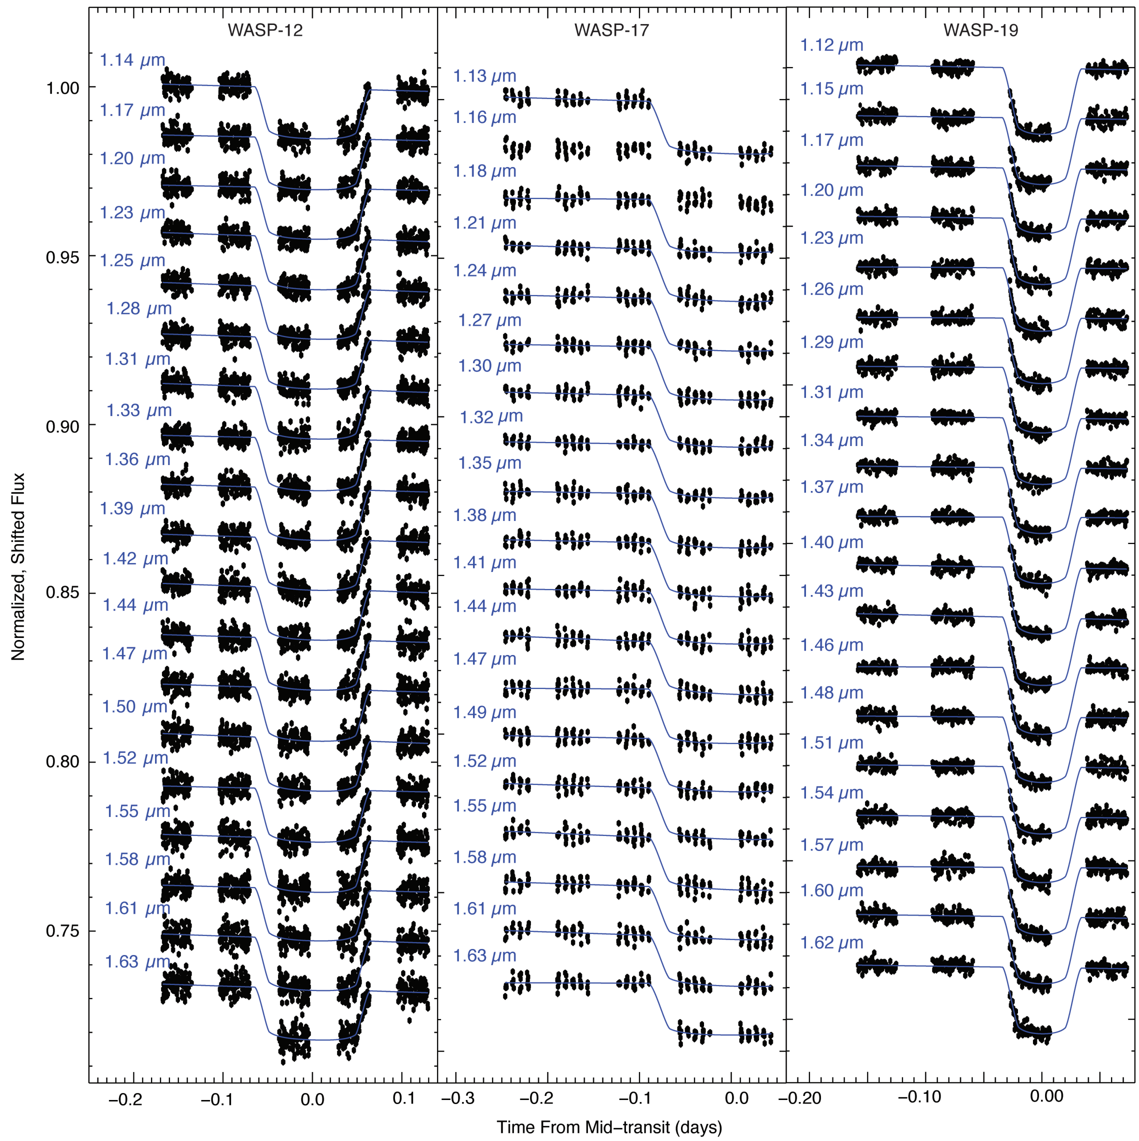

Figure 4.14 The final results for all the bins for each target are shown in black, after removing time series components based on the scaled residuals from the band-integrated light curve, as well as any scaled components based on the spectral shift in the x and y directions that were deemed statistically significant (see §4.4.4). 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.

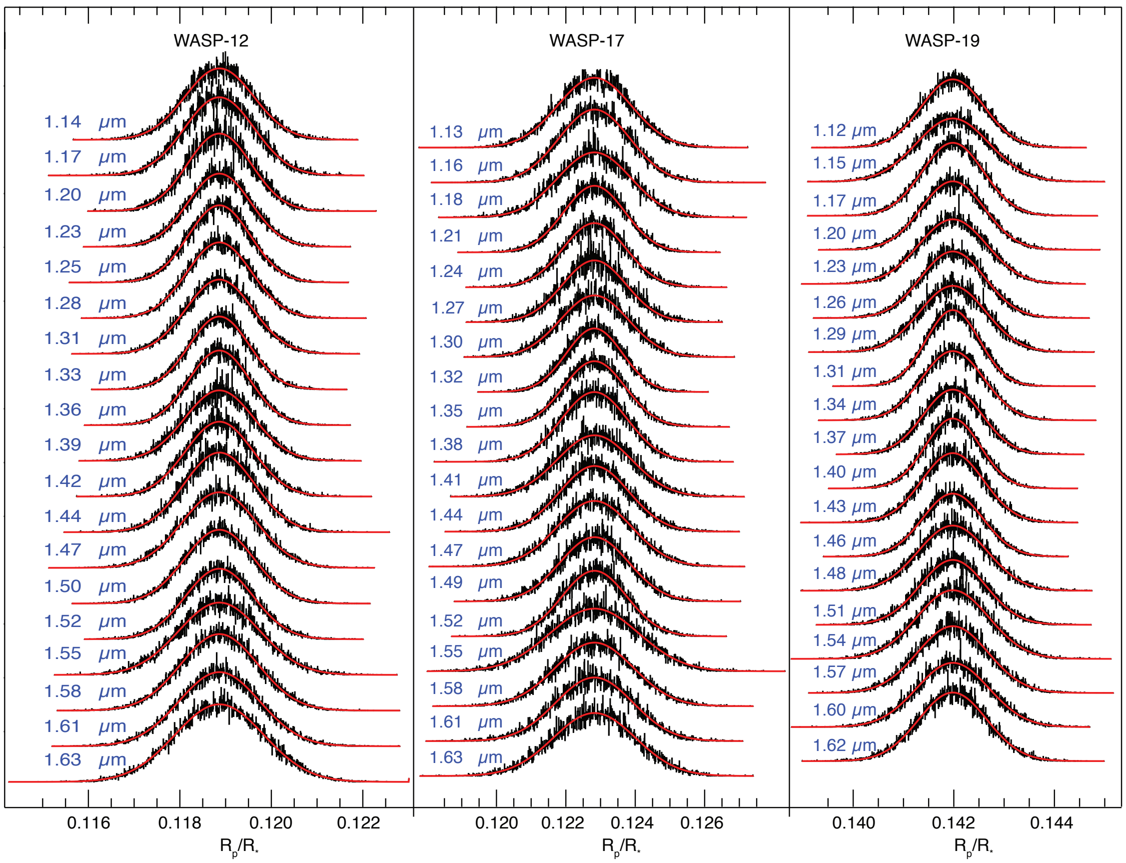

Figure 4.15 Posterior distributions from MCMC for R_p/R_* for every bin, for each of the three targets. All of the final distributions are symmetric and well-approximated by a Gaussian fit (red).

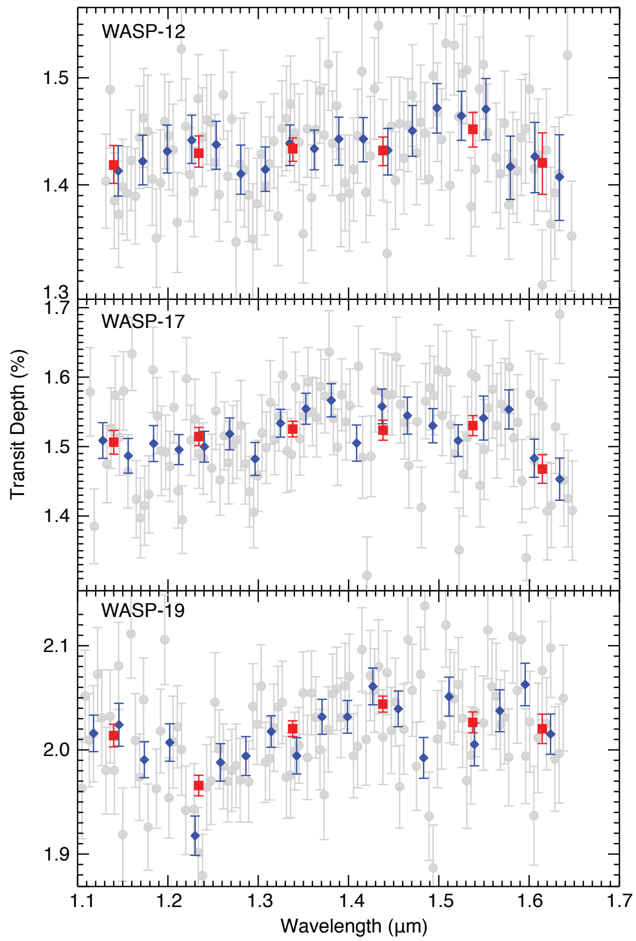

Figure 4.16 Final spectra for each of our targets. The individual channel depths are shown in grey, with the results for 0.027 μm-wide (blue) and 0.1 μm-wide (red) bins overplotted. The differences between the channels and the 0.027 μm-wide bins are consistent with photon-noise variations, but the 0.1 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.

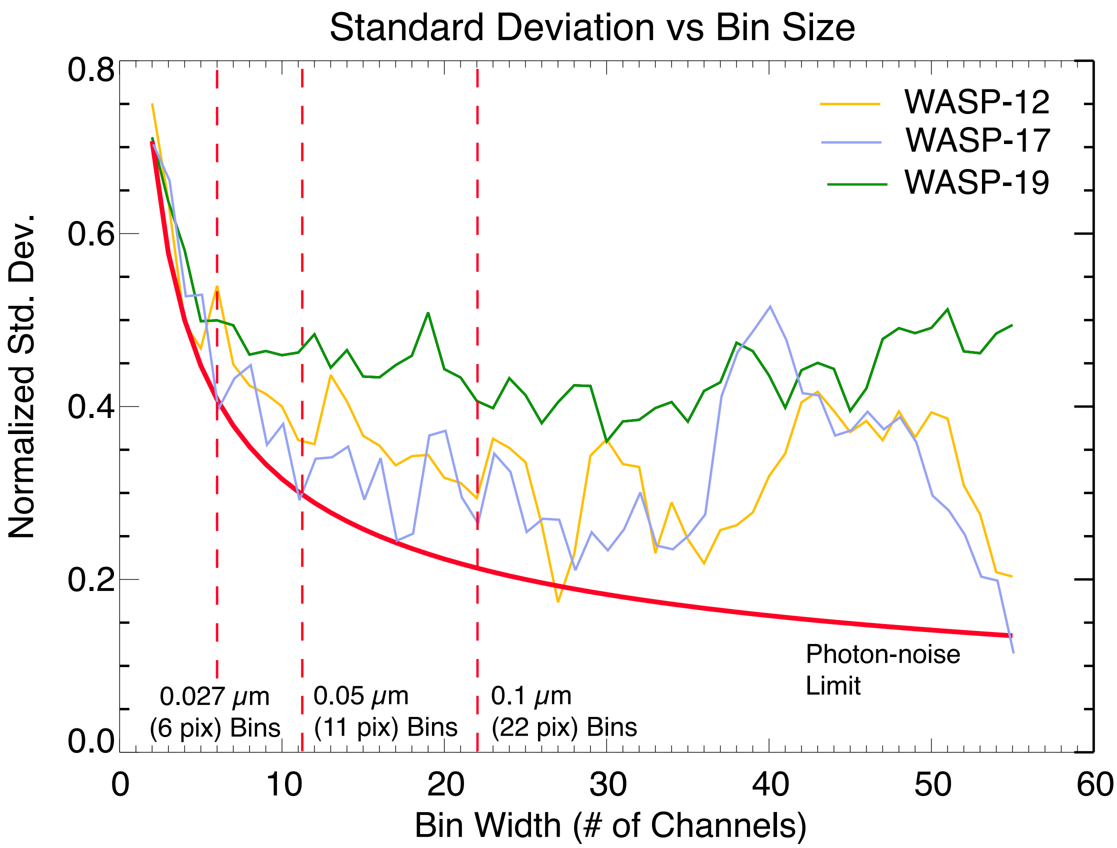

Figure 4.17 After fitting for the transit depths using individual channels, we binned the spectra using bin sizes between 2 and 55 points and then calculated the standard deviation of each binned spectrum; the results for each star are plotted as well as the expected relationship based on photon-noise statistics alone. The standard deviation for all the targets is approximately photon-limited up to 5-channel bins, but then levels off. We use a 6-channel bin size for our final results; spectra using both the 6-channel bins and 22-channel (0.1 μm) bins are shown in Figure 4.16 for comparison.

Table 4.4. 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

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.

Table 4.5. Uncertainty Analysis

Parameters	WASP-12 b	WASP-17 b	WASP-19 b
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[1] σ_{td} (ppm)	52	67	45
σ_{td} from MCMC, Data	53	144	65
σ_{td} from MCMC, Sim.[2]	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[1] σ_{td} (ppm)	174	249	170
σ_{td} from MCMC, Data	180	257	180
σ_{td} from MCMC, Sim.[2]	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

[1]Calculated from the residual rms and the number of points during transit and out of transit

[2]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.

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 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 4.5. We find that the rms of the data is $1.2 - 1.44\times$ the expected photon noise for band-integrated time series, but only $1.11 - 1.22\times$ 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 §4.4.3. 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 (§4.4.3). For WASP-17 the uncertainty for the synthetic data is more than $2\times$ larger than the predicted uncertainty 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.

4.5 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.

4.5.1 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 4.18 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 4.18) 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 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 as 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.

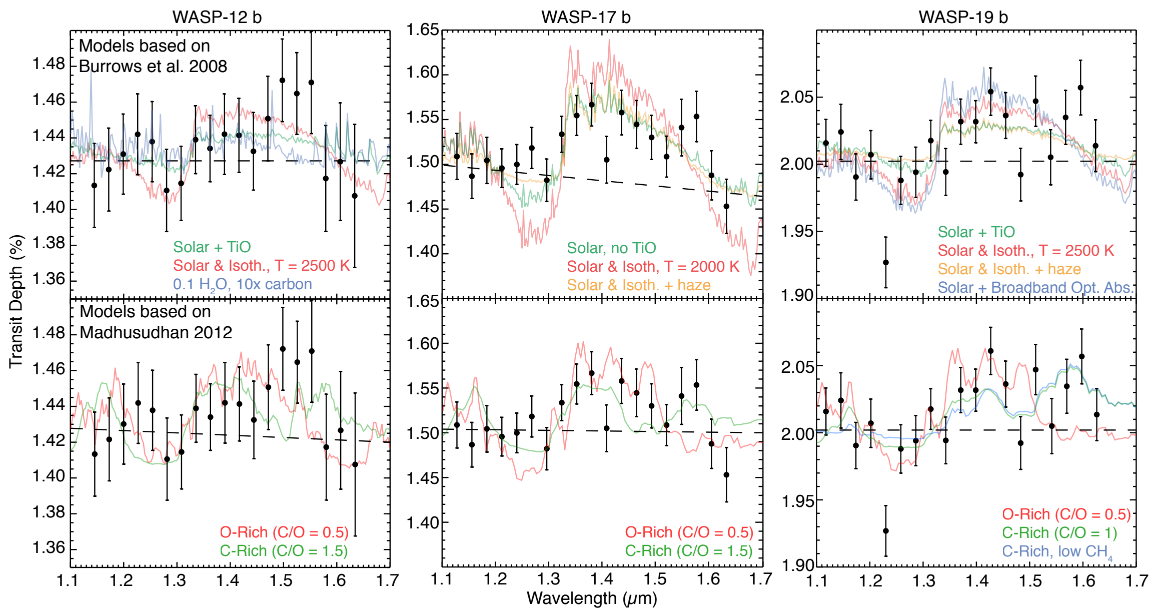

Figure 4.18 Transit depths for each of the 19 bins for each target, with models based on the framework of Burrows et al. (top) and Madhusudhan et al. (bottom). Standard models from Burrows et al. provide a good fit for WASP-17 b and a reasonable fit for WASP-12 b, but for WASP-19 b the models do not 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; 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.

The Burrows models, which are characterized by broad H_2O absorption at 1.4 μm that slopes consistently downward towards 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 shortwards of 1.3 μm, and a haze hypothesis may gain additional support from the fact that the best fits to the models are improved ($\Delta\chi^2 < 0$) in every case by including a linear trend to the models; we discuss the implications of these results in §4.5.2. 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 ($\chi^2_{red} = 0.57$), 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_2O 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 4.18) 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 carbon-to-oxygen (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≥ 1. The models plotted roughly correspond to either an oxygen-rich chemistry (C/O ~ 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_2O feature at 1.4 μm overlaps with CH_4 at 1.36 μm and HCN at 1.42 μm-1.51 μm, while the H_2O 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_2H_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 importance of using bins appropriately sized to be sensitive to the possibility of narrower spectral features in the data; Figure 4.16 demonstrates that using bin sizes larger than ~ 0.03 μm smoothes the data significantly and has the potential to erase the signatures of small-scale fluctuations in the data.

4.5.2 Comparison to Previous Results

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. (2014). Figure 4.19 shows our final spectrum for WASP-12 binned to match Stevenson et al. (2014) and plotted with the results from these two studies. While it is always difficult to pin-point 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. (2014) demonstrated that by using two different transit modeling methods, small differences could be introduced in the spectrum; similarly, we have shown in §4.4.2 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 wavelength 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 $1 - 2\sigma$ 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 which shows a drop shortwards 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 4.7), 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 over time (see §4.4.1); 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. (2014), but we believe the larger uncertainties are warranted based on the uncertainty in the contribution from WASP-12 BC.

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

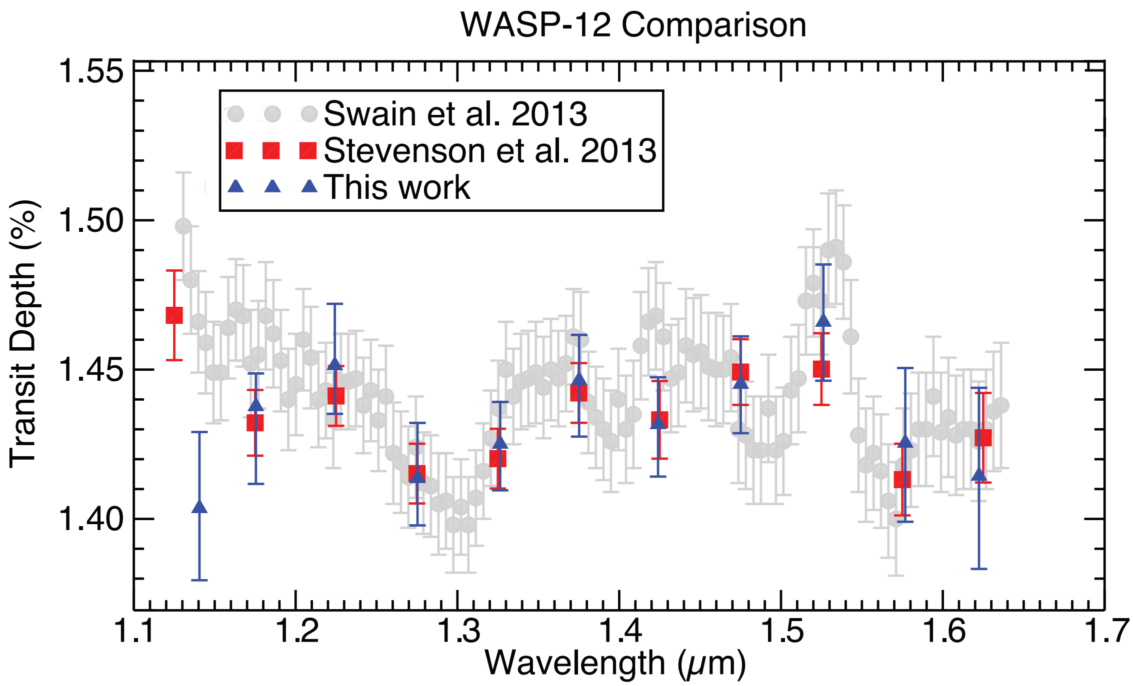

Figure 4.19 Results from Swain et al. (2013) shown in grey, results from Stevenson et al. (2014) shown in red, and from this work in blue. Results from this work have been binned to the same size and number of bins as those used by Stevenson et al. (2014), with edge bins offset due to different choices of spectral trimming. Results from Stevenson et al. (2014) 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.

smaller scale height - WASP-17 b has a scale height that is 3.4$\times$ larger than HD209458 b. In Figure 4.20 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; we calculate a change in the baseline radius of $\sim$1.63$\times$10^4 km across our bandpass for the best-fitting hazy model. If we assume that this spectral slope is due to a change in effective radius with wavelength due to Rayleigh scattering, we can use Eqn. 4 from Lecavelier Des Etangs et al. (2008) to compare our spectral slope to similar results for the spectral slope of HD189733 b across optical and IR wavelengths (Pont et al., 2008; Lecavelier Des Etangs et al., 2008; Pont et al., 2013). WASP-17 b is hotter than HD189733 b by $\sim$400K, and the gravity is lower by a factor of $\sim$7; combining these factors leads to a change in altitude across our bandpass of $\sim$4.65$\times$10^3 km –10$\times$ larger than for HD189733 b, but still a factor of 3.5$\times$ smaller than our best-fit value. Considering the lack of a detectable slope in the data for HD209458 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.

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 3 (0.1 μm); they also subtracted the band-integrated residuals from each bin, but then used

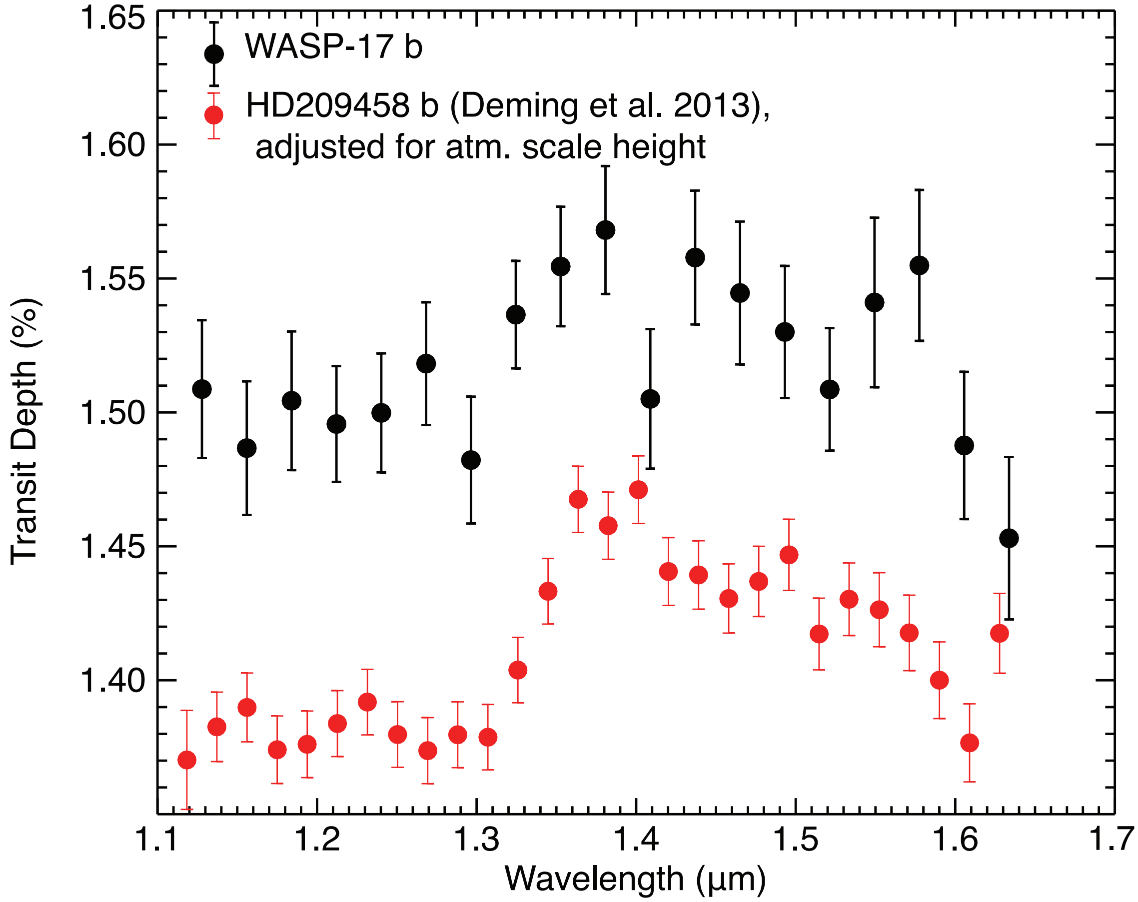

Figure 4.20 Results for WASP-17 b (black) compared with results for HD209458 b from Deming et al. (2013) in red. The spectrum for HD209458 b has been scaled to compensate for the difference in scale height for the two planets. The spectra match very well, suggesting commonality between the spectra for cooler, smaller planets.

the divide-oot method on each bin separately and fit for transit depth and a linear trend. In Figure 4.21 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.

4.6 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

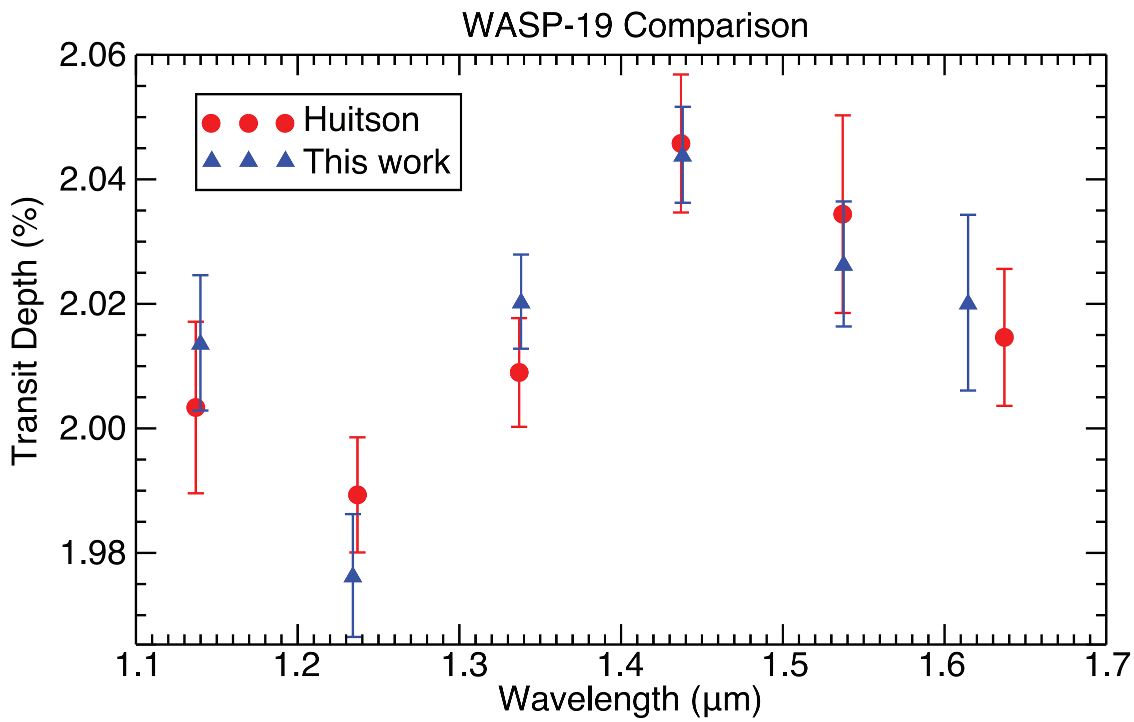

Figure 4.21 Results from Huitson et al. (2013) in red, with results from this work over plotted in blue, binned to the same size, with edges offset due to different choices of spectral binning. The spectra are largely consistent, but comparison with our smaller bin size suggests that the Huitson et al. (2013) may be missing statistically significant features in the spectrum.

model selection strategy that relies on the Bayesian Information Criterion 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., 2014) 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.

Chapter 5: Exoplanet Eclipse Spectroscopy Using WFC3: WASP-33

5.1 Introduction

One of the most intriguing areas of study in the field of exoplanet characterization is the temperature structure of exoplanet atmospheres. Hot Jupiters represent an extreme end of the exoplanet distribution: they orbit very close to their host stars, which subjects them to an intense amount of stellar radiation. Also due to their proximity, they likely become tidally locked on astrophysically short timescales (Guillot et al., 1996), and are heated only on the side facing the star. This results in strong zonal winds (Showman et al., 2008) that redistribute the heat, with the dynamics of this redistribution dictated by the physical and thermal structure of the planet's atmosphere.

Temperature inversions were an early prediction from atmospheric models of highly irradiated planets (Hubeny et al. (2003)), which demonstrated that radiative absorption in the upper atmosphere due to high-temperature absorbers commonly seen in low-mass stars and brown dwarfs such as TiO and VO could produce an increase the the amount of energy deposited above 0.1 bars. Evidence for the existence of thermal inversions began with the first secondary eclipse measurements of HD209458b taken with the IRAC camera on *Spitzer* (Knutson et al. (2008)), which revealed that the planets thermal emission is higher than expected in regions with higher opacity due to features of H_2O and CO (4.5 and 5.6 μm) compared with nearby bands measuring the deeper thermal continuum (3.6 and 8 μm). However, the presence or absence of an inversion appears to defy predictions based on the level of incident radiation or the overall equilibrium temperature of the atmosphere (Machalek et al. (2008), Fressin et al. (2010), and others). More recent models suggest heavy molecules such as TiO and VO may not remain suspended in the upper atmosphere of Jupiter-mass planets (Spiegel et al. (2009)), and searches for spectral signatures of TiO in the optical have been unsuccessful (Sing et al. (2013)). Recent theories have postulated several additional atmospheric processes that could play a role in the formation of inversions, such as the production of photochemical sulfur-based hazes (Zahnle et al. (2009)) or the inhibition of oxide formation due to a super-solar C/O ratio (Madhusudhan (2012)). Furthermore, a possibility explored by Knutson et al. (2010) is that the absorbing molecular species may be destroyed by photodissociation, and created or destroyed by photochemistry.

But these theories remain largely untested because of the broadband nature of the Spitzer/IRAC filters, making the conclusions largely model-dependent and subject to possible systematic offsets or uncertainties. The inference of thermal inversions from IR photometry is based solely on our ability to determine whether there is a larger-than-expected flux from molecular bands compared with the continuum. Madhusudhan & Seager (2010) showed that with only a few data points, this interpretation is heavily dependent on the assumed composition of the planet and the accuracy of the uncertainties ascribed to each measurement. A subsequent Bayesian retrieval analysis on a subset of well-observed planets covering a wide range of effective temperatures by Line & Yung (2013) showed that the data is inconsistent with thermal inversions for many of the planets expected to have an inversion due to the aforementioned theories for the physical origin of the phenomenon. Warm Spitzer photometry has now measured two-band eclipse depths for a large number of planets, but while these measurements can provide some indication of a potential inversion, such data cannot uniquely identify inverted atmospheres because of degeneracies between atmospheric composition and structure (Madhusudhan & Seager (2010), Stevenson et al. (2010), Moses et al. (2013)). Additionally, Hansen et al. (2014) has recently suggested that if the Spitzer uncertainties are significantly higher than reported, all of these results are essentially consistent with featureless blackbody spectra.

It is therefore critical that we further investigate planets that provide the best chance for confirming the presence of temperature inversions, in order to better constrain the actual temperature structure of these planets and clarify the role of various stellar and planetary characteristics in defining this structure. Here we present new occultation observations of WASP-33 b, which orbits a star with significant UV flux and is one of the largest and hottest planets known, using the Wide Field Camera 3 (WFC3) on HST. WASP-33 is an A-type δ-Scuti star, and its planet, WASP-33 b, is one of the most highly irradiated planets discovered to date, orbiting once every 1.22 days (Cameron et al., 2010; Herrero et al., 2011). WASP-33 b is unique, being the only exoplanet yet discovered to orbit a δ-Scuti star. Multiple observations of the host star over wavelengths ranging from the visible to the infrared have shown oscillations with a range of frequencies, and amplitudes on the order of 1 mmag. Previous occultation observations in the infrared (Deming et al., 2012b) concluded that WASP-33 b might host a temperature inversion with a solar composition atmosphere, or a non-inverted atmosphere with enhanced carbon abundance. The inversion scenario is advocated by de Mooij et al. (2013), based on WASP-33 b's apparent inefficient heat redistribution, which was also noted by Smith et al. (2011). Our spectroscopic observations with WFC3 cover a wavelength range from 1.1 to 1.7 μm, which spans a strong water absorption band and therefore provides the opportunity to confirm the presence of an inversion.

We describe the observations in Section 2, data reduction in Section 3, removal of stellar oscillations and analysis strategies in Section 4, and discussion of results in Section 5.

5.2 Observations

Two occultations of WASP-33 were observed on on November 25, 2012 and January 14, 2013. The observations were taken with the G141 grism on WFC3's infrared channel, providing slitless spectra covering the wavelength range 1.1 μm to 1.7 μm at a maximum resolving power of 130 at 1.4 μm (Dressel, 2012). Each target was allocated 5 HST orbits,

Table 5.1. Observations of WASP-33

	Visit 1	Visit 2
Time of first scan (MJD)	56256.405	56306.455
Planetary orbital phase at first scan	0.328	0.354
Time of last scan (MJD)	56256.687	56306.746
Planetary orbital phase at last scan	0.549	0.583
Number of scans	119	119
Number of HST orbits	5	5
Detector subarray size	256	256
Detector reads per scan	7	7
Duration of scan (s)	51.7	51.7
Signal level on detector (electrons pixel $^{-1}$)	7.3×10^4	7.3×10^4

each lasting approximately 90 minutes followed by 45 minute gaps due to Earth occultations of the telescope. This was sufficient to cover a single planetary occultation while including periods of the orbit both before and after occultation.

Both sets of observations were taken using the 256 x 256 sub-array with 7 non-destructive reads per exposure and sampled using the RAPID sampling sequence. The data were observed in spatial scan mode (McCullough & MacKenty, 2012) to decrease systematic patterns in the data that can result from persistent levels of high flux on individual pixels. See Table 5.1 for details.

5.3 Data Reduction

We used the series of single-exposure "ima" images produced by the WFC3 **calwf3** pipeline for our data analysis. The "ima" files are fully reduced data products with the exception of a step to combine multiple reads. The final stage "flt" files provided by the Space Telescope Science Institute are not appropriate for use in spatial scan mode, since the additional pipeline processing for combining multiple reads does not account for the motion of the source on the detector in spatial scan mode. We followed the methodology of Deming et al. (2013) to produce 2D spectral frames from the "ima" files provided on MAST by differencing subsequent reads, and then added the differenced frames to create one scanned image. We used our own strategy from Mandell et al. (2013) to search for and correct bad pixels within the combined spectral frames, and collapse the images into 1D spectra. We used the modified coefficients from Wilkins et al. (2014) to produce the wavelength and wavelength-dependent flat-field calibrations. We did not perform background subtraction, because the differencing method used to combine reads removes most surrounding background, and the amount that remains is negligible compared to the stellar flux.

We identified in the direct image the nearby star noted by Adams et al. (2013), which lies 1.9" from WASP-33. However, this nearby star has a magnitude of only $\Delta K_s = 5.69$, which means it accounts for a flux difference of $\sim$0.5%, well below our measured uncertainties. Due to this negligible effect on our data and the difficulty of removing such an object, especially from spatial scan data, we decided not to attempt any corrections for the additional flux during our analysis.

We trimmed roughly 70 pixels from either end of the spectral extent, to remove the parts of the spectrum with low sensitivity. After trimming the edges of the spectrum, the spectrum covers the region between approximately 1.13 and 1.63 μm. We also identified a strong feature at 1.28 μm due to the Paschen β stellar feature, and took care to isolate this feature when defining our spectral bins due to potential shifts in the spectrum which could cause sharp changes in flux.

Table 5.2 From top to bottom, the evolution of the reported frequencies and amplitudes for WASP-33's pulsation spectrum. This table has been directly copied from von Essen 2013 and modified, and should be credited as such.

Frequency (c/d)	Amplitude (mmag)	Central Wavelength (μm)
Herrero et al. (2011)		
21.004 ± 0.004	0.98 ± 0.05	0.66
21.311 ± 0.004	∼0.86	0.66
Smith et al. (2011)		
26.9 ± 0.4	1.479 ± 0.069	0.91
18.8 ± 0.6	0.567 ± 0.134	0.91
34.3 ± 0.4	0.766 ± 0.115	0.91
21.6 ± 0.6	0.605 ± 0.105	0.91
Deming et al. (2012b)		
∼21.1	∼1.3	3.6
∼21.1	∼1.3	4.5
∼20.2	∼2.3	2.15
∼9.8	∼1.6	1.25
∼26.6	∼2.1	2.15
∼11.4	∼2.1	2.15
de Mooij et al. (2013)[1]		
22.5 ± 0.1	0.95 ± 0.04^{N1}	2.19
33.3^{N1}	0.41 ± 0.04	2.19
27.3 ± 0.2	0.56 ± 0.06^{N2}	2.19
33.2^{N2}	0.17 ± 0.05	2.19
22.0^{N1}	0.11 ± 0.06	2.19
17.1^{N1}	0.13 ± 0.06	2.19
Kovács et al. (2013)[2]		
15.21643 ± 0.00004	0.758 ± 0.085^{HN}	0.81
20.16229 ± 0.00004	0.733 ± 0.080^{HN}	0.81
21.06339 ± 0.00004	0.719 ± 0.078^{HN}	0.81
15.21517 ± 0.00001	$0.477 \pm 0.054^{H+F}$	various
20.16230 ± 0.00001	$0.739 \pm 0.053^{H+F}$	various
21.06346 ± 0.00001	$0.728 \pm 0.049^{H+F}$	various
von Essen et al. (2013)		
20.1621 ± 0.0023	1.03 ± 0.03	0.49
21.0606 ± 0.0023	1.01 ± 0.03	0.49
9.8436 ± 0.0023	0.86 ± 0.03	0.49
24.8835 ± 0.0017	0.45 ± 0.03	0.49
20.5353 ± 0.0013	0.77 ± 0.03	0.49
34.1252 ± 0.0027	0.53 ± 0.03	0.49
8.3084 ± 0.0025	0.68 ± 0.03	0.49
10.8249 ± 0.0030	0.69 ± 0.03	0.49
This work		
19.88 ± 0.32	0.62 ± 0.05^{N1}	1.4
29.65 ± 0.48	0.32 ± 0.04^{N1}	1.4
14.40 ± 0.56	0.65 ± 0.16^{N2}	1.4
22.16 ± 0.57	0.50 ± 0.07^{N2}	1.4

[1]Following the nomenclature of de Mooij et al. (2013), the superscripts $N1$ and $N2$ refer to "Night I" and "Night II".

[2]The superscript HN stands for "HATNet" and $H + F$ for "HATNet+FUP".

5.4 Analysis and Results

5.4.1 Identifying and Fitting Trends

Stellar Oscillations

WASP-33 is known to be an oscillating δ-Scuti star whose pulsation frequencies have been measured over multiple campaigns (Herrero et al., 2011; Smith et al., 2011; Deming et al., 2012b; de Mooij et al., 2013; Kovács et al., 2013; von Essen et al., 2013; Sada et al., 2012), with most groups finding agreement for pulsations at the 1 mmag level. A wide range of pulsation frequencies are found (see Table 5.2), but since multiple oscillation modes are to be expected, and that the strength of these modes will vary with wavelength, observations taken across a range of spectral bands and at various times should not be expected to have perfect agreement. The incomplete temporal sampling caused by HST orbits complicated characterization of the oscillation modes in our own data, and so we explored different avenues for constraining the detectable pulsation frequencies and removing the stellar oscillations.

We divided frequency space into regions based on the frequencies identified by previous observing campaigns and allowed our MCMC models to fit, iteratively or simultaneously, between 1 and 3 sine curves restricted to those regions of frequency space. We used Bayesian information criterion (BIC; Schwarz (1978), Liddle (2004)) to determine the best combination of sine terms. We find that two sine curves achieve the best results, and that the frequencies and amplitudes identified are robust whether we fit the sine curves simultaneously or in sequence. The results are shown in Table 5.2, and our best-fit frequencies agree roughly with previously determined values.

Red noise remains in the residuals after removal of the two sine curves representing the stellar oscillation modes, indicating that we are unable to fully characterize either the stellar oscillations or underlying instrument systematics. These instrument systematics should be weak in spatial scan mode, but potentially still present (Deming et al., 2013). Because the first orbit of WFC3 observations tends to be more noisy than subsequent orbits (Mandell et al., 2013; Deming et al., 2012b), we do not include this orbit for the band-integrated fitting process (including fitting for the stellar oscillations), though we do incorporate it later in our wavelength-dependent relative analysis.

Spectral Shifting With Time

For each image, we also calculated the shift in the horizontal (i.e. spectral) direction referenced to a template exposure in the time series, allowing us to correct for any modulation in channel flux due to spectral features.

In our initial analysis procedure we fit the slope of the pixels at the short-wavelength and long-wavelength edges of the spectrum, where the sensitivity curve causes the shape of the spectrum to change most dramatically, to determine the spectral shift in each exposure. The slopes measured at both edges of the spectrum were averaged to decrease the effective uncertainty of the measurement. We used the zeroth-order coefficient of this fit to determine the shift of each spectrum relative to the first exposure in the time series.

We compared this strategy to that employed by Deming et al. (2013), which used instead the central region of the spectrum to measure the shits. In this method, the template spectrum comprises an average of the exposures in the time series immediately preceding and following eclipse. We interpolated the template spectrum onto a wavelength grid shifted in either direction up to a pixel and a half, stepping in 0.001 pixel increments, and saved each shifted spectrum. For each exposure, we compared the observed spectrum with each shifted template spectrum and calculated the rms; the shift corresponding to the lowest rms is saved as our best-fit spectral shift.

We compared the two spectral shift measurement strategies, finding that if the edges of the spectrum are included for the latter method, then the resulting shifts match the "edge only" measurements of our method very closely. If instead only the central region of the spectrum is used to determine the shifts, the measured shifts change and result in significant change in the final transmission spectrum. This result suggests that the shape or placement of the grism sensitivity function changes as a function of time or placement on the detector, either due to changes in the optical path or as part of the thermal breathing modes of the telescope.

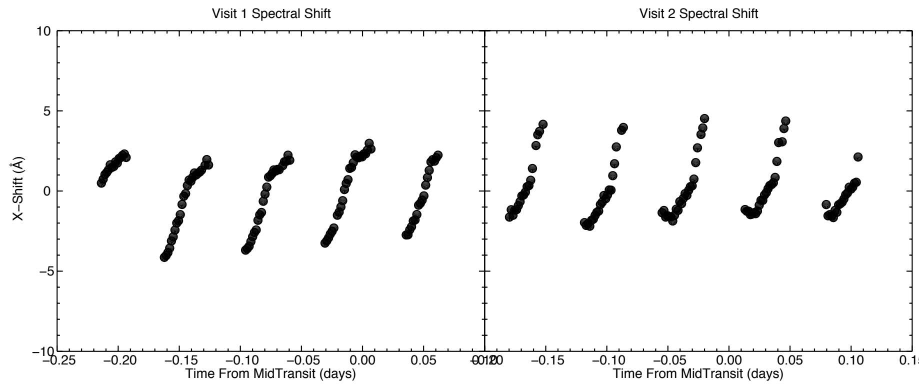

Figure 5.1 White light shifts in the spectral direction for both visits across the eclipse duration.

The use of the "center-only" spectral shifts resulted in a lower residual rms for the resulting light curves, suggesting that the fit was improved by using shifts derived from the same portion of the spectrum as the light curves themselves. We further extended this analysis by determining the specific spectral shift for the specific portion of the spectrum associated with each binned light curve, and using that set of shifts in the systematic decorrelation procedure. This has the advantage of using the set of shifts that best describe the region of the spectrum used by each bin. As seen in Figures 5.3 and 5.2, the overall trend of the shifts with time does change with wavelength. Given this, we advocate careful inspection of the shift of the spectrum on the detector with time, and for this work, we use these binned shifts for our final analysis.

In all cases, the shifts take the form of a repeating, inter-orbit pattern, as well as a visit-long slope. We remove the visit-long effect from the xshifts, since a visit-long trend in flux is also seen in previous WFC3 data sets and we choose to instead fit for this slope as an independent parameter in our eclipse curve fit. The horizontal shifts used in our analysis are plotted in Figure 5.1.

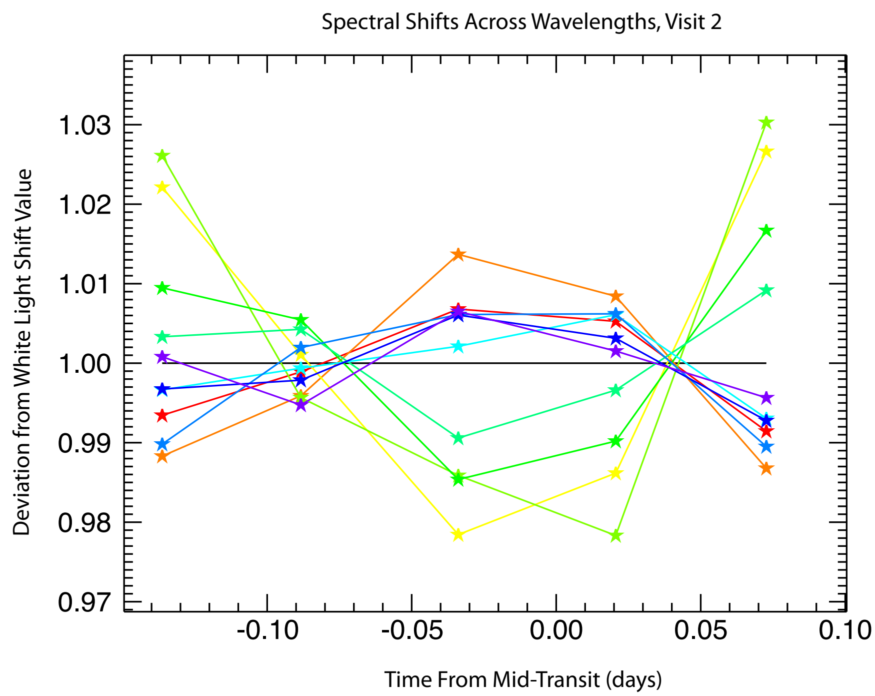

Figure 5.2 Shifts in the spectral direction by wavelength, binned in time, for Visit 2.

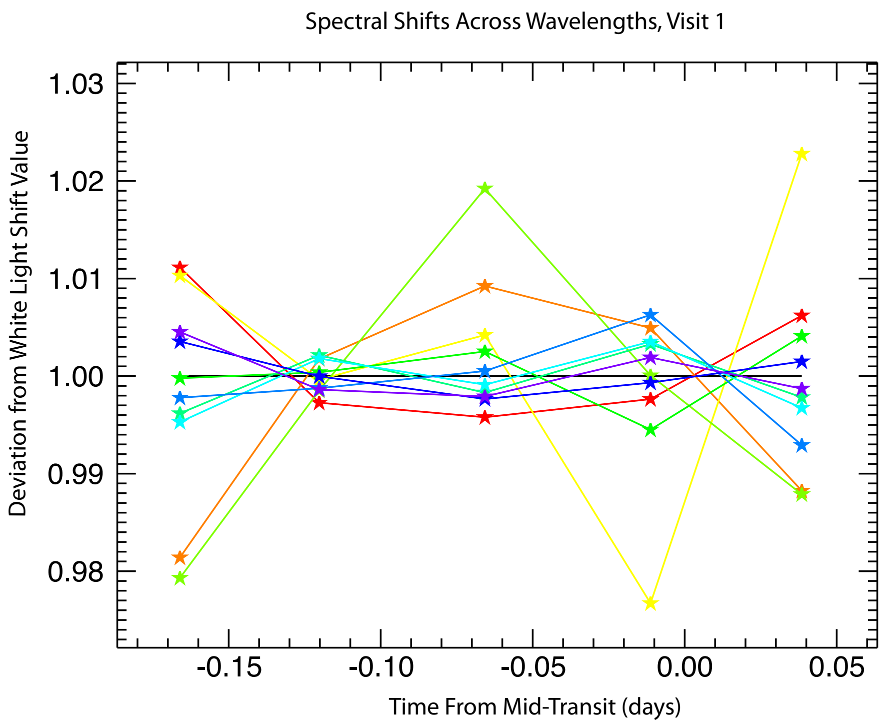

Figure 5.3 Shifts in the spectral direction by wavelength, binned in time, for Visit 1.

As an alternative strategy, we can choose to interpolate each exposure's spectrum onto a shifted wavelength grid according to its best-measured shift value, as determined by the rms, and use these shifted spectra for light curve fitting. In this case, we do not use the scaled xshifts as a parameter in our MCMC fitting. We find that the transmission spectrum is minimally affected by the method of correction (interpolation vs xshift scaling), so long as the same wavelength ranges (center, center + edges, or binned central region) are used to measure the shifts for both methods.

5.4.2 Band-Integrated Eclipse Curve Fitting

Our analysis strategy relies on the fact that almost all of the time-dependent trends present in the band-integrated light curve are consistent across wavelengths, since the systematics are related to the general exposure parameters, correlated with the illumination of each pixel, and/or are wavelength independent, as in the case of the stellar oscillations, though the amplitudes of some of these trends may change with wavelength. We therefore decided to determine the band-integrated eclipse curve parameters first, and then use the residuals from this band-fit as a component in our eclipse model when fitting individual spectral channels, with the amplitude of this component allowed to vary (Mandell et al., 2013; Stevenson et al., 2014). This method allows us to incorporate any common-mode systematic trends into our fit, providing a more robust measurement across spectral channels of the relative change in eclipse depth, which is the most important factor when measuring the depth of spectral absorption features.

We fit the light curve with a Markov Chain Monte Carlo (MCMC) routine with a Metropolis-Hastings 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 eclipse data to date as well as two sine curves to account for the stellar oscillations. We lock all orbital parameters, leaving open for fitting only the eclipse depth, slope, and sine terms. Orbital parameters are listed in Table 5.3. For each

Table 5.3. Orbital and Stellar Parameters for WASP-33 [a]

Parameter	Value
Period (days)	1.2198709
i (°)	86.2 ± 0.2
R_p/R^*	0.1143 ± 0.0002
a/R^*	3.69 ± 0.01
Semi-major axis (AU)	$0.0259^{+0.0002}_{-0.0005}$
e	0.00 ± 0.00
M_* ($M_\odot$)	$1.561^{+0.045}_{-0.079}$
Spectral type	A5
H band Magnitude	7.5
[Fe/H]	0.1 ± 0.2 [c]

[a]Values from Kovács et al. (2013) except where otherwise noted.

[c]From Cameron et al. (2010).

light curve we ran three MCMC chains with 100,000 links for analysis, with an additional initial burn period of 25,000 links, and find that this is sufficient for convergence for all open parameters.

Our band-integrated time series is shown in various stages in Figure 5.4. While in practice we fit all parameters simultaneously, we show here the effect of removing systematics one at a time, and overplotting models for the slope, stellar oscillations, and finally the eclipse model itself on iterative versions of the residuals. Due to the offset in the phase of the stellar oscillations from one visit to the next, we fit each visit separately. We find agreement for the two visits' eclipse depths at the $\approx 1.5\sigma$ level, and we take for our final, best fit eclipse depth the weighted mean of both visits. We tabulate our best fit eclipse depths for the band integrated light curves in Table 5.4.

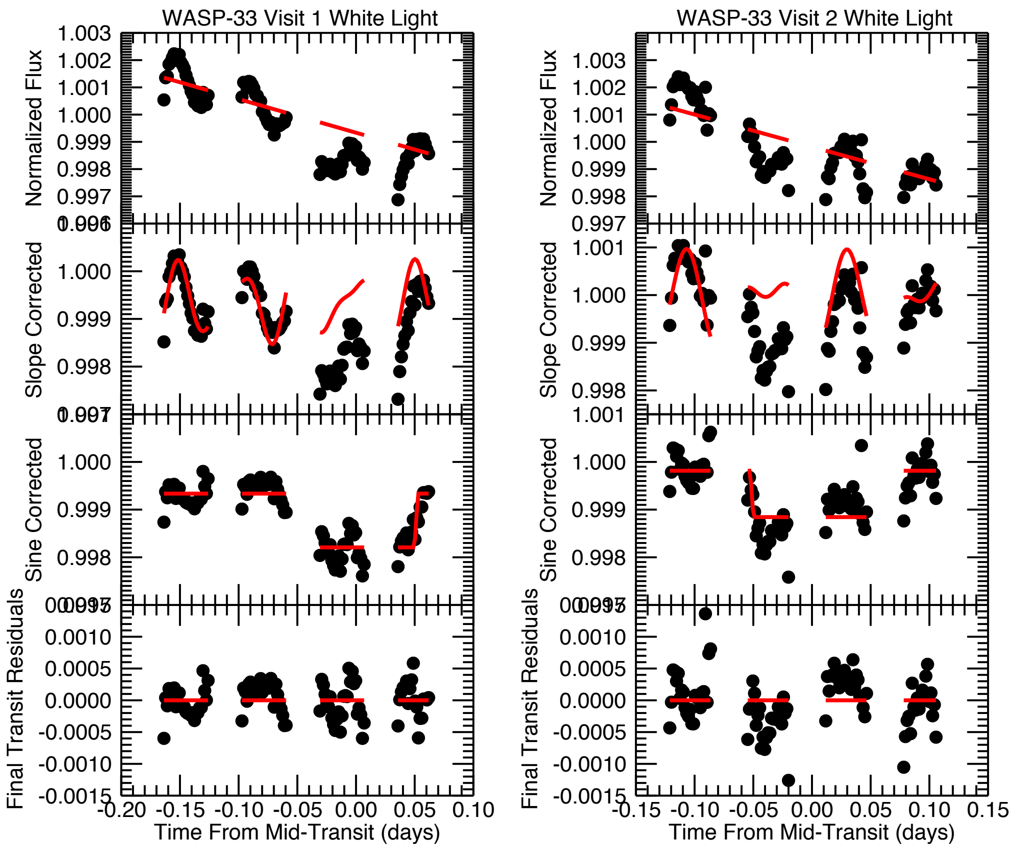

Figure 5.4 Band integrated curves for both visits in black with various model fits in red. From top to bottom, the plots show a) the normalized data with a model for the visit-long slope effect; b) the slope corrected data with a model for the stellar oscillations; c) the slope and oscillation corrected data with a model for the eclipse; d) the residuals for the full fit. Parameters are fit concurrently in MCMC, and are shown here in stages for clarification about the relative contributions of each parameter.

Table 5.4. Band Integrated Results

Visit	Eclipse Depth (%)
Visit 1	0.113 ± 0.009
Visit 2	0.097 ± 0.010
Combined	0.105 ± 0.006

5.4.3 Fitting the Spectrally Binned Light Curves

Once we determine an adequate fit to the band-integrated light curves, we use the residuals of these fits to remove the common-mode systematics in each individual spectral channel (or bin of channels) and thereby calculate a more accurate change in eclipse depth across the spectrum. We continue to fit each visit separately, but re-form the residuals to use the eclipse depth derived from the weighted mean of both visits. This affects the overall offset we see between visits. We utilize a strategy whereby we incorporate the residuals from the band-integrated light curve into our model for the eclipse of each channel or spectral bin, scaling the residuals by some factor which we leave open as an additional fitting parameter. Additionally, we introduce another set of systematic trends into the light curve model (with a scaling factor as a free parameter) based on our measurements of the horizontal shift of the spectrum on the detector over time. The scaling factor for this contribution is generally quite small, 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

$$LC_{final} = LC_{eclipse}(1 + X_1 * Res_{BI} + X_2 * Shift_X) \tag{5.1}$$

where $LC_{eclipse}$ is the light curve model calculated using the Mandel & Agol (2002) prescription (including a polynomial trend as well as two sine curves), Res_{BI} are the residuals from the band-integrated light curve, $Shift_X$ are the measured shifts in time spectrum for each exposure, and the X coefficients are scaling parameters determined through our MCMC fitting.

For the light curve for each channel or spectral bin we followed the band-integrated methods for fitting using MCMC. We locked the same orbital parameters as with the band-integrated light curve, and used a BIC to determine whether the sine amplitudes, second order polynomial coefficient, residuals scaling, and spectral shift scaling terms

Table 5.5. Visit 1: Open parameters for each bin as determined by BIC

Wavelength (μm)	Eclipse Depth	Out of Eclipse Flux	Linear Slope	Sine Amplitudes	Residuals Scaling	XShifts Scaling	2nd Order Coeffecient
1.16	✓	✓	✓			✓	✓
1.20	✓	✓	✓	✓		✓	
1.24	✓	✓	✓			✓	✓
1.28	✓	✓	✓			✓	
1.32	✓	✓	✓			✓	
1.37	✓	✓	✓			✓	
1.41	✓	✓	✓			✓	✓
1.46	✓	✓	✓			✓	✓
1.51	✓	✓	✓			✓	✓
1.56	✓	✓	✓			✓	✓
1.61	✓	✓	✓			✓	✓

Table 5.6. Visit 2: Open parameters for each bin as determined by BIC

Wavelength (μm)	Eclipse Depth	Out of Eclipse Flux	Linear Slope	Sine Amplitudes	Residuals Scaling	XShifts Scaling	2nd Order Coeffecient
1.16	✓	✓	✓			✓	
1.20	✓	✓	✓			✓	
1.24	✓	✓	✓			✓	
1.28	✓	✓	✓	✓		✓	
1.32	✓	✓	✓			✓	
1.37	✓	✓	✓			✓	
1.41	✓	✓	✓	✓		✓	
1.46	✓	✓	✓			✓	
1.51	✓	✓	✓			✓	
1.56	✓	✓	✓			✓	
1.61	✓	✓	✓			✓	

should be varied in our final analysis. The results of open model parameters for individual bins are shown in Tables 5.5 and 5.6.

We find that in general that the same terms make significant contributions for all the bins in a single visit. The light curves are typically best fit by a model including a linear slope term, unscaled band-integrated residuals and sine amplitudes, and a scaled version of the spectral shifts. A fit for a second-order polynomial term for the visit-long trend (as suggested by Stevenson et al. (2014)) passed the BIC for most bins in Visit 1, but did not pass for any bins in Visit 2. In all cases we choose the open parameters based on the BIC, on a bin-by-bin basis.

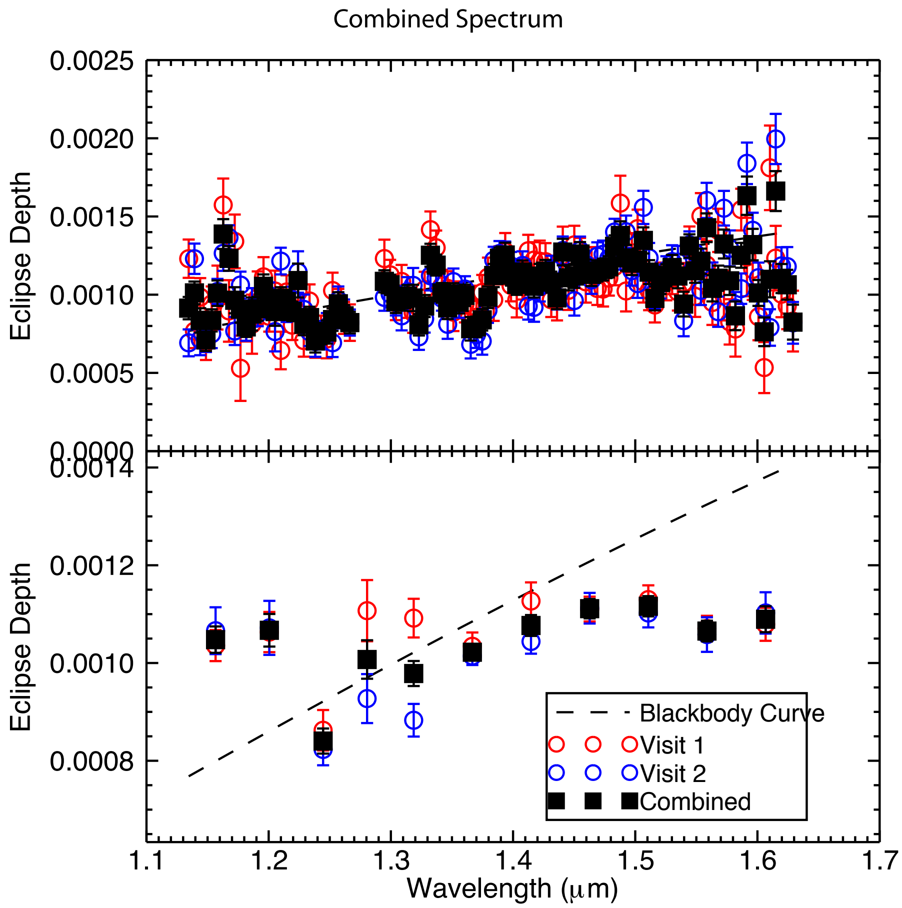

Figure 5.5 Top: Channel spectrum. Bottom: Binned spectrum. For both plots, Visit 1 is in red, Visit 2 in blue, and the combined visits are in black. Combined visits use a weighted mean. The largest discrepancies between visits are seen near the hydrogen line at 1.28 μm.

Finally, we use an uncertainty-weighted mean to combine both visits in our final stage of analysis, and these results are presented in Figure 5.5 and in Tables 5.7 (for the channels) and 5.8 (for the bins).

Table 5.7. Spectral Results for Channel Data

Wavelength (μm)	Eclipse Depth (%)	Wavelength (μm)	Eclipse Depth (%)	Wavelength (μm)	Eclipse Depth (%)
1.135	0.0914 ± 0.00714	1.304	0.0955 ± 0.0071	1.474	0.116 ± 0.00695
1.139	0.101 ± 0.00725	1.309	0.0967 ± 0.0069	1.478	0.118 ± 0.00678
1.144	0.0835 ± 0.00779	1.314	0.0991 ± 0.0070	1.483	0.131 ± 0.00699
1.149	0.0710 ± 0.00706	1.318	0.0979 ± 0.0084	1.488	0.138 ± 0.00903
1.153	0.0836 ± 0.00709	1.323	0.0808 ± 0.0068	1.493	0.124 ± 0.00817
1.158	0.101 ± 0.00815	1.328	0.0928 ± 0.0070	1.497	0.120 ± 0.00767
1.163	0.139 ± 0.00953	1.332	0.125 ± 0.00732	1.502	0.122 ± 0.00735
1.168	0.123 ± 0.00791	1.337	0.119 ± 0.00667	1.507	0.135 ± 0.00835
1.172	0.0964 ± 0.00833	1.342	0.101 ± 0.00632	1.511	0.114 ± 0.00764
1.177	0.0907 ± 0.00858	1.346	0.0915 ± 0.0071	1.516	0.0976 ± 0.00764
1.182	0.0794 ± 0.00641	1.351	0.100 ± 0.00698	1.521	0.108 ± 0.00736
1.186	0.0876 ± 0.0106	1.356	0.0951 ± 0.0070	1.525	0.112 ± 0.00708
1.191	0.0910 ± 0.00761	1.361	0.0996 ± 0.0068	1.530	0.117 ± 0.00805
1.196	0.105 ± 0.00758	1.365	0.0782 ± 0.0075	1.535	0.112 ± 0.00783
1.200	0.0933 ± 0.00637	1.370	0.0794 ± 0.0066	1.540	0.0940 ± 0.00791
1.205	0.0892 ± 0.00909	1.375	0.0839 ± 0.0069	1.544	0.131 ± 0.00747
1.210	0.0977 ± 0.00705	1.379	0.0983 ± 0.0069	1.549	0.115 ± 0.00763
1.215	0.0956 ± 0.00663	1.384	0.112 ± 0.00740	1.554	0.121 ± 0.00809
1.219	0.0893 ± 0.00623	1.389	0.123 ± 0.00683	1.558	0.143 ± 0.00917
1.224	0.109 ± 0.0108	1.394	0.125 ± 0.00786	1.563	0.104 ± 0.00839
1.229	0.0811 ± 0.00697	1.398	0.113 ± 0.00673	1.568	0.111 ± 0.00805
1.233	0.0854 ± 0.00693	1.403	0.107 ± 0.00669	1.573	0.132 ± 0.00924
1.238	0.0704 ± 0.00736	1.408	0.115 ± 0.00730	1.577	0.109 ± 0.00905
1.243	0.0751 ± 0.00683	1.412	0.109 ± 0.00663	1.582	0.0865 ± 0.00919
1.248	0.0761 ± 0.00736	1.417	0.105 ± 0.00722	1.587	0.125 ± 0.00796
1.252	0.0841 ± 0.00718	1.422	0.112 ± 0.00779	1.591	0.163 ± 0.0123
1.257	0.0942 ± 0.00736	1.427	0.115 ± 0.00716	1.596	0.132 ± 0.0101
1.262	0.0865 ± 0.00687	1.431	0.108 ± 0.00670	1.601	0.101 ± 0.00829
1.266	0.0822 ± 0.00689	1.436	0.0979 ± 0.0071	1.606	0.0763 ± 0.00916
1.271	0.0452 ± 0.0101	1.441	0.127 ± 0.00849	1.610	0.110 ± 0.0115
1.276	0.0181 ± 0.00960	1.445	0.111 ± 0.00811	1.615	0.166 ± 0.0128
1.281	0.229 ± 0.0252	1.450	0.112 ± 0.00760	1.620	0.110 ± 0.00930

Table 5.7 (cont'd)

Wavelength (μm)	Eclipse Depth (%)	Wavelength (μm)	Eclipse Depth (%)	Wavelength (μm)	Eclipse Depth (%)
1.285	0.164 ± 0.0184	1.455	0.126 ± 0.00764	1.624	0.106 ± 0.00955
1.290	0.107 ± 0.00934	1.460	0.115 ± 0.00742	1.629	0.0825 ± 0.0112
1.295	0.109 ± 0.00708	1.464	0.113 ± 0.00738		
1.299	0.106 ± 0.00647	1.469	0.115 ± 0.00681		

Table 5.8. Spectral Results for Binned Data

Wavelength (μm)	Eclipse Depth (%)
1.155	0.105 $\pm$ 0.00270
1.199	0.107 $\pm$ 0.00334
1.243	0.0840 $\pm$ 0.00257
1.279	0.101 $\pm$ 0.00393
1.318	0.0979 $\pm$ 0.00256
1.366	0.102 $\pm$ 0.00157
1.414	0.108 $\pm$ 0.00214
1.462	0.111 $\pm$ 0.00198
1.510	0.112 $\pm$ 0.00207
1.558	0.107 $\pm$ 0.00212
1.606	0.109 $\pm$ 0.00266

5.4.4 Error Analysis

Our errors are initially drawn from the MCMC posterior probability distributions. In order to estimate the impact of our red noise, we use a modified version of the residuals permutation method (Gillon et al., 2007). In addition to shifting the residuals (the "prayer-bead" permutation), we also invert our residuals (multiplying by -1) and/or reverse them in time, and then shift each of those permutations via the "prayer-bead" method. This yields 4 $\times$ N permutations, where N is the number of exposures. We deemed this extra step useful because of the otherwise limited number of possible permutations, which did not yield clear results from a traditional residuals permutation analysis. For each channel or bin of channels, we use whichever is higher, the uncertainty from MCMC or residuals permutation. For Visit 1, we find that uncertainties from residuals permutation are on average 1.47 times higher than uncertainties from MCMC, while for Visit 2 uncertainties from residuals permutation are 1.20 times higher.
We also compare the photon noise to the RMS of our white light, channel, and binned data, and the measured MCMC + residuals permutation uncertainty to a predicted

Table 5.9. Error Analysis. We compare the photon-noise and rms statistics for each source with the predicted and actual uncertainties for the overall eclipse depth, both for the band-integrated time series and the binned spectral channels. The difference between the predicted uncertainty and the simulated-data uncertainty from MCMC reflects the incomplete coverage over the light curve, while the difference between the actual data and the simulated data reflects the impact of additional noise beyond the photon noise.

Parameters	Visit 1	Visit 2
Data points during eclipse	38	47
Data points out of eclipse	81	72
Channels		
Photon noise (ppm)	419.	420.
RMS of residuals (ppm)	440.	437.
Predicted[1] σ_{ed} (ppm)	79.0	78.0
σ_{ed} from MCMC+RP Data (ppm)	203.	121.
RMS/photon noise	1.05	1.04
MCMC Data/Pred.	2.57	1.55
0.042-Micron Bins (11 Total)		
Photon noise (ppm)	181.	162.
RMS of residuals (ppm)	191.	187.
Predicted[1] σ_{ed} (ppm)	34.0	30.1
σ_{ed} from MCMC+RP Data (ppm)	76.2	52.6
RMS/photon noise	1.07	1.17
MCMC Data/Pred.	2.33	1.79
Band Integrated Time Series		
Photon noise (ppm)	47.3	46.6
RMS of residuals (ppm)	272.	331.
Predicted[1] σ_{ed} (ppm)	9.50	9.58
σ_{ed} from MCMC+RP Data (ppm)	112.	90.5
RMS/photon noise	5.75	7.10
MCMC Data/Pred.	11.8	9.45

[1]Calculated from the photon noise and the number of points during eclipse

eclipse depth uncertainty. This prediction is based on the photon noise (or RMS) and the number of exposures in eclipse versus out of eclipse. We present our results in Table 5.9. In general we find that while a substantial amount of red noise remains in our band integrated light curves after removal of our best fit models, by incorporating this red noise in our residuals scaling process we are able to closely approach the photon noise limit for our channels and bins. We find an RMS $\sim$1.05 times the photon noise. We find that our measured eclipse depth uncertainties are between $\sim$1.5-2.5 times the predicted uncertainty, and we ascribe this to the contributions from the uneven sampling of HST orbits and the residual red noise in the light curves. In previous studies we found that the lack of complete and evenly spaced temporal coverage during the eclipse is likely the cause of this failure to meet the predicted uncertainty, even for photon-limited results. In this same comparison of measured versus predicted uncertainties, we note that Visit 1 is further from the predicted uncertainty than Visit 2. Visit 1 has substantially fewer points post-eclipse than Visit 2, which can be detrimental to the uncertainty on eclipse fitting, and additionally Visit 1 has a higher rate of spectral drift. While we perform corrections for this drift, it remains an additional source of uncertainty. Given these factors, we feel confident that the uncertainty we measure accurately reflects the sources of uncertainty in the data.

5.5 Discussion

The hot Jupiter WASP-33 b is one of the most irradiated hot Jupiters known and hence, is among the most favorable candidates to host a thermal inversion in its dayside atmosphere. Studies in the past have suggested that extremely irradiated hot Jupiters should host thermal inversions due to strong absorption of incident stellar light by absorbers such as TiO and VO (Hubeny et al. (2003), Fortney et al. (2008)). While Spiegel et al. (2009) have suggested that TiO and VO may not be aloft in some hot Jupiter atmospheres due to downward drag by gravitational settling and condensation overtaking upward vertical mixing, the extreme irradiation of WASP-33 b should still be a

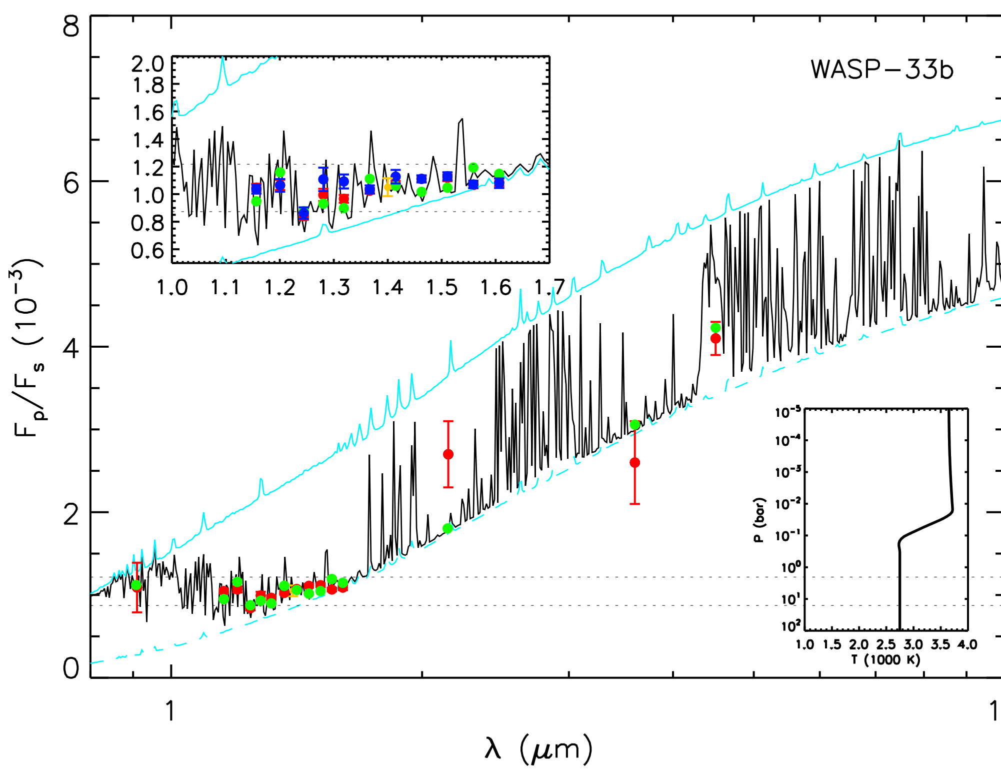

Figure 5.6 Transmission spectrum for Visit 1 (blue) with the weighted mean (red) shown for the WFC3 data points. Spitzer and z-band data also shown in red. Green points are binned model points. This model, shows an oxygen-rich atmosphere with a thermal inversion, and provides a significantly better fit to the data than a model with no inversion, or a blackbody model.

secure case for TiO to be abundant and cause a thermal inversion. On the other hand, Knutson et al. (2010) proposed that the formation of inversions may be correlated with chromospheric activities of the hot stars, implying that hot Jupiters orbiting active stars are less likely to host thermal inversions; though their study did not include A-stars like WASP-33. Finally, Madhusudhan et al. (2011a) and Madhusudhan (2012) suggested that high C/O ratios could also deplete inversion-causing compounds such as TiO and VO in hot Jupiters, thereby precluding the formation of thermal inversions. Thus, given its extreme thermal emission, WASP-33 b presents a valuable opportunity to constrain the various hypotheses regarding thermal inversions in hot Jupiters. Previously reported photometric observations from Spitzer and ground-based facilities have been unable to conclusively constrain the presence of a thermal inversion in WASP-33 b (Deming et al. (2013); de Mooij et al. (2013)).

We use the observed emission spectrum of WASP-33 b to constrain the thermal and chemical properties of the dayside atmosphere of the planet. We model the planetary atmosphere and retrieve its properties using the retrieval technique of Madhusudhan et al. (2011a) and Madhusudhan (2012). The model computes line-by-line radiative transfer for a plane-parallel atmosphere with the assumptions of hydrostatic equilibrium and global energy balance, as described in Madhusudhan & Seager (2009). The composition and pressure-temperature (P-T) profile of the dayside atmosphere are free parameters in the model. The model includes all the major opacity sources expected in hot Jupiter atmospheres, namely H_2O, CO, CH_4, CO_2, C_2H_2, HCN, TiO, VO, and collision-induced absorption (CIA) due to H_2-H_2, as described in Madhusudhan (2012). Our molecular line lists are obtained from Freedman et al. (2008), Freedman (personal communication, 2009), Rothman (2005), Karkoschka & Tomasko (2010), and Karkoschka (personal communication, 2011). Our CIA opacities are obtained from Borysow et al. (1997) and Borysow (2002). A Kurucz model Castelli & Kurucz (2004) is used for the stellar spectrum, and the stellar and planetary parameters are adopted from Cameron et al. (2010). We use our WFC3 observations together with previously published photometric data (Deming et al. (2013); de Mooij et al. (2013); Smith et al. (2011)) to obtain joint constraints on the chemical composition and temperature structure of the planet. We explore the model parameter space using a Markov Chain Monte Carlo algorithm (Madhusudhan et al., 2011a) and derive regions of model space that best explain the data. Our model space includes models with and without thermal inversions, and models with oxygen-rich as well as carbon-rich compositions.

We find that the sum-total of observations are best explained by an oxygen-rich atmosphere with a thermal inversion (see Figure 5.6). Previous photometric observations were consistent with two distinct models (Deming et al., 2013): (a) a model with oxygen-rich composition with a strong thermal inversion, and (b) a model with a carbon-rich composition but with no thermal inversion. In our current work, we use our WFC3 observations to break degeneracies between models and provide strong evidence for a temperature inversion caused by TiO.

TiO is responsible for the spectral signatures at wavelengths short of 1.2 μm, and is supported by both the bluest two bins in our WFC3 data, as well as the z band photometric point from Smith et al. (2011). This strongly supports not just evidence for a temperature inversion, but also argues that the temperature inversion is due to TiO, and further represents the first observational evidence for TiO in an exoplanet atmosphere.

Previous studies, both theoretical and observational, have suggested that the hottest exoplanets may be the most inefficient at redistributing heat to their night sides. This study would confirm those findings, as our best-fit model for WASP-33 b, the hottest known exoplanet, has a low day-night redistribution (7%), compared to the non-inverted model, which has a roughly 50% redistribution efficiency.

We allow a constant offset on the WFC3 spectrum as a free parameter in our model fits. The best fit inverted model has a chi-square of 113, and the best-fit non-inverted model has a chi-square of 365. A blackbody (BB) spectrum has a chi-square of 380, implying that the O-rich, inverted model provides the best fit to all the data and that the spectrum is not a blackbody.

5.6 Conclusion

In this paper we present our analysis of WFC3 observations of two occultations of WASP-33 b, a hot Jupiter orbiting a δ-Scuti star. We reduce and analyze the spectroscopic time series for both visits, and correct for stellar oscillations of the star, as well as for motion of the target on the detector. We bin our spectrum, and achieve an RMS $\sim$1.05 times the photon noise. We compare our final emission spectrum to atmospheric models testing a range of carbon to oxygen ratios and temperature profiles, and find good evidence for WASP-33 b having an oxygen-rich atmosphere that hosts a temperature inversion. We also present the first observational evidence for TiO in an exoplanet atmosphere. This is consistent with and improves upon previous observations that could not discern between competing models. The spectrum is inconsistent with a blackbody or non-inverted atmosphere.

Chapter 6: Paper 3: Re-analysis of Selected Occultations

6.1 Overview

Project 3 analyzes a collection of occultations observed with WFC3 as part of the same large observing campaign by Drake Deming (HST General Observer Program 12181) used in Paper 1 (with additional observations of WASP-12 from Mark Swain's observations). While all of these targets have previously been published by various different research groups (Swain et al. (2013), Ranjan et al. (2014)), the lack of consensus in the past between groups makes a uniform analysis valuable. The goal here is to present consistent reduction and analysis methods in the hunt for molecular features and evidence of temperature inversions. Comparison to the results already in the literature informs the reliability of different analysis methods.

All three planets orbit extremely close to their host stars, and are to some extent bloated or over-inflated. They all orbit G type stars, though with a range of metallicities. In Project 1, we were unable to distinguish in the transmission spectrum between oxygen-rich or carbon-rich models for WASP-12 b, though WASP-19 b showed marginal support for a carbon-rich interpretation, supporting past observational studies. Bean et al. (2013) found that ground-based emission measurements of WASP-19 b indicated an isothermal atmosphere, with their best fit given by a carbon-rich model without a temperature inversion. Ranjan et al. (2014) found a featureless spectrum for CoRoT-1 b in transit. They further found no evidence for a temperature inversion in the emission spectrum of WASP-4 b.

As in Paper 1, these observations were conducted in WFC3's stare mode, instead of the spatial scan mode used for the observations taken for Paper 2. The reduction and analysis procedures are largely unchanged from Project 1, though certain improvements have been

Table 6.1. Stellar and Orbital Parameters Used For Model Fitting and Comparison

Parameters	CoRoT-1 b[a]	WASP-4 b[b]	WASP-12 b[c]	WASP-19 b[d]
Period (days)	1.51	1.34	1.09	0.789
i (°)	85.1 ± 0.5	88.8 ± 0.61	82.5 ± 0.8	79.5 ± 0.500
R_p/R^*	0.135 ± 0.0095	0.122 ± 0.0297	0.117 ± 0.00068	0.139 ± 0.0457
T_c	55943.293	55527.672	55663.199	55743.532
a/R^*	4.920 ± 0.235	4.322 ± 1.05	3.03 ± 0.0220	3.57 ± 0.0460
e	0.0	0.0	0.0447 ± 0.00430	0.00770 ± 0.00680
ω (°)	0.0	0.0	94.4 ± 0.0300	43.0 ± 67.0
Semi-major axis (AU)	0.0254 ± 0.0004	0.02312 ± 0.00033	0.02309 ± 0.00096	0.01616 ± 0.00024
M_* ($M_\odot$)	0.95 ± 0.15	0.93 ± 0.05	1.38 ± 0.18	0.904 ± 0.040
$M_p \times \sin i$ (M_J)	1.03 ± 0.12	1.237 ± 0.062	1.378 ± 0.181	1.114 ± 0.04
Spectral type	G0V	G8	G0	G8V
H-band Magnitude	12.218	10.842	10.228	10.602
[Fe/H]	0.06 ± 0.07	0.03 ± 0.09	0.3 ± 0.1	0.02 ± 0.09

[a]Values from Barge et al. (2008).

[b]Values from Hoyer et al. (2013).

[c]Values from Southworth et al. (2012).

[d]Values from Lendl et al. (2013).

made to the pipeline as the work has matured and new insights were gained during Paper 2.

6.2 Observations

The observations of WASP-4 and WASP-19 analyzed here were conducted in November 2010 and June 2011, respectively, while the observations of WASP-12 were obtained in April of 2011. CoRoT-1 was observed in January and February of 2012. Precise observation dates and exposure information are listed in Table 5.1. The observations were taken with the G141 grism on WFC3's infrared channel, providing slitless spectra covering the wavelength range 1.1 μm to 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

Table 6.2. Observation Parameters

	CoRoT-1 (Visit 1)	CoRoT-1 (Visit 2)	CoRoT-1 (Visit 3)	WASP-4	WASP-12	WASP-19
Date of Observation	2012-2-05	2012-1-27	2012-1-17	2010-11-27	2011-04-15	2011-06-09
Integration Time	100.65	100.65	100.65	36.02	7.62	21.66
Subarray Mode	128	128	128	128	256	128
CALWF3 version	3.1.2	3.1.2	3.1.2	3.1.2	3.1.2	3.1.2
NSamp	16	16	16	7	11	5
Timing Sequence	SPARS10	SPARS10	SPARS10	SPARS10	RAPID	SPARS10
Peak Pixel Value[1]	6.3×10^4	7.4×10^4	7.5×10^4	9.5×10^4	3.8×10^4	7.3×10^4

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

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 occultation while including some out-of-eclipse data as well. Observations can be conducted on subarrays of the full detector in order to decrease overhead time. Multiple reads are taken per exposure, and the timing of these reads is a preference set by the observer.

Observations of CoRoT-1 were taken using the 128 x 128 sub-array with 16 non-destructive reads per exposure and sampled using the SPARS10 sampling sequence. This resulted in a total integration time of 100.65 seconds per exposure and 25 exposures per orbit, with a total of 90 exposures taken over 4 HST orbits. Observations of WASP-4 were taken using the 128 x 128 sub-array with 7 non-destructive reads per exposure and sampled using the SPARS10 sampling sequence. This resulted in a total integration time of 36.02 seconds per exposure and 55 exposures per orbit, with a total of 260 exposures taken over 5 HST orbits. Observations of WASP-12 were taken using the 256 x 256 sub-array with 11 non-destructive reads per exposure and sampled using the RAPID sampling sequence. This resulted in a total integration time of 7.62 seconds per exposure and 95 exposures per orbit, with a total of 476 exposures taken over 5 HST orbits. Observations of WASP-19 were taken using the 128 x 128 sub-array with 5 non-destructive reads per exposure and sampled using the SPARS10 sampling sequence. This resulted in a total integration time of 21.66 seconds per exposure and 69 exposures per orbit, with a total of 266 exposures taken over 4 HST orbits.

6.3 Data Reduction

We use the `.ima` files comprising the individual detector reads as our initial data product. We combine individual detector reads into single images, correct for bad pixels, subtract a background region, select an extraction box, and collapse the box to achieve 1D spectra. See §4.3 for details on individual steps. We simplify our method for calculating the wavelength solution and wavelength-dependent flat field for each set of observations by using the modified coefficients from Wilkins et al. (2014).

For all three targets, we trim the spectrum between 1.13 and 1.63 microns, leaving 106 channels for analysis.

6.4 Analysis and Results

We observe the instrumental systematics seen in most previous stare-mode 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. The systematics are less clear in the lower-signal eclipse light curves than in the transits observed for Project 1, and differences in the morphology of the shape from target to target may be seen in

Figure 6.1. As in Project 1, 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. We also test a quadratic fit, as advocated by Stevenson et al. (2014), but find that our uncertainties are larger, making the improvement in rms statistically insignificant. 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 §4.4.3 and §4.4.4 for a detailed description), and with this fitting strategy we are able to incorporate the first noisier orbit.

In addition to the hook effect, we also measure effects due to the shifting of the source on the detector with time, in both 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, 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 1D 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

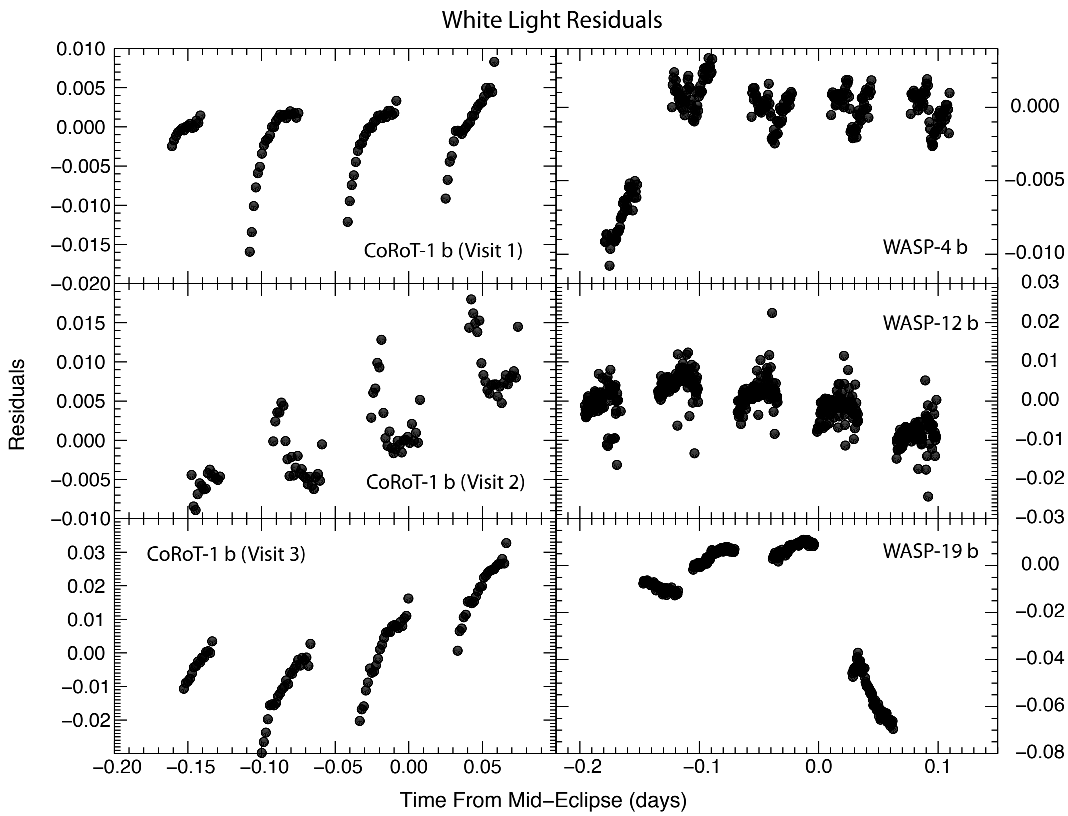

Figure 6.1 White light residuals depicting the systematics inherent to the data. The morphology of these systematics varies widely depending on observational strategies utilized.

was more difficult to calculate, since the sensitivity function of the grism does not allow for an analytical fit. During Project 1, we 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. However, during Project 2, we found small yet measurable changes in the morphology of the spectral shifts with wavelength. This prompted us to measure changes in the spectrum only across the wavelength region used by each bin. The revised method uses for its template an average of the exposures in the time series immediately preceding and following eclipse. We interpolate the averaged, template spectrum onto a wavelength grid shifted in either direction up to a pixel and a half, stepping in 0.001 pixel increments, and save each shifted spectrum. For each exposure, we compare the observed spectrum with each shifted template spectrum, and calculate the rms. The shift corresponding to the lowest rms is saved, and we can use these xshifts in place of the above method. For each bin of channels, we apply a mask in wavelength space, whereby the rms is calculated using only the channels contributing to that bin. This therefore results in a different set of xshifts for each bin. The spectral shifts may be seen in Figure 6.2.
All of the variables change relatively coherently within an orbit, and then reset at the beginning of the next orbit.

6.4.1 Background Source Correction

We examined each object carefully for background sources, in particular WASP-12, already known to have a nearby companion. We follow the same methodology for corrections used in Project 1 (see 4.4.2 for details). Briefly, we use the PSF of an uncontaminated source, in this case WASP-19, as a template. We position one template

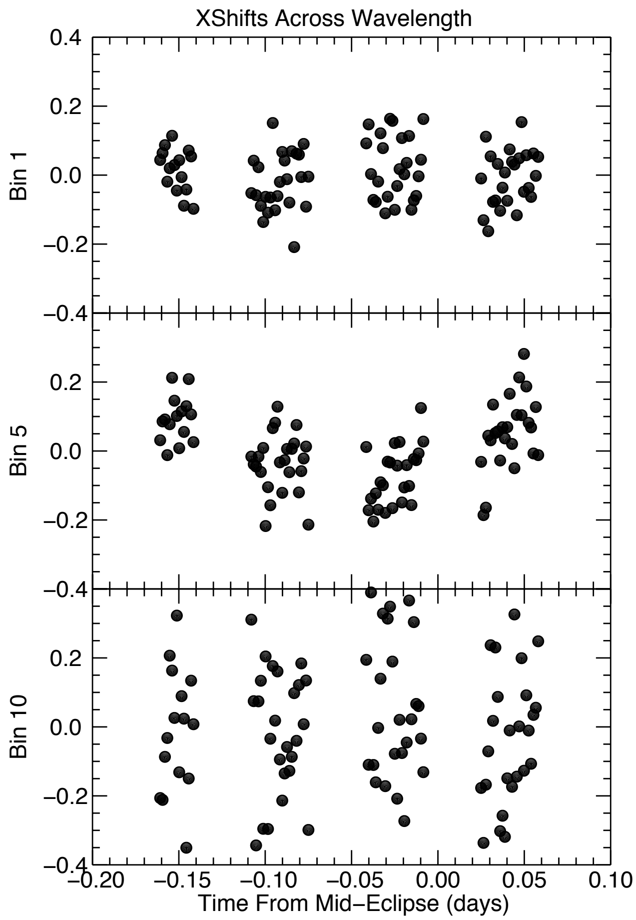

Figure 6.2 Xshifts for various bins for CoRoT-1's third visit. Note that the best-fit shifts vary across wavelength.

PSF under the peak of the target and one under the peak of the contaminating source, and jointly scale them to match the observed PSF. We then subtract the modeled flux contribution of the contaminating source from the spectrum before proceeding further in the analysis process. See Figure 6.3 for an individual channel fit.

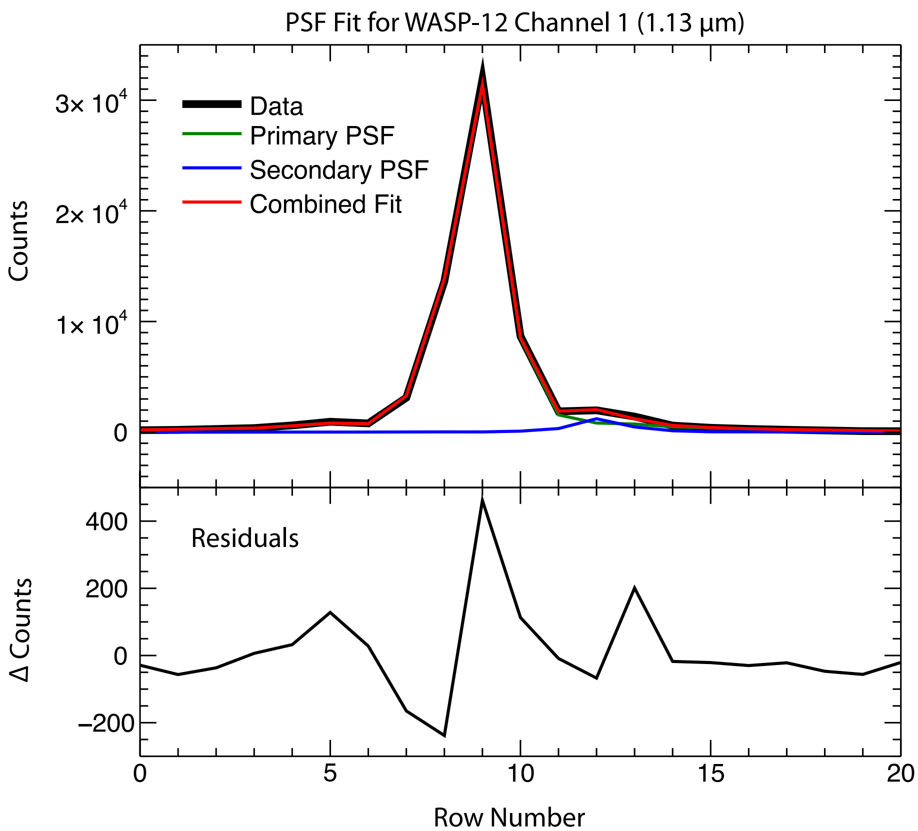

Figure 6.3 Top: Data and the best-fit PSF model for a single channel for WASP-12, using the template-PSF method. The data are shown in black, the fit to the main peak is shown in green, and the fit to the contamination peak is shown in blue; the combined fit is shown in red. Bottom: Remaining residuals after removing the model; the remaining flux under the region of contamination was used as the uncertainty in the contamination flux.

6.4.2 Band-Integrated Transit Curve Fitting

As in the previous projects, we assume that the ramp systematic is consistent across wavelengths, and so we determine the band-integrated transit curve parameters on the divide-oot, white light curve first, in order to use the residuals of the best fit and the uncorrected light curve as a template for the systematics. All of the orbital parameters in our transit light curve model were locked to the literature values (see Table 6.1), since we are only analyzing single transits and lack full-transit coverage. 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 6.4, with the best-fit transit curve overlaid.

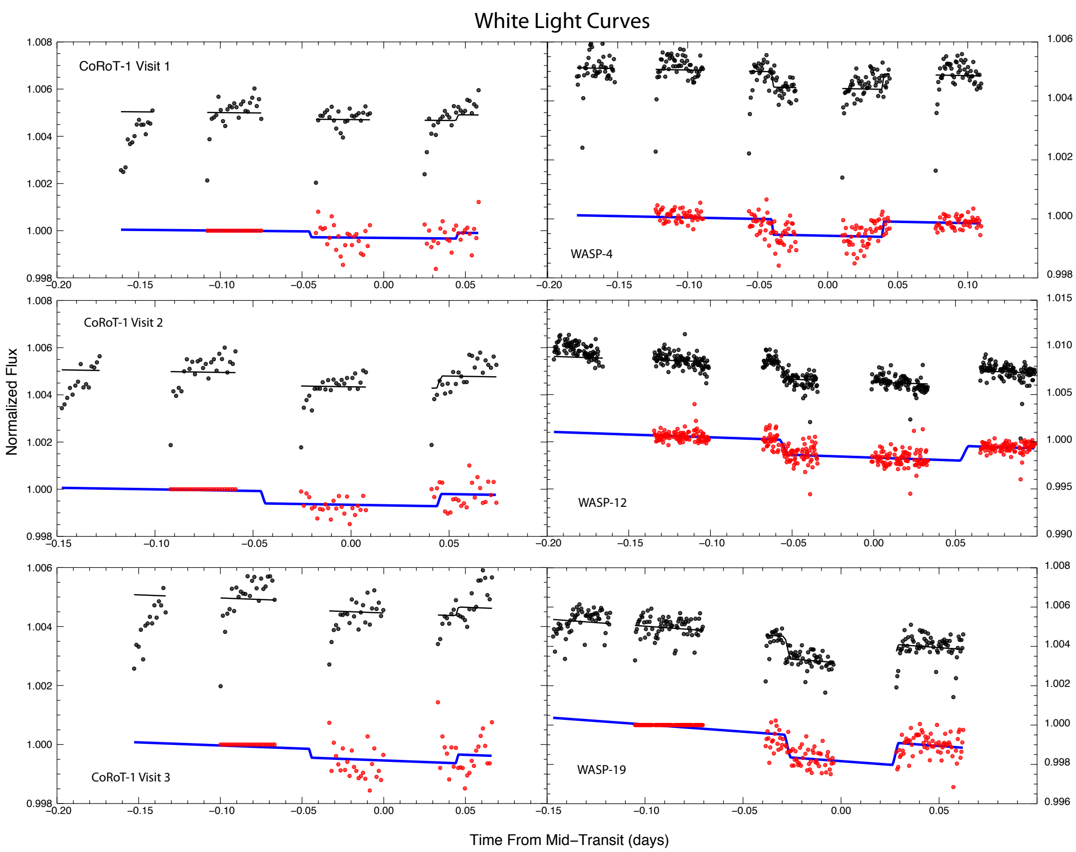

Figure 6.4 White light curves for all targets. Uncorrected data are shown in black, systematics-corrected (divide-oot) light curves are shown in red, and a best-fit transit model is shown in blue. Artificial offsets have been added to the uncorrected data for plotting purposes.

6.4.3 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. (2012a) and Stevenson et al. (2014) (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

$$LC_{final} = LC_{transit} * (a + bt + C_1 * \mathrm{Res}_{BI} + C_2 * \mathrm{Shift}_y + C_3 * \mathrm{Shift}_x) \quad (6.1)$$

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 fit for the best model using MCMC, as above. We locked the same parameters as with the band-integrated light curve, and, as in previous projects, use a BIC to determine whether the systematics scaling terms should be left open. The results of the BIC can be seen in Table 6.3.

Results for all four targets are shown in Figure 6.5. We show some deviation from the

Table 6.3. Open parameters for each bin as determined by BIC

Wavelength	CoRoT-1 (Visit 1)	CoRoT-1 (Visit 2)	CoRoT-1 (Visit 3)	WASP-4	WASP-12	WASP-19
1.16	YShifts Xshifts	XShifts	Resids XShifts	XShifts	Resids XShifts	Resids YShifts XShifts
1.21	XShifts	XShifts	XShifts	XShifts	Resids XShifts	Resids XShifts
1.26	Resids XShifts	XShifts	XShifts	Resids YShifts XShifts	Resids XShifts	Resids XShifts
1.31	XShifts	XShifts	XShifts	XShifts	Resids XShifts	Resids XShifts
1.36	XShifts	XShifts	XShifts	Resids YShifts XShifts	Resids XShifts	Resids YShifts XShifts
1.41	Resids XShifts	XShifts	YShifts	Resids YShifts XShifts	Resids XShifts	Resids YShifts XShifts
1.46	XShifts	XShifts	XShifts	XShifts	Resids XShifts	Resids YShifts XShifts
1.51	XShifts	XShifts	XShifts	XShifts	Resids XShifts	Resids XShifts
1.56	XShifts	Resids XShifts	XShifts	XShifts	Resids XShifts	Resids XShifts
1.61	YShifts XShifts	XShifts	XShifts	XShifts	Resids YShifts	Resids YShifts XShifts

blackbody curves, particularly for CoRoT-1 b. Other targets do not show statistically significant deviations from a flat spectrum.

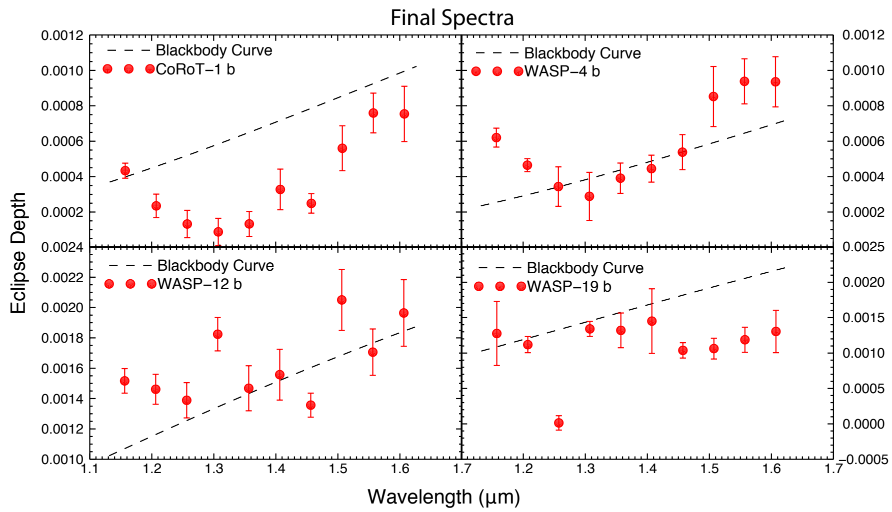

Figure 6.5 Final spectra for all four targets in red with error bars. Predicted blackbody curves shown as a dotted black line.

6.4.4 Uncertainty Analysis

We derive initial uncertainties from the widths of our MCMC posteriors. Additionally, we run a residuals permutation analysis, as described in Project 1, to account for uncertainty due to red noise that might be improperly captured by the MCMC uncertainties.

6.5 Discussion

WASP-4 b and WASP-12 b have been previously analyzed by Ranjan et al. (2014) and Swain et al. (2013), respectively. We compare our results to theirs, plus models, below. For WASP-4 b, we approach, but do not replicate, previous results. See Figures 6.6 and 6.7 for details. According to the models, our results more likely indicate a solar composition atmosphere with no inversion, or a simple blackbody. However, WASP-4 b also has extremely high flux rates on the detector, which likely indicates some amount of uncorrectable systematic response. We show results here for comparison's sake of the data reduction process, but we do not advocate interpretation of the final spectrum with any reliability.

For WASP-12 b, we achieve a similar spectrum to previous results (Figures **??** and 6.8), which indicate a likely carbon rich atmosphere, although this result is marginal.

For WASP-19 b, though we ran a full analysis on the data set, in the end we conclude that the sudden shift during the final orbit (see Figure 6.1) renders the data unusable for analysis. We do not address the final emission spectrum for this planet.

For a comparison of all our occultations, see Figure 6.9. An overall offset has been applied. All three spectra show strong similarities, though the scatter is larger for WASP-12 b. A slight increase may be seen at 1.4μm, and in all cases the flux rises steeply at the long wavelength end of the range.

WASP-12 b shows a flatter spectrum than the other two targets, though the uncertainties are larger, in part due to the contaminating source. CoRoT-1 b and WASP-4 b show very similar spectral shapes. This is reasonable since they occupy similar regions of parameter space in terms of size and temperature, as well as distance from their stars and type of stars.

The increase at short wavelengths may be indicative of TiO, which could give rise to temperature inversions in these planets. The similarity in spectral shape across targets is a good sign for the robustness of our methodology, and points to a consistency in reduction that has been lacking until now.

6.6 Conclusion

We reduce and reanalyze data for four hot Jupiter occultations observed with WFC3, though we present results for only three of these planets. We carefully examine the spectroscopic time series for all visits, and correct for instrument systematics as well as for motion of the target on the detector. We bin our spectra, and compare both to models

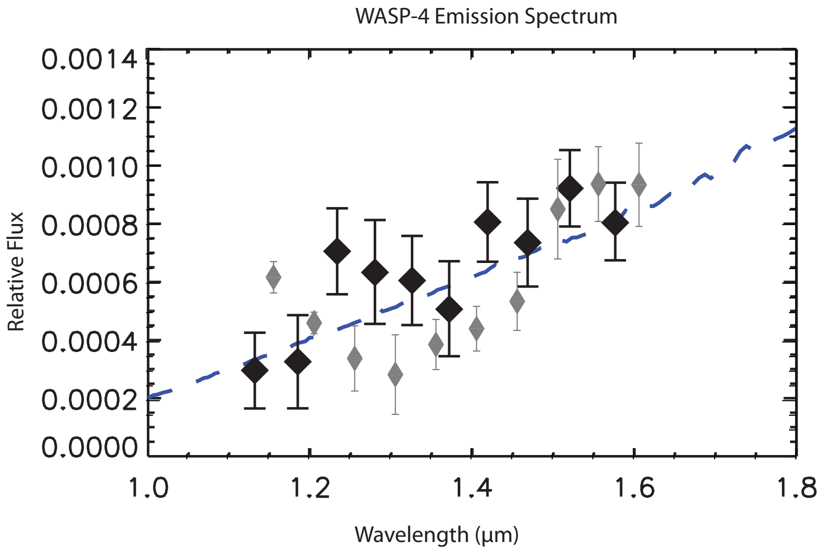

Figure 6.6 WASP-4 b emission spectrum: results from Ranjan et al. (2014) in black, with our data shown in grey. Our final spectrum is similar to previous reduction efforts, though variable at the $\tilde{1}\sigma$ level.

and to previously published results. While we are unable to break certain model degeneracies between carbon and oxygen rich models, or models with or without temperature inversions, we nonetheless find that we are generally consistent with previous findings, and that our reduction strategies are robust.

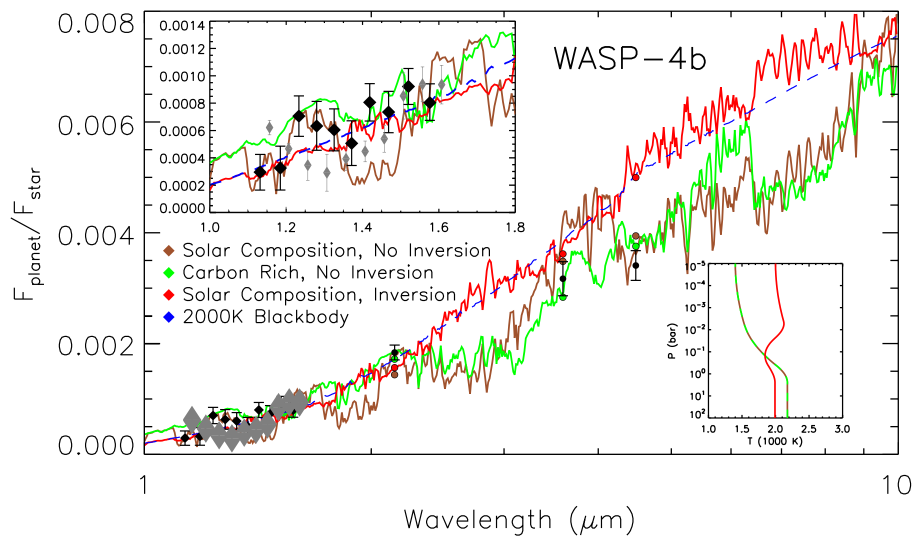

Figure 6.7 WASP-4 b emission spectrum: results from Ranjan et al. (2014) in black, with our data shown in grey. Also shown here are models for various types of atmospheres. Out results are not precise enough to distinguish between models based on different chemical compositions or structures.

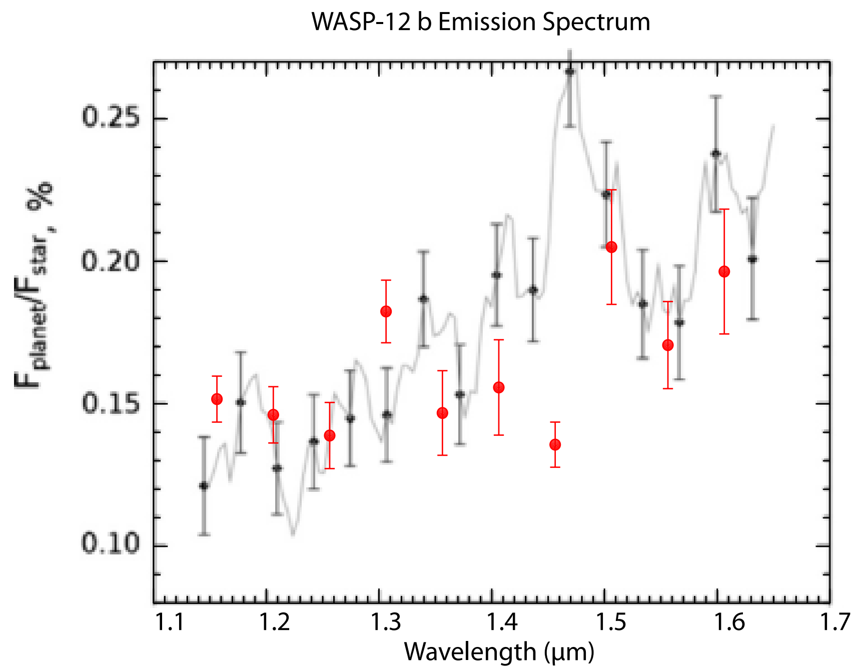

Figure 6.8 WASP-12 b emission spectrum: results from Swain et al. (2013) shown in grey, with our data shown in red. The two reduction methods are consistent within uncertainties across most of the spectrum. It is unclear what causes the discrepancy at 1.45 μm.

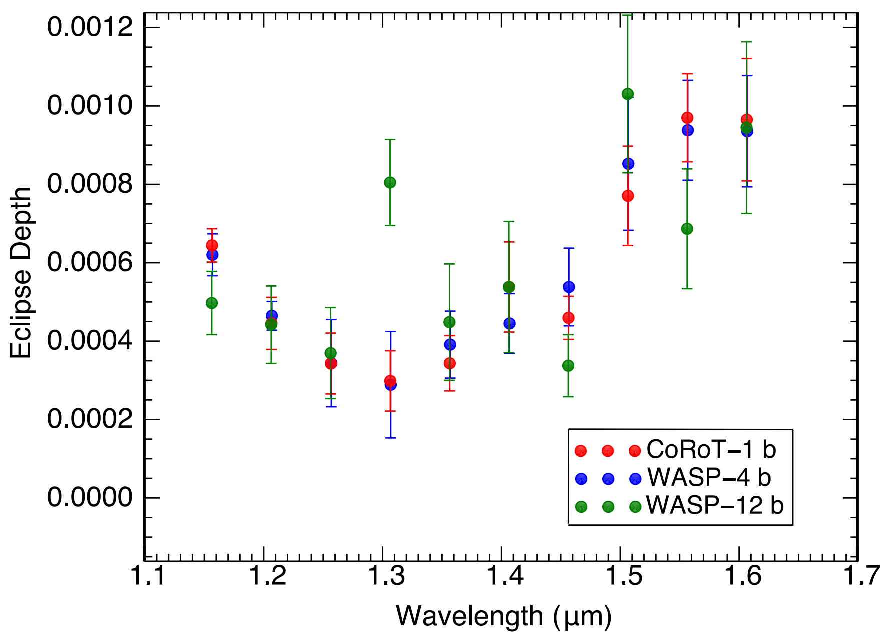

Figure 6.9 Final emission spectrum for all three targets, offset in vertical space to show similarities in spectral morphology. The three planets show strong similarities, especially in the rise at short wavelengths, likely due to features caused by TiO in the planets' atmospheres.

Bibliography

Adams, E. R., Dupree, a. K., Kulesa, C., & McCarthy, D. 2013, The Astronomical Journal, 146, 9

Albrecht, S., Winn, J. N., Johnson, J. A., et al. 2012, The Astrophysical Journal, 757

Anderson, D. R., Hellier, C., Gillon, M., et al. 2010, The Astrophysical Journal, 709

Anderson, D. R., Smith, A. M. S., Lanotte, A. A., et al. 2011, Monthly Notices of the Royal Astronomical Society, 416

Anderson, D. R., Smith, A. M. S., Madhusudhan, N., et al. 2013, Monthly Notices of the Royal Astronomical Society, 430

Barge, P., Baglin, A., Auvergne, M., et al. 2008, Astronomy and Astrophysics, 482, L17

Barman, T. 2007, The Astrophysical Journal, 661, L191

Barman, T. S. 2008, The Astrophysical Journal, 676, L61

Barman, T. S., Macintosh, B., Konopacky, Q. M., & Marois, C. 2011, The Astrophysical Journal, 733, 65

Batygin, K., & Stevenson, D. J. 2010, The Astrophysical Journal, 714, L238

Bayliss, D. D. R., Winn, J. N., Mardling, R. A., & Sackett, P. D. 2010, The Astrophysical Journal, 722, L224

Bean, J. L., Désert, J.-M., Seifahrt, A., et al. 2013, The Astrophysical Journal, 771, 108

Bechter, E. B., Crepp, J. R., Ngo, H., et al. 2014, , 788, 2

Bergfors, C., Brandner, W., Daemgen, S., et al. 2013, Monthly Notices of the Royal Astronomical Society, 428

Berta, Z. K., Charbonneau, D., Désert, J.-M., et al. 2012, The Astrophysical Journal, 747, 35

Bertin, E., & Arnouts, S. 1996, Astronomy and Astrophysics Supplement, 117

Bodenheimer, P., Lin, D. N. C., & Mardling, R. A. 2001, The Astrophysical Journal, 548

Bonnefoy, M., Lagrange, A.-M., Boccaletti, A., et al. 2011, Astronomy & Astrophysics, 528, L15

Borysow, A. 2002, A&A, 390

Borysow, A., Jorgensen, U. G., & Zheng, C. 1997, A&A, 324

Brogi, M., Snellen, I. a. G., de Kok, R. J., et al. 2012, Nature, 486, 502

Burke, C. J., McCullough, P. R., Bergeron, L. E., et al. 2010, The Astrophysical Journal, 719, 1796

Burrows, A., Budaj, J., & Hubeny, I. 2008, The Astrophysical Journal, 678, 1436

Burrows, A., Guillot, T., Hubbard, W. B., et al. 2000, The Astrophysical Journal, 534, L97

Burrows, A., Sudarsky, D., & Hubeny, I. 2006, The Astrophysical Journal, 650, 1140

Cameron, A. C., Guenther, E., Smalley, B., et al. 2010, Monthly Notices of the Royal Astronomical Society, 407, 507

Campo, C. J., Harrington, J., Hardy, R. A., et al. 2011, The Astrophysical Journal, 727

Castelli, F., & Kurucz, R. L. 2004, arXiv:astro-ph/0

Charbonneau, D., Brown, T., Noyes, R., & Gilliland, R. 2002, The Astrophysical Journal, 568, 377

Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, The Astrophysical Journal, 529, L45

Charbonneau, D., Allen, L. E., Megeath, S. T., et al. 2005, The Astrophysical Journal, 626, 523

Claret, A. 2000, Astronomy and Astrophysics, 363, 1081

Claret, A., & Bloemen, S. 2011, Astronomy and Astrophysics, 529

Cowan, N. B., & Agol, E. 2011, The Astrophysical Journal, 726

Cowan, N. B., Machalek, P., Croll, B., et al. 2012, The Astrophysical Journal, 747

Croll, B., Lafreniere, D., Albert, L., et al. 2011, The Astronomical Journal, 141, 30

Crossfield, I. J. M., Barman, T., Hansen, B. M. S., Tanaka, I., & Kodama, T. 2012, The Astrophysical Journal, 760

Crouzet, N., McCullough, P. R., Burke, C., & Long, D. 2012, The Astrophysical Journal, 761, 7

de Kok, R., Brogi, M., Snellen, I., et al. 2013, Astronomy & Astrophysics, 554, A82

de Mooij, E. J. W., Brogi, M., de Kok, R. J., et al. 2013, Astronomy & Astrophysics, 550, A54

de Mooij, E. J. W., & Snellen, I. A. G. 2009, Astronomy and Astrophysics, 493, L35

Deming, D., Seager, S., Richardson, L. J., & Harrington, J. 2005, Nature, 434, 740

Deming, D., Wilkins, A., McCullough, P., et al. 2012a, American Astronomical Society, 219

Deming, D., Fraine, J. D., Sada, P. V., et al. 2012b, The Astrophysical Journal, 754, 106

Deming, D., Wilkins, A., McCullough, P., et al. 2013, The Astrophysical Journal, 774, 95

Désert, J.-M., Lecavelier des Etangs, A., Hébrard, G., et al. 2009, The Astrophysical Journal, 699, 478

Dressel, L. 2012, Wide Field Camera 3, HST Instrument Handbook, Version 5.0

Fabrycky, D., & Tremaine, S. 2007, The Astrophysical Journal, 669

Fazio, G. G., Hora, J. L., Allen, L. E., et al. 2004, The Astrophysical Journal Supplement Series, 154, 10

Ford, E. B. 2005, The Astronomical Journal, 129, 1706

Fortney, J., Lodders, K., Marley, M., & Freedman, R. 2008, The Astrophysical Journal, 678, 1419

Fortney, J. J., Saumon, D., Marley, M. S., Lodders, K., & Freedman, R. S. 2006, The Astrophysical Journal, 642, 495

Freedman, R. S., Marley, M. S., & Lodders, K. 2008, , 174, 504

Fressin, F., Knutson, H. a., Charbonneau, D., et al. 2010, The Astrophysical Journal, 711, 374

Gibson, N. P., Aigrain, S., Roberts, S., et al. 2012a, Monthly Notices of the Royal Astronomical Society, 419, 2683

Gibson, N. P., Pont, F., & Aigrain, S. 2011, Monthly Notices of the Royal Astronomical Society, 411, 2199

Gibson, N. P., Aigrain, S., Pont, F., et al. 2012b, Monthly Notices of the Royal Astronomical Society, 422, 753

Gillon, M., Demory, B., Barman, T., et al. 2007, Astronomy and Astrophysics, 471, L51

Gillon, M., Demory, B.-O., Triaud, A. H. M. J., et al. 2009, Astronomy and Astrophysics, 506, 359

Grillmair, C. J., Charbonneau, D., Burrows, A., et al. 2007, The Astrophysical Journal, 658, L115

Grillmair, C. J., Burrows, A., Charbonneau, D., et al. 2008, Nature, 456, 767

Guillot, T., Burrows, A., Hubbard, W. B., Lunine, J. I., & Saumon, D. 1996, The Astrophysical Journal, 459

Guillot, T., & Showman, A. P. 2002, Astronomy and Astrophysics, 385

Hansen, C., Schwartz, J., & Cowan, N. 2014, arXiv preprint arXiv:1402.6699, 1

Hebb, L., Collier-Cameron, A., Loeillet, B., et al. 2009, The Astrophysical Journal, 693

Hellier, C., Anderson, D. R., Collier-Cameron, A., et al. 2011, The Astrophysical Journal, 730, L31

Herrero, E., Morales, J. C., Ribas, I., & Naves, R. 2011, Astronomy & Astrophysics, 526, L10

Howe, A. R., & Burrows, A. S. 2012, The Astrophysical Journal, 756

Hoyer, S., Lopez-Morales, M., Rojo, P., et al. 2013, Monthly Notices of the Royal Astronomical Society, 434, 46

Hubeny, I., Burrows, A., & Sudarsky, D. 2003, The Astrophysical Journal, 594, 1011

Huitson, C. M., Sing, D. K., Pont, F., et al. 2013, Monthly Notices of the Royal Astronomical Society, 434, 3252

Ibgui, L., & Burrows, A. 2009, The Astrophysical Journal, 700

Jackson, B., Greenberg, R., & Barnes, R. 2008, The Astrophysical Journal, 678

Kaeufl, H.-U., Ballester, P., Biereichel, P., et al. 2004, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 5492, Ground-based Instrumentation for Astronomy, ed. A. F. M. Moorwood & M. Iye, 1218–1227

Kalas, P., Graham, J. R., Chiang, E., et al. 2008, Science (New York, N.Y.), 322, 1345

Karkoschka, E., & Tomasko, M. 2010, Icarus, 205

Knutson, H. A., Charbonneau, D., Allen, L. E., Burrows, A., & Megeath, S. T. 2008, The Astrophysical Journal, 673, 526

Knutson, H. a., Howard, A. W., & Isaacson, H. 2010, The Astrophysical Journal, 720, 1569

Knutson, H. A., Charbonneau, D., Allen, L. E., et al. 2007, Nature, 447

Konopacky, Q. M., Barman, T. S., Macintosh, B. a., & Marois, C. 2013, Science (New York, N.Y.), 339, 1398

Kovács, G., Kovács, T., Hartman, J. D., et al. 2013, Astronomy & Astrophysics, 553, A44

Kreidberg, L., Bean, J. L., Désert, J.-M., et al. 2014, Nature, 505, 69

Kümmel, M., Walsh, J. R., Pirzkal, N., Kuntschner, H., & Pasquali, A. 2009, Publications of the Astronomical Society of the Pacific, 121

Kuntschner, H., Bushouse, H., Walsh, J. R., & Kümmel, M. 2008, ST-ECF Instrument Science Report WFC3-2008-16, 16

Lagrange, A.-M., Gratadour, D., Chauvin, G., et al. 2009, Astronomy and Astrophysics, 493, L21

Lecavelier Des Etangs, A., Pont, F., & Vidal-Madjar A., & Sing, D. 2008, Astronomy and Astrophysics, 481, L83

Lee, J.-M., Fletcher, L. N., & Irwin, P. G. J. 2012, Monthly Notices of the Royal Astronomical Society, 420, 170

Lendl, M., Gillon, M., Queloz, D., et al. 2013, Astronomy and Astrophysics, 552

Liddle, A. R. 2004, Monthly Notices of the Royal Astronomical Society, 351, L49

Line, M. R., & Yung, Y. L. 2013, The Astrophysical Journal, 779, 3

Line, M. R., Zhang, X., Vasisht, G., et al. 2012, The Astrophysical Journal, 749, 93

Machalek, P., McCullough, P., Burke, C., et al. 2008, The Astrophysical Journal, 684, 1427

Maciejewski, G., Dimitrov, D., Seeliger, M., et al. 2013, Astronomy and Astrophysics, 551

Madhusudhan, N. 2012, The Astrophysical Journal, 758, 36

Madhusudhan, N., Knutson, H., Fortney, J., & Barman, T. 2014, Protostars and Planets VI

Madhusudhan, N., Mousis, O., Johnson, T. T. V., & Lunine, J. I. J. 2011a, The Astrophysical Journal, 743, 191

Madhusudhan, N., & Seager, S. 2009, The Astrophysical Journal, 707, 24

—. 2010, The Astrophysical Journal, 725, 261

—. 2011, The Astrophysical Journal, 729, 41

Madhusudhan, N., Harrington, J., Stevenson, K. B., et al. 2011b, Nature, 469, 64

Mandel, K., & Agol, E. 2002, The Astrophysical Journal Letters, 580, L171

Mandell, A. M., Deming, L. D., Blake, G. a., et al. 2011, The Astrophysical Journal, 728, 18

Mandell, A. M., Haynes, K., Sinukoff, E., et al. 2013, The Astrophysical Journal, 779, 128

Marois, C., Macintosh, B., Barman, T., et al. 2008, Science (New York, N.Y.), 322, 1348

Mayor, M., & Queloz, D. 1995, Nature, 378, 355

McCullough, P., & Deustua, S. 2008, Instrument Science Report WFC3 2008-33, 33

McCullough, P., & MacKenty, J. 2012, STScI Instrument Science Report WFC3, 1

McLean, I. S., Becklin, E. E., Bendiksen, O., et al. 1998, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 3354, Infrared Astronomical Instrumentation, ed. A. M. Fowler, 566–578

Moses, J. I., Madhusudhan, N., Visscher, C., & Freedman, R. S. 2013, The Astrophysical Journal, 763, 25

Pont, F., Gilliland, R. L., Knutson, H., Holman, M., & Charbonneau, D. 2009, Monthly Notices of the Royal Astronomical Society: Letters, 393, L6

Pont, F., Knutson, H., Gilliland, R. L., Moutou, C., & Charbonneau, D. 2008, Monthly Notices of the Royal Astronomical Society, 385

Pont, F., Sing, D. K., Gibson, N. P., et al. 2013, Monthly Notices of the Royal Astronomical Society, 432

Ranjan, S., Charbonneau, D., Désert, J.-M., et al. 2014, The Astrophysical Journal, 785, 148

Rasio, F. A., & Ford, E. B. 1996, Science, 274

Richardson, L. J., Deming, D., Horning, K., Seager, S., & Harrington, J. 2007, Nature, 445, 892

Rieke, G. H., Young, E. T., Engelbracht, C. W., et al. 2004, The Astrophysical Journal Supplement Series, 154, 25

Rodler, F., Kurster, M., & Barnes, J. R. 2013, Monthly Notices of the Royal Astronomical Society, 432, 1980

Rodler, F., Lopez-Morales, M., & Ribas, I. 2012, The Astrophysical Journal, 753, L25

Rogers, J. C., Apai, D., López-Morales, M., Sing, D. K., & Burrows, A. 2009, The Astrophysical Journal, 707, 1707

Rogers, T. M., Lin, D. N. C., & Lau, H. H. B. 2012, The Astrophysical Journal, 1209

Rothman, L. S. 2005, Quant Spec. & Rad. Transfer, 96

Sada, P. V., Deming, D., Jennings, D. E., et al. 2012, Publications of the Astronomical Society of the Pacific, 124, 212

Schwarz, G. 1978, The Annals of Statistics, 6, 461

Seager, S., Richardson, L. J., Hansen, B. M. S., et al. 2005, The Astrophysical Journal, 632, 1122

Showman, A. P., Cooper, C. S., Fortney, J. J., & Marley, M. S. 2008, The Astrophysical Journal, 685, 1324

Sing, D. K., & López-Morales, M. 2009, Astronomy and Astrophysics, 493, L31

Sing, D. K., Lecavelier des Etangs, A., Fortney, J. J., et al. 2013, Monthly Notices of the Royal Astronomical Society, 436, 2956

Smith, a. M. S., Anderson, D. R., Skillen, I., Cameron, A. C., & Smalley, B. 2011, Monthly Notices of the Royal Astronomical Society, 416, 2096

Smith, R. M., Zavodny, M., Rahmer, G., & Bonati, M. 2008, High Energy, 7021

Snellen, I. a. G., de Kok, R. J., de Mooij, E. J. W., & Albrecht, S. 2010, Nature, 465, 1049

Southworth, J., Hinse, T. C., Dominik, M., et al. 2012, Monthly Notices of the Royal Astronomical Society, 426

Spiegel, D. S., Silverio, K., & Burrows, A. 2009, The Astrophysical Journal, 699, 1487

Stevenson, K. B., Bean, J. L., Fabrycky, D., & Kreidberg, L. 2014, , 796, 32

Stevenson, K. B., Bean, J. L., Seifahrt, A., et al. 2014, The Astronomical Journal, 147, 161

Stevenson, K. B., Harrington, J., Nymeyer, S., et al. 2010, Nature, 464, 1161

Swain, M., Deroo, P., Tinetti, G., et al. 2013, Icarus, 225, 432

Swain, M. R., Line, M. R., & Deroo, P. 2014, The Astrophysical Journal, 784, 133

Swain, M. R., Vasisht, G., & Tinetti, G. 2008, Nature, 452, 329

Swain, M. R., Vasisht, G., Tinetti, G., et al. 2009a, The Astrophysical Journal Letters, 690, L114

Swain, M. R., Tinetti, G., Vasisht, G., et al. 2009b, The Astrophysical Journal, 704, 1616

Tinetti, G., Deroo, P., Swain, M. R., et al. 2010, The Astrophysical Journal, 712, L139

Tinetti, G., Vidal-Madjar, A., Liang, M.-C., et al. 2007, Nature, 448, 169

Tregloan-Reed, J., Southworth, J., & Tappert, C. 2013, Monthly Notices of the Royal Astronomical Society, 428

Triaud, A. H. M. J., Collier-Cameron, A., Queloz, D., et al. 2010, Astronomy and Astrophysics, 524

von Essen, C., Czesla, S., Wolter, U., et al. 2013, Astronomy & Astrophysics, 561, A48

Waldmann, I. P., Tinetti, G., Deroo, P., et al. 2013, The Astrophysical Journal, 766, 7

Weidenschilling, S. J., & Marzari, F. 1996, Nature, 384

Werner, M. W., Roellig, T. L., Low, F. J., et al. 2004, The Astrophysical Journal Supplement Series, 154, 1

Wilkins, A. N., Deming, D., Madhusudhan, N., et al. 2014, The Astrophysical Journal, 783, 113

Zahnle, K., Marley, M. S., Freedman, R. S., Lodders, K., & Fortney, J. J. 2009, The Astrophysical Journal, 701, L20

Curriculum Vitae

Korey Haynes received her Bachelor of Arts in Physics with an Emphasis in Astronomy from Macalester College in 2005. She conducted a year of astronomy research at Macalester College with Professor John Cannon. She entered George Mason University in 2010, and received her Masters in Applied Physics in 2012. She graduated with a Doctorate of Philosophy in Physics in 2014. She joined the staff at Astronomy Magazine as Associate Editor in 2014.