THERMAL PHASE VARIATIONS OF WASP-12b: DEFYING PREDICTIONS

NICOLAS B. COWAN[1,10], PAVEL MACHALEK[2,3], BRYCE CROLL[4,5], LOUIS M. SHEKHTMAN[1], ADAM BURROWS[6], DRAKE DEMING[7,8],
TOM GREENE[3], AND JOSEPH L. HORA[9]

[1] Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) and Department of Physics & Astronomy,
Northwestern University, 2131 Tech Dr, Evanston, IL 60208, USA; n-cowan@northwestern.edu
[2] SETI Institute, 189 Bernardo Ave., Suite 100, Mountain View, CA 94043, USA
[3] NASA Ames Research Center, Moffett Field, CA 94035, USA
[4] Department of Astronomy & Astrophysics, University of Toronto, 50 George St., Toronto, ON, M5S 3H4, Canada
[5] Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[6] Department of Astrophysical Sciences, Princeton University, Princeton, NJ 05844, USA
[7] Planetary Systems Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA
[8] Department of Astronomy, University of Maryland, CSS Bldg., Rm. 1204, Stadium Dr., College Park, MD 20742, USA
[9] Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138, USA

ABSTRACT

We report *Warm Spitzer* full-orbit phase observations of WASP-12b at 3.6 and 4.5 μm. This extremely inflated hot
Jupiter is thought to be overflowing its Roche lobe, undergoing mass loss and accretion onto its host star, and has
been claimed to have a C/O ratio in excess of unity. We are able to measure the transit depths, eclipse depths,
thermal and ellipsoidal phase variations at both wavelengths. The large-amplitude phase variations, combined
with the planet's previously measured dayside spectral energy distribution, are indicative of non-zero Bond albedo
and very poor day–night heat redistribution. The transit depths in the mid-infrared—$(R_p/R_*)^2 = 0.0123(3)$ and
$0.0111(3)$ at 3.6 and 4.5 μm, respectively—indicate that the atmospheric opacity is greater at 3.6 than at 4.5 μm,
in disagreement with model predictions, irrespective of C/O ratio. The secondary eclipse depths are consistent
with previous studies: $F_{\rm day}/F_* = 0.0038(4)$ and $0.0039(3)$ at 3.6 and 4.5 μm, respectively. We do not detect
ellipsoidal variations at 3.6 μm, but our parameter uncertainties—estimated via prayer-bead Monte Carlo—keep
this non-detection consistent with model predictions. At 4.5 μm, on the other hand, we detect ellipsoidal variations
that are much stronger than predicted. If interpreted as a geometric effect due to the planet's elongated shape,
these variations imply a 3:2 ratio for the planet's longest:shortest axes and a relatively bright day–night terminator.
If we instead presume that the 4.5 μm ellipsoidal variations are due to uncorrected systematic noise and we fix
the amplitude of the variations to zero, the best-fit 4.5 μm transit depth becomes commensurate with the 3.6 μm
depth, within the uncertainties. The relative transit depths are then consistent with a solar composition and short
scale height at the terminator. Assuming zero ellipsoidal variations also yields a much deeper 4.5 μm eclipse depth,
consistent with a solar composition and modest temperature inversion. We suggest future observations that could
distinguish between these two scenarios.

Key words: infrared: planetary systems – planets and satellites: atmospheres – planets and satellites: composition
– planets and satellites: individual (WASP-12b) – planet–star interactions – techniques: photometric

Online-only material: color figures

1. INTRODUCTION

Thermal phase variations are a powerful way to constrain
the climate on exoplanets. Such observations have been made
for non-transiting short-period planets (Cowan et al. 2007;
Crossfield et al. 2010), but are most potent when combined with
transit and eclipse observations for edge-on systems, because
of the additional knowledge of the planet's inclination, mass,
and radius (Knutson et al. 2007, 2009a, 2009b). Secondary
eclipse depths provide a constraint on the planet's dayside
temperature. Thermal phase variations probe the day–night
temperature contrast and hence the planet's heat redistribution
efficiency. If the observational cadence and signal-to-noise ratio
are sufficiently high, phase variations are also sensitive to the
offset between the noon meridian and the planet's hottest local
stellar time, hence constraining wind speed and direction.

By considering eclipse depths at a variety of wavelengths
for a sample of 24 transiting planets, Cowan & Agol (2011b)

estimated their dayside effective temperatures, hence placing
a joint constraint on the Bond albedo and heat recirculation
efficiency of these planets. That study found that typical hot
Jupiters exhibit a variety of albedo/recirculation efficiencies,
but planets with sub-stellar equilibrium temperatures greater
than $T_0 \approx 2700$ K all seem to have lower albedo and/or
recirculation efficiency. In other words, the hottest transiting
giant planets have a qualitatively different climate than the
merely hot Jupiters, but it is not known whether this is due to a
difference in albedo, circulation, or both. Direct measurements
of hot Jupiter geometric albedos from optical secondary eclipse
observations span more than an order of magnitude and do not
resolve this degeneracy.

In this paper, we break the albedo-recirculation degeneracy
for WASP-12b (Hebb et al. 2009), one of the very hottest
known exoplanets, with a dayside temperature of ∼3000 K:
the amplitude of thermal phase variations is a direct measure
of the planet's day–night temperature contrast and hence heat
transport efficiency. If the nightside temperature is high, then the
planet's albedo must be exceedingly low to be consistent with

---

[10] CIERA Postdoctoral Fellow.

1

its high dayside temperature. If, on the other hand, the nightside temperature is low, then the planet has an albedo in the tens of percent.

WASP-12b has been a fascinating planet since its discovery. The discrepant timing of its secondary eclipse indicated that the planet had a slight eccentricity (López-Morales et al. 2010), but subsequent eclipse (Campo et al. 2011; Croll et al. 2011) and radial velocity (Husnoo et al. 2011) observations have all but ruled this out. Nevertheless, the planet's short-period orbit (1.1 days; just outside its star's Roche limit) and inflated radius ($1.8\,R_J$) led to the prediction that it is tidally distorted (Ragozzine & Wolf 2009; Leconte et al. 2011; Budaj 2011), and undergoing Roche lobe overflow followed by accretion onto its host star (Li et al. 2010; Lai et al. 2010). The putative early ingress of an ultraviolet transit observed by Fossati et al. (2010) seems to support this prediction, but may also be explained in terms of a leading bow shock from material streaming off the planet (Vidotto et al. 2010; Llama et al. 2011).

More recently, Madhusudhan et al. (2011) used the wavelength dependance of mid-infrared eclipse depths of Croll et al. (2011) and Campo et al. (2011) to constrain the atmospheric composition of WASP-12b, and found it has a carbon to oxygen ratio (C/O) greater than unity, unlike solar system planets, or the assumed composition of extrasolar planets. Those findings rested heavily on the relative eclipse depths at 3.6 and $4.5\,\mu m$. Our observations of eclipses and transits at these wave bands should be able to reinforce or rule out the high C/O scenario.

2. OBSERVATIONS AND REDUCTION

We acquired observations of WASP-12 (spectral type F9V) with IRAC (Fazio et al. 2004) on the *Spitzer Space Telescope* (Werner et al. 2004) at $3.6\,\mu m$ (2010 November 17–18) and $4.5\,\mu m$ (2010 December 11–12) as part of the Warm Mission. We used the sub-array mode and acquired images every 2 s (1.92 s effective exposure time), observing the system for approximately 1.3 days at each wave band, from slightly before a secondary eclipse to shortly after the following secondary eclipse. This yields 902 data cubes (64 frames of 32×32 pixels) at each wave band.

Due to a scheduling error, we did not observe all of the second eclipse's egress at $3.6\,\mu m$. This does not severely affect our science objectives since we simultaneously fit both eclipses at a given wavelength; even at $3.6\,\mu m$ we have nearly two full eclipse light curves to work with.

We use the basic calibrated data files and convert MJy/str to electron counts by multiplying the flux values by *GAIN*× *EXPTIME/FLUXCONV*, using parameter values from the header of the *fits* files. We use the *BMJD_OBS* and *FRAMTIME* parameter values to compute the BJD time at the middle of every exposure.

We start by considering the pixel-by-pixel time series for each 64-frame data cube, replacing *NaN*'s with the pixel's median over that data cube; if the entire time series for a given pixel is bad, it is flagged as a bad pixel and ignored in the subsequent analysis. At $4.5\,\mu m$, the first row of pixels ($y = 0$) is consistently bad.

Deming et al. (2011) noted that *Warm Spitzer* sub-array data cubes exhibit a frame-dependent background flux. At $3.6\,\mu m$, the 1st and 58th frames are consistent background outliers (both high), and there is a clear drop in background flux throughout each data cube (see Figure 5 of Deming et al. 2011). At $4.5\,\mu m$, the same two frames (1 and 58) are outliers (high and low backgrounds, respectively), and there is a slight *increase* in

background flux throughout each data cube. We elect to ignore the 1st and 58th frames of each data cube (3% of our data). To correct for the gradual change in background flux, we perform initial background subtraction on each frame of the data cube using the IDL routine MMM and excluding the 16 central pixels of the detector (those closest to the star).

We then perform a two-step sigma clipping on each pixel's time series, replacing 4σ outliers by the pixel's median in that data cube. The sigma clipping at the pixel level affects 0.028% and 0.035% of our science time series data at 3.6 and $4.5\,\mu m$, respectively.

To determine the centroid of our target, we first identify the brightest of the central 16 pixels in each frame, then fit a two-dimensional (2D) Gaussian to the 7×7 pixel box centered on this brightest pixel using the IDL routine GAUSS2DFIT.[11]

We perform aperture photometry on the individual frames of the data cubes using the IDL routine APER with a sky annulus of 7–12 pixels in radius (as did Campo et al. 2011, who observed the same system with the same instrument). Since we perform an initial background subtraction early in our reduction pipeline (see above), our results are not very sensitive to changes in the sky annulus. We verified that nudging the inner/outer radius of the annulus by 1–2 pixels does not significantly affect the goodness-of-fit or astrophysical parameters.

Bad pixels are ignored in the background estimation; images with a bad pixel within the aperture are ignored. To determine the optimal aperture for our analysis, we re-run our entire data reduction and analysis pipeline for a range of apertures, from 1.5 to 5.0 pixels, in increments of 0.5 pixels. We find that for both wave bands, the root-mean-squared scatter in the residuals is minimized for an aperture of 2.5 pixels, which we therefore adopt for the remainder of our study. This is smaller than the apertures of 3.75 and 4.0 pixels (for 3.6 and $4.5\,\mu m$ images, respectively) used by Campo et al. (2011). While using a larger aperture might reduce the photon-counting uncertainty, a small aperture makes it easier to correct for our dominant source of systematic uncertainty, the intra-pixel sensitivity variations (IPSVs) described in Section 3.2. Since the *Spitzer* heater cycling was different for the two observing campaigns (see Section 3.2), it is possible that the nature of this systematic was different for the Campo et al. (2011) observations.

Finally, we perform an iterative 4σ clipping on the flux time series, removing any outliers. This affects 0.01% and 0.02% of the science time series data at 3.6 and $4.5\,\mu m$, respectively. The $3.6\,\mu m$ cut is more generous because of the greater systematic flux variations, as described below.

The raw photometry is shown in Figure 1. (Note that we perform all of our analysis on the unbinned data, but we bin the data for plotting.) The transits are easy to spot in the middle of the observations. The eclipses and phase variations are also visible by eye at $4.5\,\mu m$. For the $3.6\,\mu m$ light curve, the first eclipse and the phase variations are difficult to distinguish from detector systematics by eye. We estimate the system flux in mJy by converting back to MJy/str and using the pixel scale parameter values, *PXSCAL1* and *PXSCAL2*. Our system fluxes—23.0(5) mJy and 14.7(1) mJy, at 3.6 and $4.5\,\mu m$, respectively—are approximately 10% lower than those of Campo et al. (2011), even when we adopt their larger apertures. Fazio et al. (2004) expected

---

[11] We follow Agol et al. (2010), who compared many centroiding algorithms and found this one to be optimal. Using flux-weighted centroiding instead of PSF-fitting results in slightly worse χ^2, commensurate correlated noise as measured using β (see first Section 4.1), and consistent astrophysical parameters.

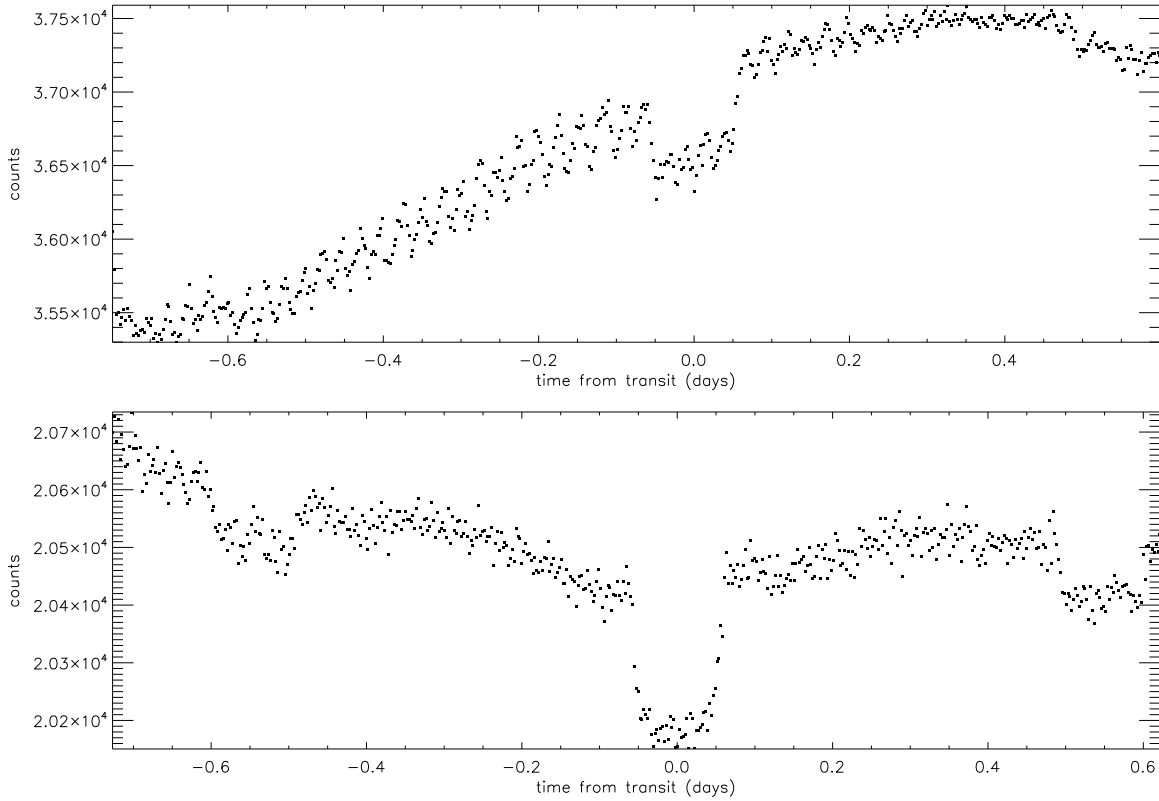**Figure 1.** Raw photometry in electron counts at 3.6 μm (top) and 4.5 μm (bottom). The data have been binned for ease of viewing.

the absolute photometric precision of IRAC to be better than 10%.

As part of our *Warm Spitzer* observations, we also obtained mapping data in sub-array mode for WASP-12, at each 3.6 and 4.5 μm, immediately following the time series described above. For the mapping data we acquired images every 0.4 s (0.36 s effective exposure time); we obtained 450 data cubes in each wave band. The purpose of these data was to map the central four pixels of the detector by scanning over them in 0.2 pixel and 0.1 pixel increments in the *x*- and *y*-directions, respectively.

The data reduction for the mapping data is identical to that for the science light curve, except that we stack the frames of each data cube using a pixel-by-pixel median. Aperture photometry is performed on these 450 stacked images rather than on the individual frames. This is necessary because of the shorter integration times for these data. (We also tried this data reduction—performing photometry cube by cube rather than frame by frame—on our phase curve data. Our best-fit model parameters were consistent using this approach, but the detector systematics were harder to correct for, leading to larger parameter uncertainties.)

The pixel-level sigma clipping affects 0.01% of both the 3.6 and 4.5 μm mapping data. There are no outliers in the flux time series for the mapping data.

3. MODEL

Our model has 9 free astrophysical parameters, plus up to 11 free parameters to characterize the detector response. The model parameters are listed in Table 1 and described below.

Table 1
Model Variables

Name	Symbol
Stellar flux	F_*
Orbital period[a]	P
Impact parameter	b
Geometric factor	a/R_*
Time of transit (BMJD)	t_0
Area ratio	$(R_p/R_*)^2$
Mean planet flux	$\langle F_p/F_* \rangle$
Thermal phase amplitude	A_{therm}
Phase offset	α_{max}
Ellipsoidal amplitude	A_{ellips}
t Linear[b]	m_t
x Linear[c]	a_1
y Linear[c]	b_1
x Quadratic[c]	a_2
y Quadratic[c]	b_2
x Cubic[c]	a_3
y Cubic[c]	b_3
x Quartic[c]	a_4
y Quartic[c]	b_4
x Quintic[c]	a_5
y Quintic[c]	b_5

Notes.
[a] We fix the orbital period to the value from Maciejewski et al. (2011).
[b] This parameter is only used when fitting occultations independently of the rest of the light curves.
[c] These parameters are only used in the polynomial IPSV fits described in Section 3.2.3.

3.1. Astrophysical Model

Occultations. Transits are modeled using the IDL implementation of Mandel & Agol (2002) with fixed nonlinear limb darkening. To determine the limb darkening coefficients, we fit a four-parameter Claret (2000) model to a Kurucz stellar model with [Fe/H] = 0.3, T_{eff} = 6250 K, and $\log g$ = 4.5 (Kurucz 1979, 2005). We set the eccentricity to zero and fix the orbital period to the value from Maciejewski et al. (2011). The impact parameter, b, and the geometrical factor, a/R_*, are allowed to vary freely. Eclipses are modeled using the uniform-disk version of the Mandel & Agol (2002) expressions, and both eclipses are modeled simultaneously with the same parameters. This means that we do not look for eclipse depth variability (searches for such variability have so far only resulted in upper limits: Agol et al. 2010; Knutson et al. 2011). We account for light-travel time within the system, but this is only a matter of 22.5 s and does not affect our analysis. By the same token, we neglect eclipse mapping effects for the planet (Williams et al. 2006; Rauscher et al. 2007; Agol et al. 2010), since we are insensitive to the resulting offset in-eclipse time of less than a minute, let alone the detailed ingress/egress morphology.

Diurnal phases. The planet's temperature and hence brightness vary as a function of local stellar time. This inhomogeneous intensity is modeled with three parameters: the orbit-averaged planet/star flux ratio, $\langle F_p/F_* \rangle$, the semi-amplitude of thermal phase variations, A_{therm}, and the offset of the phase peak from superior conjunction, $\alpha_{\max}$ ($\alpha_{\max} < 0$ corresponds to a peak prior to superior conjunction and therefore to an eastward-advected hot spot and super-rotating winds). The phase variations have a sinusoidal shape, corresponding to a sinusoidal longitudinal brightness profile for the planet (Cowan & Agol 2008).[12] Note that in the limit of poor recirculation, the longitudinal temperature profile should be more akin to a half-sine (uniformly dark on the night side), leading to thermal phase variations similar to the Lambert phase function: broader minimum, briefer maximum. We neglect reflected starlight, which does not contribute appreciably at these wavelengths.

Ellipsoidal variations. Because of the planet's small semi-major axis and inflated radius, it is likely that it—and possibly its host star—is ellipsoidal in shape rather than spherical. This leads to changes in cross-sectional area throughout the planet's orbit. To a good approximation, these variations are sinusoidal with a period half of the orbital period; the maxima occur at quadrature, when we are seeing the long axes of the star and planet, and minima at superior and inferior conjunction, when we are seeing the short axes of the two bodies.[13]

Li et al. (2010), Leconte et al. (2011), and Budaj (2011) all predict that the projected area of WASP-12b should vary by approximately 10% (peak to trough) due to its prolate shape, whether it is modeled as a prolate ellipsoid or a partially filled Roche lobe. Given the mid-infrared planet/star flux ratio of $F_p/F_* \approx 4 \times 10^{-3}$ (Campo et al. 2011), we expect to see ellipsoidal variations *in the planet* at the level of $\Delta F/F_* \approx 4 \times 10^{-4}$ (or a semi-amplitude of 2×10^{-4}).

The presence of a massive companion should also produce ellipsoidal variations *in the star*, as seen at optical wavelengths in the system HAT-P-7 (Welsh et al. 2010). Using the expressions given in Faigler & Mazeh (2011), we estimate the semi-amplitude of these variations to be $\sim 4 \times 10^{-5}$ at all wavelengths. We therefore expect that at thermal wavelengths, the ellipsoidal variations of the system should be dominated by the shape of the planet and not that of its host star.

We experimented with different functional forms for the planet's phase variations in an effort to reduce correlations between astrophysical variables. Our best model in this regard is

$$\frac{F_p}{F_*} = \left\langle \frac{F_p}{F_*} \right\rangle + A_{\mathrm{therm}} \cos(\alpha - \alpha_{\max}) - A_{\mathrm{ellips}} \cos(2\alpha), \quad (1)$$

where A_{therm} and A_{ellips} are the semi-amplitudes of diurnal and ellipsoidal phase variations, respectively, and α is the phase angle ($\alpha = 0$ at superior conjunction, $\alpha = \pi$ at inferior conjunction).

The secondary eclipse depth is related to the model variables by

$$\frac{F_{\mathrm{day}}}{F_*} = \left\langle \frac{F_p}{F_*} \right\rangle + A_{\mathrm{therm}} \cos \alpha_{\max} - A_{\mathrm{ellips}}, \quad (2)$$

and to first order the associated uncertainties can be propagated as

$$\sigma^2_{F_{\mathrm{day}}/F_*} = \sigma^2_{\langle F_p/F_* \rangle} + \cos^2 \alpha_{\max} \sigma^2_{A_{\mathrm{therm}}} \\ + A^2_{\mathrm{therm}} \sin^2 \alpha_{\max} \sigma^2_{\alpha_{\max}} + \sigma^2_{A_{\mathrm{ellips}}}. \quad (3)$$

In practice, this is an overestimate of uncertainty, because $\langle F_p/F_* \rangle$ and A_{therm} are anti-correlated.

3.2. Detector Response Model

IRAC channels 1 and 2 exhibit well-known IPSVs: photons hitting certain parts of a pixel lead to more electron counts than photons hitting other parts of the same pixel (e.g., Charbonneau et al. 2005; Morales-Calderón et al. 2006).[14] In general, the sensitivity to photons is lowest near the edges of a pixel and greatest near its center (see bottom panels of Figure 2). On its own, IPSV would not be a problem for precision time-resolved relative photometry. But over the course of observations, the point-spread function (PSF) of the target star moves on the detector. Even though the PSF spans many pixels, the IPSVs do not average out, because most of the flux falls in the PSF core.

Since they are ultimately caused by changes in the PSF position, attempts to correct for IPSVs rely on accurate centroiding, described in Section 2. The centroiding is shown in Figure 2 for 3.6 μm (left) and 4.5 μm (right). The top panels show the centroid jitter and drift over the course of the observations; the second panels show the 2D path of the centroid; the bottom panels show the intra-pixel sensitivity map of the four central pixels of the detector, constructed by applying the Ballard et al. (2010) point-by-point decorrelation to our mapping data (see Section 3.2.2).

Both x and y centroids exhibit fast jitter (period of roughly half an hour) with peak–trough amplitude of 0.05–0.1 pixels. This is the same jitter that used to have a period of roughly

---

[12] In general, the offset between thermal phase maximum and superior conjunction is *not* the same as the offset of the hottest longitude of the planet with respect to the sub-stellar meridian. For realistic longitudinal temperature profiles, the observed offset in the light curve is significantly greater than the hot spot offset (Cowan & Agol 2011a). In the case of a first-order sinusoidal phase curve, however, the two offsets are one and the same.

[13] This is a common approximation for ellipsoidal variations (e.g., Faigler & Mazeh 2011), but we discuss the exact expression in Section 5.3.1.

[14] This is entirely different than *inter*-pixel sensitivity variations, which should have been largely calibrated out by flat fielding. In any case, a scheme that corrects IPSVs implicitly corrects inter-pixel sensitivity variations as well.
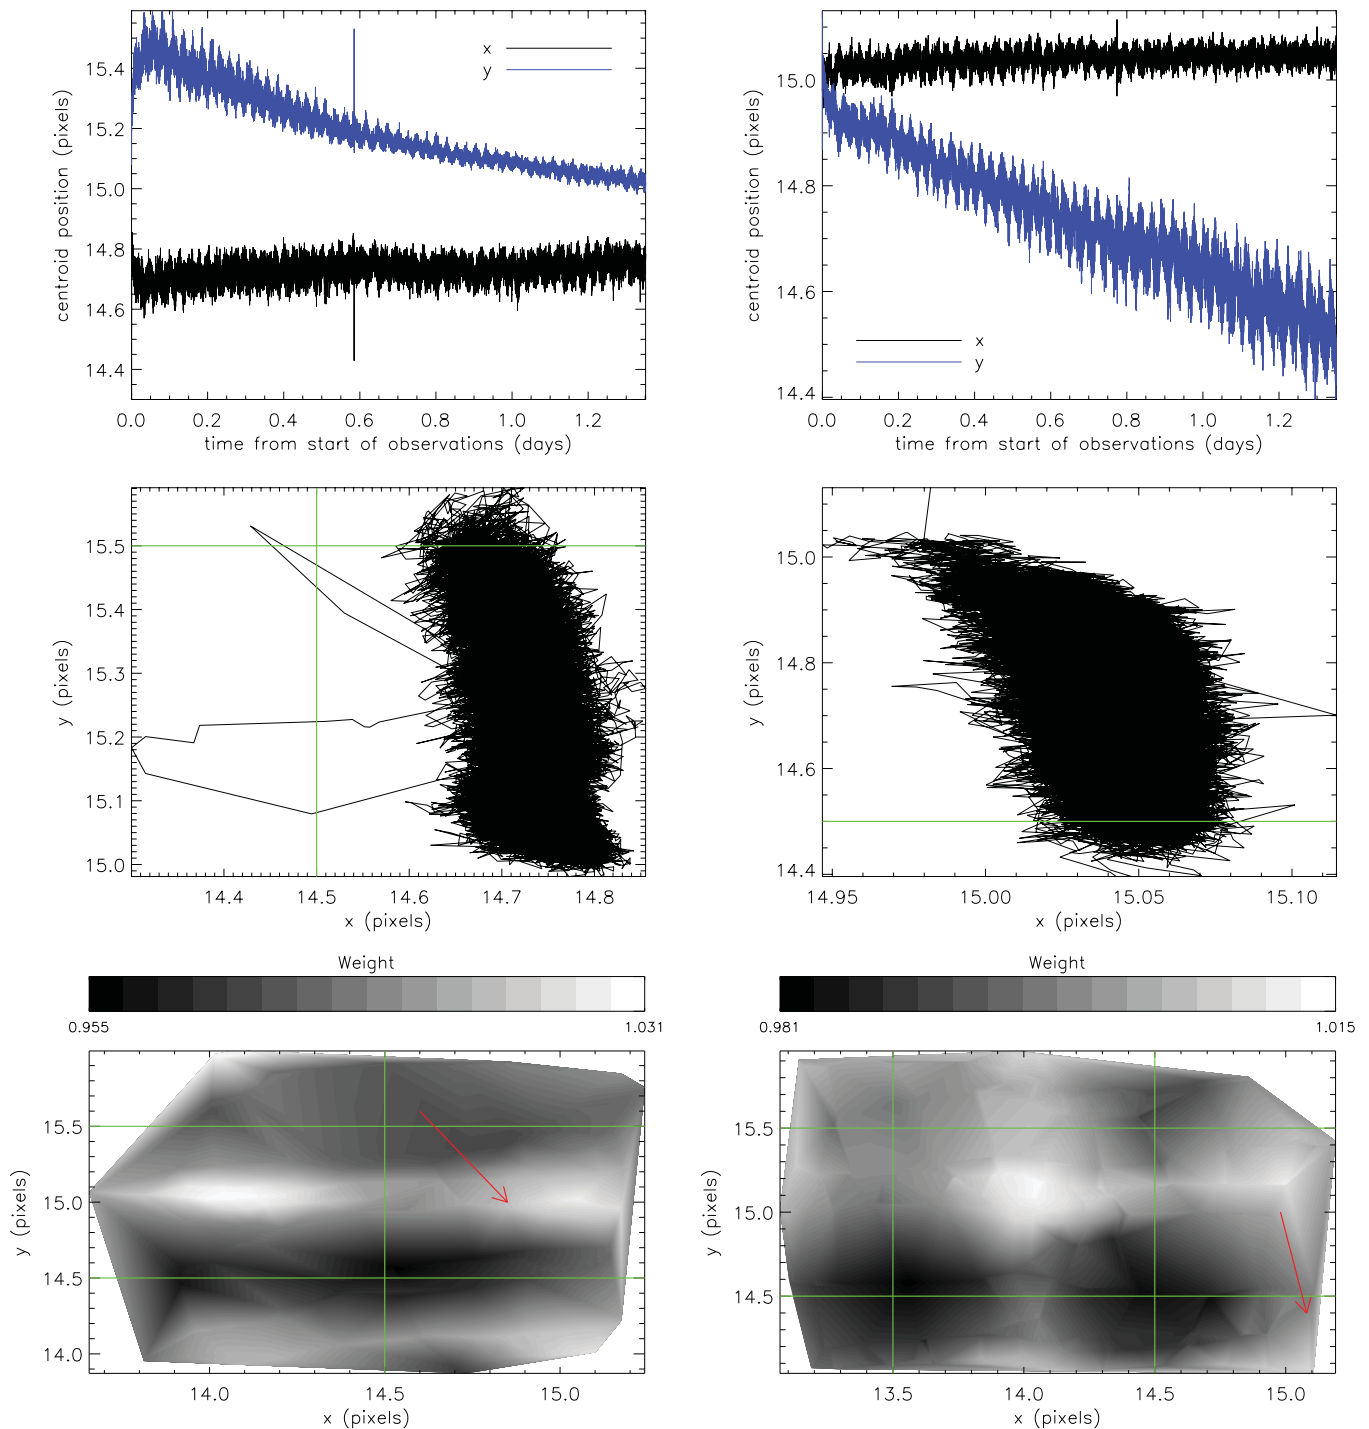
Figure 2. Point-spread function centroid movement at 3.6 μm (left) and 4.5 μm (right). The top panels show the jitter and drift of the PSF centroid. The second panels show the two-dimensional wander of the centroid. The bottom panels shows the intra-pixel sensitivity map for the central regions of the detector, constructed by applying the Ballard et al. (2010) point-by-point decorrelation to our mapping data (see Section 3.2.2), with red arrows marking the approximate drift of the PSF over the course of our observations. The green lines show pixel edges.

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

an hour: it is related to the heater cycling on *Spitzer*. The cycling frequency was doubled in fall 2010, which doubled the centroid jitter frequency and roughly halved the amplitude of the jitter.[15] The smaller amplitude of the jitter is undoubtedly an improvement, while the higher frequency may or may not be a nuisance, depending on the duration of ingress/egress for

a given planet. In our data, the 3.6 μm flux exhibits clear 1% peak-to-trough flux variations on the centroid jitter timescale, while the 4.5 μm flux does not.

There is also a longer-term centroid drift, which is greatest in the *y*-direction: 0.5 pixels of motion over the ~1 day observation at both 3.6 and 4.5 μm. The *x*-direction shows almost no long-term drift (0.05 pixels, comparable to the faster jitter). Looking at the bottom panels of Figure 2, it is easy to understand why

the IPSVs are worse at 3.6 μm than at 4.5 μm: at the shorter wave band, the PSF drifted up a steep slope in sensitivity from a pixel corner toward a center; at the longer wave band, the PSF contoured below a ridge in sensitivity.

The telescope takes a few hours to settle after pointing at a new target, resulting in larger PSF excursions for the first few data cubes of each light curve. It is difficult to correct for IPSVs in poorly sampled regions of the detector, so we remove the first 0.05 days of both the 3.6 and 4.5 μm light curves (3.8% of our data in each wave band).

The crux of the data reduction process is correcting for IPSVs, because our astrophysical signals (eclipses and phase variations) have an expected amplitude of 0.4%, comparable to—or smaller than—these systematics. We tried a variety of techniques, of which we describe the most promising below. We first present two methods for removing the IPSVs *before* fitting our astrophysical model of the system. These techniques are attractive because they allow one to produce a systematics-corrected light curve independent of any astrophysical model assumptions. We then present an attempt to simultaneously fit the IPSV and the astrophysical brightness variations.

3.2.1. Gaussian Decorrelation

We follow Ballard et al. (2010) in using point-by-point positions and fluxes to generate an intra-pixel sensitivity map with x and y Gaussian smoothing lengths of σ_x and σ_y, respectively. Stevenson et al. (2011) recently introduced a similar—but faster—method using bilinear interpolation. Since we are not attempting to iteratively fit the astrophysical and IPSV model, we use the simpler Ballard et al. (2010) method. We experimented with different smoothing lengths and chose the combinations that yield the smallest χ^2 value for the final model fit; the astrophysical parameters are not very sensitive to changes in the smoothing length. At 3.6 μm we use $\sigma_x = 0.017$ and $\sigma_y = 0.0043$, as in Ballard et al. (2010); at 4.5 μm we use $\sigma_x = \sigma_y = 0.015$. The resulting pixel maps are shown in the top panels of Figure 3.

We then divide the raw photometry by the weight function and fit our astrophysical model. The second panels of Figure 3 show the corrected data with best-fit astrophysical model and residuals. The bottom panels show the scatter in the residuals as a function of binning, along with a red line indicating the Gaussian-noise limit of root-mean-squared scatter scaling as $\sqrt{N}$. The normalization of this theoretical curve is based on the Poisson error for our electron counts (see first the raw photometry in Figure 1). The best-fit astrophysical parameters are listed in Table 2 (3.6 μm) and Table 3 (4.5 μm).

3.2.2. Mapping Data

The purpose of the mapping data was to deliberately map the central four pixels of the detector by scanning over them in 0.2 pixel and 0.1 pixel increments in the x- and y-directions, respectively. Since these observations are much shorter than the planet's orbital time, and were scheduled to avoid transits or eclipses, the changes in flux are in principle entirely due to the centroid position on the detector.

In practice, the 3.6 μm light curve observations ended approximately 3.8 minutes before the end of eclipse egress, and the 3.6 μm mapping observations immediately followed. We therefore remove the first 4 minutes (0.003 days) of the 3.6 μm mapping data.

We use the mapping data centroids and fluxes to generate a weight map at the locations of the science centroids, again

following Ballard et al. (2010). We adopt larger smoothing kernels set by the Nyquist sampling frequency of the regularly spaced mapping centroids ($\sigma_x = 0.1$, $\sigma_y = 0.05$). We then use this weight function to correct the science light curve as above.

Using the mapping data to generate pixel maps has the advantage that we are not self-calibrating our science data, and hence are not liable to throw the baby out with the bath water. On the other hand, it does not successfully remove the systematics: the model fits are far worse than either the Gaussian decorrelation discussed above or the polynomial models discussed below (the discrepancy in χ^2 is a factor of $\sim$10 at 3.6 μm and $\sim$4 at 4.5 μm). This means that the IPSV must have fine spatial structure (as seen by Ballard et al. 2010), and/or some additional flux or time dependence.

3.2.3. Polynomial IPSV Model

Here, we model the IPSVs as a polynomial in the centroid x and y. We simultaneously fit our astrophysical model and the IPSVs by treating the x and y centroids as independent variables in our function. We model the IPSVs as

$$\frac{F_{\text{obs}}}{F_{\text{astro}}} = \frac{1 + \sum_{i=1}^{n}\left[a_i(x - \bar{x})^i + b_i(y - \bar{y})^i\right]}{\langle 1 + \sum_{i=1}^{n}\left[a_i(x - \bar{x})^i + b_i(y - \bar{y})^i\right]\rangle}, \quad (4)$$

where F_{obs} is the observed flux, F_{astro} is the astrophysical model, and $\bar{x}$ and $\bar{y}$ are the mean centroid positions. Formally, cross-terms are necessary to describe an arbitrary 2D function, but we find that including cross-terms does not significantly improve the χ^2 or affect the astrophysical parameters. This is a testament to the fact that the bulk of the PSF motion is in the y-direction.

We experiment with polynomials up to sixth order. To test whether each additional pair of coefficients (one each for x and y) improved the fit, we use the Bayesian Information Criterion (BIC; Schwarz 1978), which imposes a penalty term on the χ^2 for additional free parameters: BIC $= \chi^2 + k \ln N$, where k is the number of free parameters and $N \approx 52{,}000$ is the number of data. Since $\ln(52{,}000) \approx 11$, an additional parameter is acceptable if it improves the χ^2 by at least 11. The BIC for our 3.6 μm data improved with the addition of parameters up to and including fifth order. The BIC for our 4.5 μm data was not improved by the addition of terms beyond cubic.

Unlike some previous studies, we do not include a linear ramp *in time*. Since the bulk of the PSF motion is a monotonic drift in the y-direction, we found the ramp in time to be highly correlated with the linear and quadratic coefficients of the y sensitivity, and the fit was not significantly improved.[16]

We show the resulting fits in Figures 4 and 5. For each figure, the top panel shows: the systematics-corrected light curve and best-fit astrophysical model (top inset), the residuals after subtraction of the best-fit thermal phases, along with the best-fit ellipsoidal variations model (middle inset), and the residuals after removing ellipsoidal variations (bottom inset). (Ellipsoidal variations of the planet do not affect in-eclipse data since the planet is hidden from view; hence we remove the in-eclipse data from this panel.) The bottom left panel shows the weight function used to correct the data. The bottom right panel shows the scatter in the residuals as a function of binning; the red line shows the photon noise limit; the vertical dotted line denotes the timescale of ingress/egress.

---

[16] We do include a linear ramp in time when performing isolated fits to transits or eclipses, since (1) those shorter time series do not provide enough leverage to properly fit the IPSV, and (2) the linear ramp can act as a proxy for the thermal phase variations, which are not explicitly fit for in these cases.
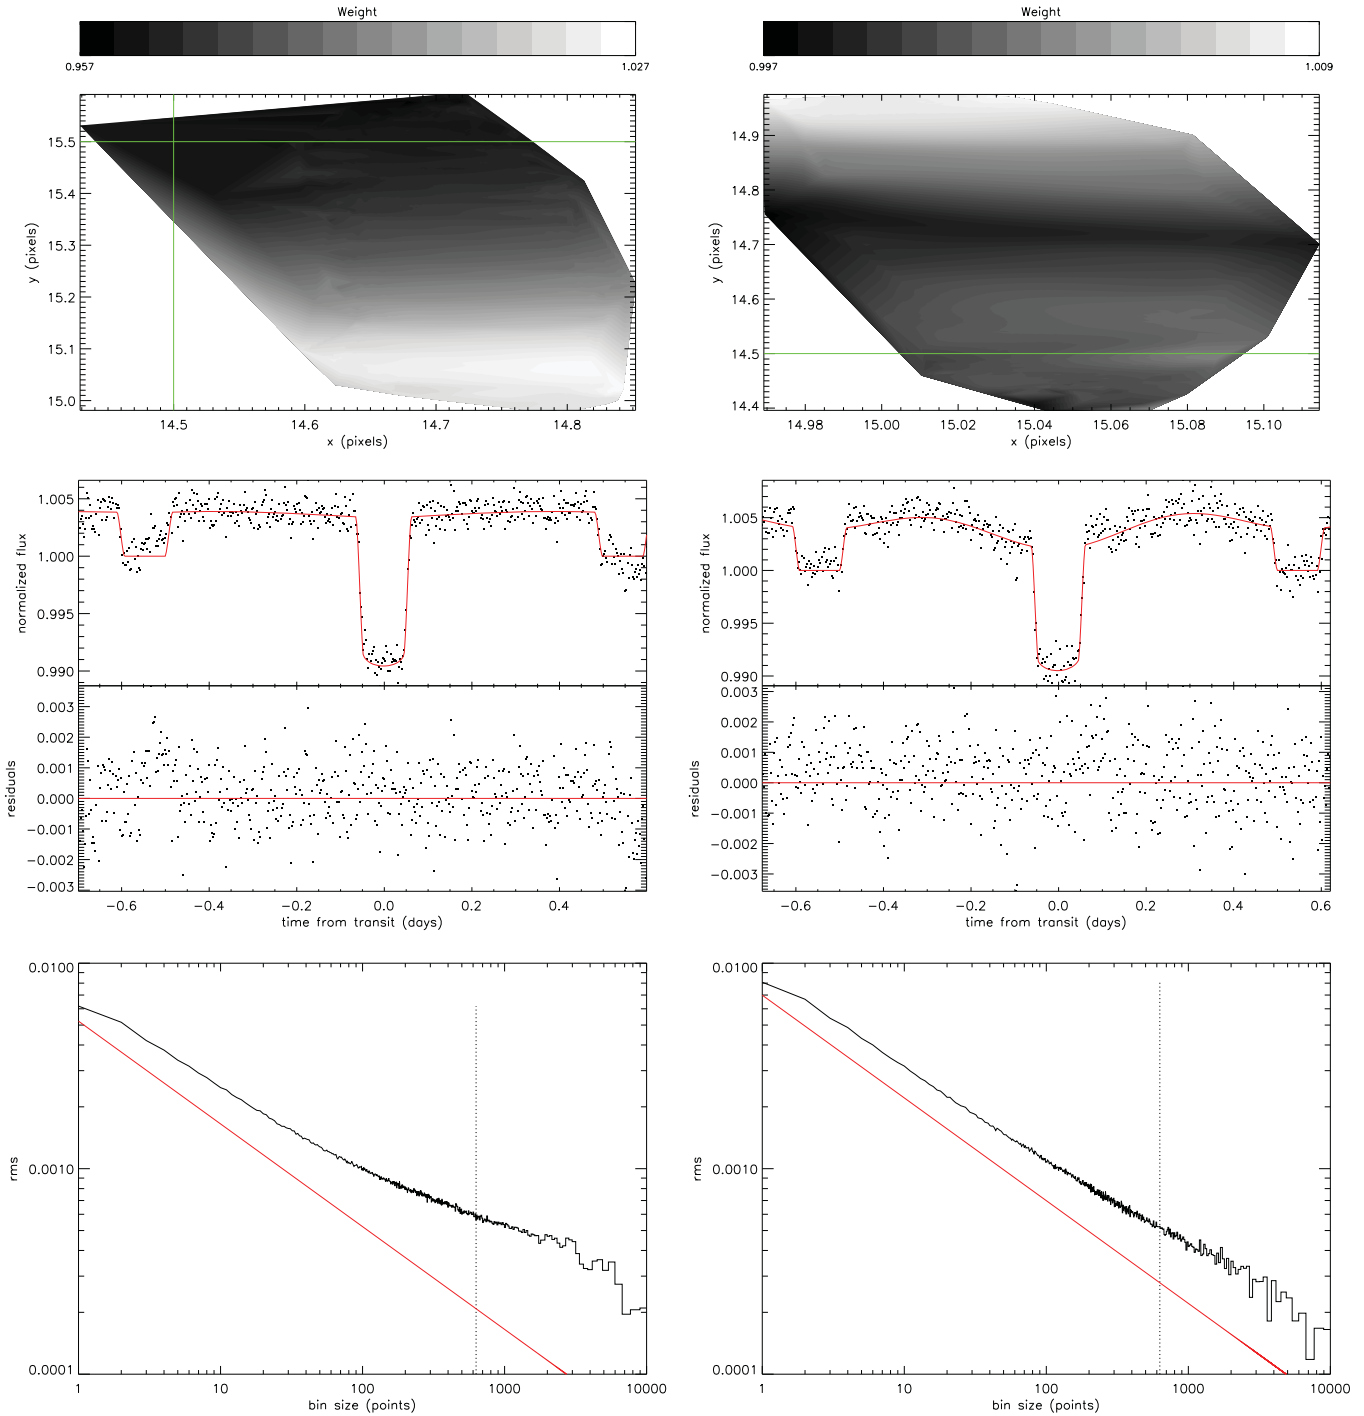
Figure 3. WASP-12b at 3.6 μm (left) and 4.5 μm (right), where we have corrected for IPSVs using the Ballard et al. (2010) point-by-point decorrelation, as described in Section 3.2.1. The upper panels show the IPSV map determined from our science data, with pixel edges shown in green. The second panels show the corrected data with best-fit astrophysical model and residuals. The bottom panels show the scatter in the residuals as a function of binning; the red line shows the photon noise limit; the vertical dotted line denotes the timescale of ingress/egress.

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

4. MODEL FITTING AND ERROR ANALYSIS

We use the IDL implementation of a Levenberg–Marquardt (L–M) χ^2 gradient descent routine, `MPFITFUN`, to find the best-fit model parameters. The covariance matrix of the model parameters provides a first guess at the parameter uncertainties.

L–M or Markov Chain Monte Carlo (MCMC) error estimation depends on the photometric uncertainties: larger error bars on the data lead to larger uncertainties on the model parameters. For the initial fits, we optimistically set the error bars on our data at the Poisson limit, $1/\sqrt{N_{\mathrm{counts}}}$; this means that the reduced χ^2 of our best-fit model fits is somewhat larger than 1 (see Tables 2 and 3), and it means that either L–M or MCMC will underestimate parameter uncertainties. To alleviate this problem, we then scale the data uncertainties to give a reduced χ^2, χ_R^2, of unity (by multiplying uncertainties by the square root of the best χ_R^2). In

Table 2
3.6 μm Parameters

Calibration Method	χ_R^2	b	a/R_*	$(R_p/R_*)^2$	F_{day}/F_*	$2A_{therm}$	α_{max}	A_{ellips}
Point-by-point decorrelation	1.392[a]	0.5(1)	2.8(2)	0.0125(4)	0.0038(4)	0.0004(3)	0(29)°	1(1) × 10^{-4}
Polynomial in x and y[b]	1.384[a]	0.3(2)	3.1(2)	0.0123(3)	0.0033(4)	0.0038(6)	−53(7)°	1(1) × 10^{-4}

Notes.
[a] When comparing these values, it is worth remembering that with approximately 52,000 degrees of freedom, a model is significantly better if it improves χ_R^2 by at least 0.004.
[b] Our fiducial analysis.

Table 3
4.5 μm Parameters

Calibration Method	χ_R^2	b	a/R_*	$(R_p/R_*)^2$	F_{day}/F_*	$2A_{therm}$	α_{max}	A_{ellips}
Point-by-point decorrelation	1.326[a]	0.5(1)	2.9(2)	0.0112(4)	0.0039(3)	0.0019(3)	−12(6)°	1.1(1) × 10^{-3}
Polynomial in x and y[b]	1.324[a]	0.5(1)	2.9(2)	0.0111(4)	0.0039(3)	0.0040(3)	−16(4)°	1.2(2) × 10^{-3}

Notes.
[a] When comparing these values, it is worth remembering that with approximately 52,000 degrees of freedom, a model is significantly better if it reduces χ_R^2 by at least 0.004.
[b] Our fiducial analysis.

the present case, this entails inflating the error bars by a constant factor of 10%–20%, depending on the wave band and model in question. Scaling photometric errors to obtain a reduced χ^2 of unity renders the χ^2 useless for comparing different models; we therefore quote the best reduced χ^2 prior to inflating the error bars. After adjusting the photometric uncertainties, we normalize the light curve to the in-eclipse (star-only) value, and fix F_* to unity in order to avoid correlations between F_* and $\langle F_p/F_* \rangle$ in the final fits.

It is well known that simply using the covariance matrix from the L–M fit does not provide a robust error estimate, so we estimate parameter uncertainties using a variety of other techniques. We experimented with MCMC and boot-strap MC error estimates and found them to be slightly larger than—but comparable to—those from the L–M. We found that our most conservative error estimates (typically larger by a factor of two or more) are obtained by considering the residuals of our best-fit model: either binning of residuals or resampling of residuals using a prayer-bead MC. Throughout the manuscript we always adopt the largest uncertainty for a given parameter, but it is worth noting that even our most conservative error estimates may still be 15%–30% smaller than what one would obtain with a wavelet analysis (Carter & Winn 2009).

4.1. Binning of Residuals

Pont et al. (2006) proposed a simple method to account for red noise by considering how the scatter in residuals decreases with bin size (bottom panels of Figures 3–5). At the left end of the plots, where we are considering point-to-point scatter in the residuals, the scatter is only 10%–20% greater than the Poisson counting limit shown in red ($\propto \sqrt{M/N(M-1)} \approx \sqrt{N}$, where N is the number of observations per bin and M is the number of bins). But the observed scatter does not follow the theoretical relation as the data are binned. The most important timescale for transit and eclipse parameter estimation is the duration of ingress/egress, which we denote by a vertical dotted line in those panels (21 minutes for WASP-12b). The scatter on this timescale determines the accuracy we can expect to achieve for transit or eclipse depths.

Following Winn et al. (2007, 2008), we define the factor β as the actual scatter (black line) divided by the theoretical Poisson limit (red line) on the 21 minute timescale (vertical dotted black line); our residuals have $\beta = 2$–3. To account for red noise in parameter uncertainties, we simply inflate the L–M parameter uncertainties by β. (Inflating the photometric error bars by β and recomputing the covariance matrix using L–M takes longer and produces slightly smaller parameter uncertainties.) The binning of residuals method turns out to be the most conservative error estimate for transit- and eclipse-specific model parameters (b, a/R_*, $(R_p/R_*)^2$, etc.).

4.2. Boot-strap and Prayer-bead Monte Carlo

We also estimate parameter uncertainties using two resampling techniques: boot-strap and prayer-bead MCs. Both of these techniques use the scatter in the residuals of our best-fit model as an estimate of photometric uncertainty. In both cases the residuals are shifted in time and added back to the best-fit model to produce a new instance of the light curve, which is then fit using the L–M.[17] The standard deviation in the sequence of model parameters is our estimate of their 1σ uncertainty. Note that—by construction—resampling techniques cannot improve parameter estimates: the best-fit parameters are those determined by fitting the original time series.

For the boot-strap MC, the residuals are randomly shuffled so—much like the L–M and MCMC techniques—the boot-strap error analysis is insensitive to the ordering of residuals. Such error estimates will therefore only be accurate insofar as residuals are uncorrelated in time. Indeed, we found that error estimates from the boot strap were comparable to those from the L–M or MCMC analyses.

The prayer-bead analysis maintains the relative ordering of the residuals and simply shifts them all by the same amount (wrapping around the start/end of the data), so that correlated noise present in the residuals is preserved. A prayer-bead

---

[17] Since we use L–M to find best-fit solutions, one might think that the prayer-bead and boot-strap analyses implicitly depend on photometric error estimates. But in fact, the L–M algorithm settles on the same solution irrespective of error bars (within reason). The prayer-bead and boot-strap parameter uncertainties are therefore effectively independent of the photometric uncertainties.
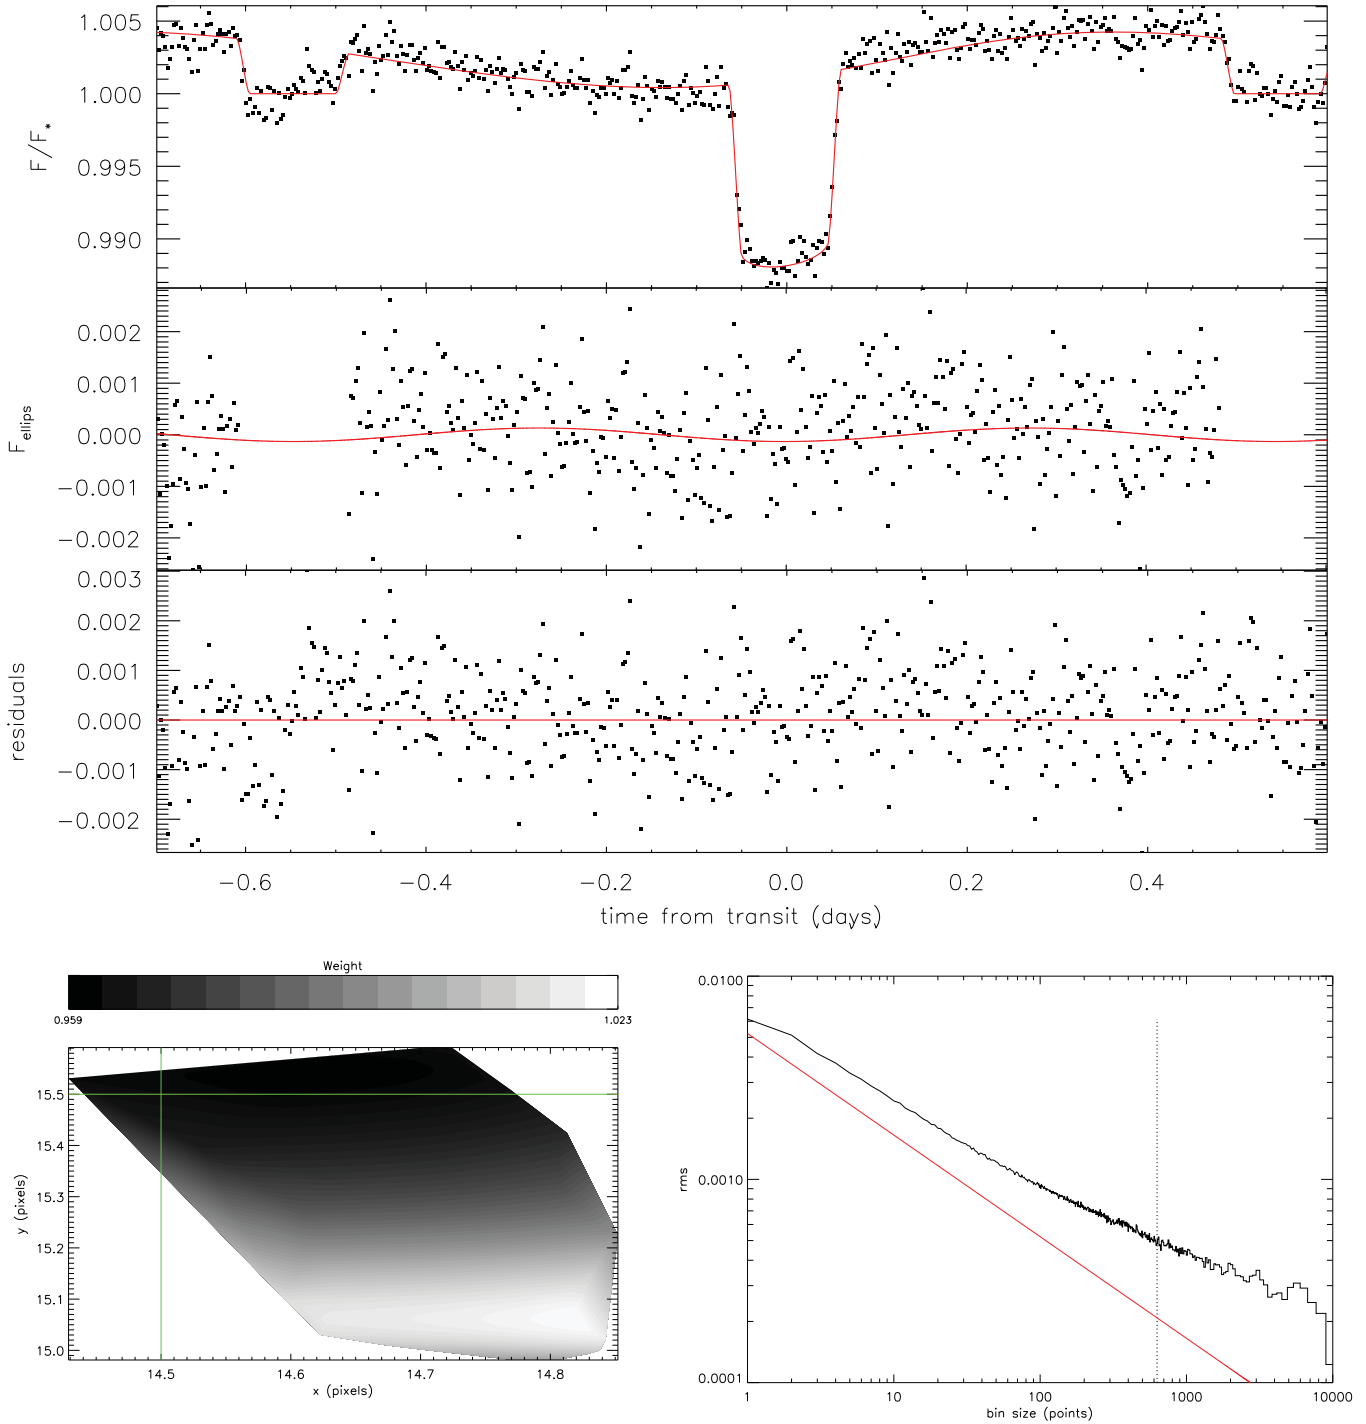
Figure 4. WASP-12b at 3.6 μm, where we have treated the IPSVs as a polynomial function in both *x* and *y* centroid, as described in Section 3.2.3. The top panel shows the systematics-corrected light curve and best-fit astrophysical model (top inset), the residuals after subtracting the best-fit transit, eclipse and thermal phase model, along with the best-fit ellipsoidal variations model (middle inset), and the residuals after removing ellipsoidal variations (bottom inset). Ellipsoidal variations of the planet do not affect in-eclipse data since the planet is hidden from view; hence we remove the in-eclipse data from this panel. The bottom left panel shows the weight function used to correct the data, with pixel edges shown in green. The bottom right panel shows the scatter in the residuals as a function of binning; the red line shows the photon noise limit; the vertical dotted line denotes the timescale of ingress/egress.

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

analysis is not appropriate if the nature of the noise is expected to change throughout the observations (e.g., the 8 μm ramp seen in cryogenic *Spitzer* observations). But there is no evidence for such changes in behavior in *Warm Spitzer* data in general, or in our time series in particular. The prayer bead can only have as many iterations as there are data points, but this is not a problem for the current study given our approximately

52,000 images; we run the prayer bead for 10,000 iterations with randomly chosen offsets and verify that the uncertainties have converged by comparing to 100- and 1000-iteration prayer-bead MCs. It is also worth noting that with such a long data set, we are more likely to observe rare instances of bad behavior in the detector, making the prayer-bead technique particularly conservative. Indeed, prayer-bead MC provides the
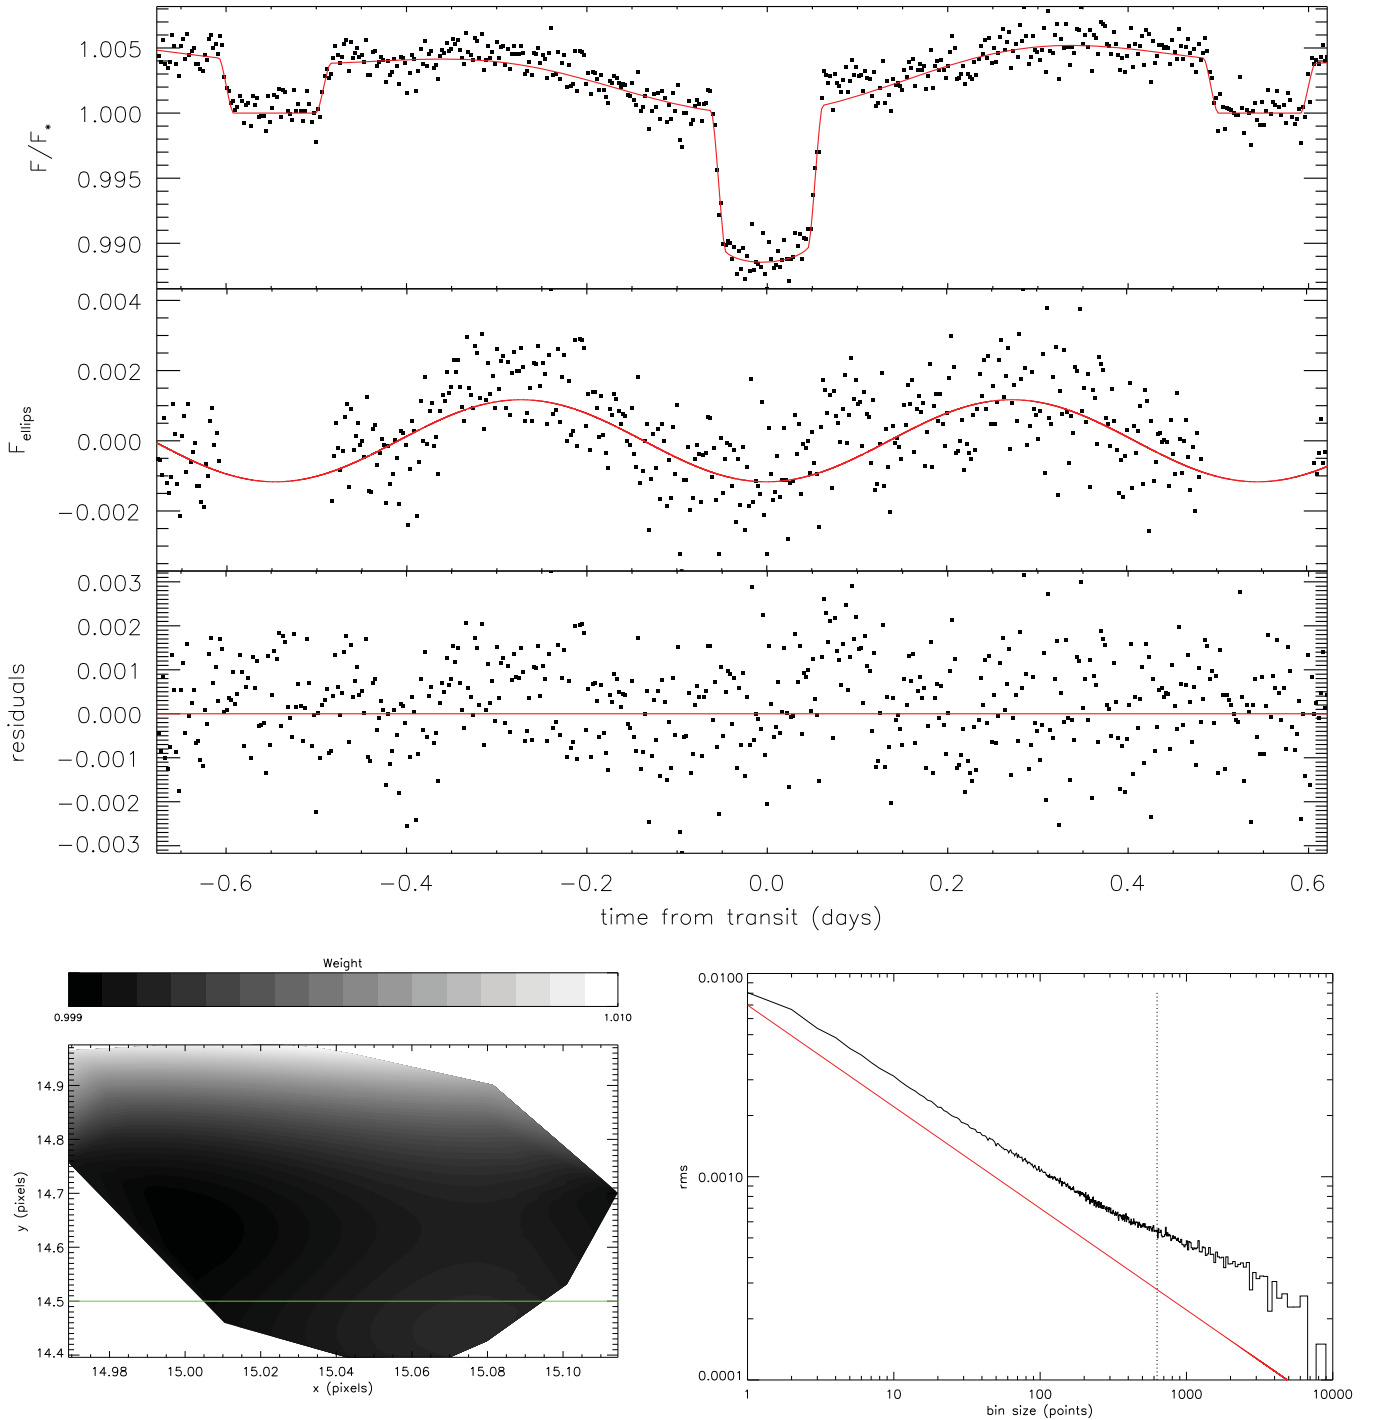
Figure 5. WASP-12b at 4.5 μm, where we have treated the IPSVs as a polynomial function in both x and y centroid, as described in Section 3.2.3. The top panel shows the systematics-corrected light curve and best-fit astrophysical model (top inset), the residuals after subtracting the best-fit transit, eclipse and thermal phase model, along with the best-fit ellipsoidal variations model (middle inset), and the residuals after removing ellipsoidal variations (bottom inset). Ellipsoidal variations of the planet do not affect in-eclipse data since the planet is hidden from view; hence we remove the in-eclipse data from this panel. The bottom left panel shows the weight function used to correct the data, with pixel edges shown in green. The bottom right panel shows the scatter in the residuals as a function of binning; the red line shows the photon noise limit; the vertical dotted line denotes the timescale of ingress/egress.

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

largest error bars for the phase variation parameters: A_{therm}, α_{max}, A_{ellips}.

5. RESULTS

The two methods used to remove IPSVs are fundamentally different. The Gaussian decorrelation uses local information to correct the flux with no assumption about larger-scale trends; it is able to correct for small-scale variations in sensitivity, but requires very high densities of centroids. The polynomial fit assumes a functional form for the smoothly varying sensitivity, but is better able to correct regions that are less-well sampled. It is not clear which method is better suited to a given data set. Ballard et al. (2010) found the decorrelation to be better (in

terms of χ^2) than the polynomial fit when analyzing their binned 4.5 μm time series. We run both methods on unbinned data and find that the two methods perform equally well at 4.5 μm, while at 3.6 μm the polynomial fit is better.

Given their very different underlying philosophies, it is encouraging that the two methods yield similar weight functions (compare the pixel maps in Figure 3 to those in Figures 4 and 5). In fact, the transit depths, eclipse depths, and ellipsoidal variations recovered by the two techniques are generally consistent. The thermal phase variation amplitude and offset differ significantly, however.

5.1. Transits

Our values for the impact parameter and geometrical factor are broadly consistent with published values. Note that we simultaneously fit the transits and eclipses, and eclipses are notoriously bad at constraining these orbital parameters.[18]

Combining our transit times for epochs 925 and 947 (BMJD of 55,518.0407(4) and 55,542.0521(4), respectively) with those of Chan et al. (2011) and the ephemeris of Maciejewski et al. (2011), we obtain a BJD discovery epoch transit center of $t_0 = 2454508.9768(2)$ and period $P = 1.0914207(4)$ days. Since our transits occurred slightly earlier than predicted, our best-fit orbital period is 3.5σ shorter than that of Maciejewski et al. (2011), but it is notoriously difficult to compare transit times at different wavelengths analyzed by different groups because of subtle differences in star-spot coverage, treatment of limb darkening, etc. (Désert et al. 2011b).

More importantly, we are simultaneously fitting an entire orbit's worth of data including many different astrophysical effects, while the optical transit data were fit on their own, irrespective of longer-term astrophysical trends. In order to make a more appropriate comparison, we try fitting the transits independently of the rest of the data (considering only those data within 0.15 days of the transit center). This yields later transit centers, leading to an ephemeris more consistent with that of Maciejewski et al. (2011): $t_0 = 2454508.9767(2)$, $P = 1.0914210(4)$.

Because of WASP-12b's high temperature and inflated radius, the variations in transit depth with wavelength should be 3× larger than for "typical" hot Jupiters (using the scaling relation from Winn 2010). Based on observations and models of HD 189733b (Fortney et al. 2010), one might therefore expect the transit depth at 4.5 μm to be $\sim 10^{-3}$ deeper than at 3.6 μm. This is borne out by the dotted black line in Figure 6, which shows a Burrows et al. (2007, 2008) model transit spectrum assuming solar composition and a day-like temperature–pressure (T-P) profile. If the planet's terminator has a night-like T-P profile (shorter scale height; solid black line in Figure 6), the difference in transit depth could be as small as 2×10^{-4}—but always with a transit depth greater at 4.5 μm than at 3.6 μm.

In Figure 6, we compare our transit depths at 3.6 and 4.5 μm with previously published optical values: 0.0138(2) (B and z' band, Hebb et al. 2009), 0.01380(16) (R band, Maciejewski et al. 2011), and 0.0125(4) (V band, Chan et al. 2011), yielding a three-band transmission spectrum of the planet. Curiously, the transit depth at 4.5 μm is considerably shallower than at 3.6 μm. If we adopt the larger Maciejewski et al. (2011) optical transit depth, then our mid-IR transit depths could be indicative

---

[18] If we only consider the 0.3 days of data centered on each transit, we find impact parameters of $b = 0.46$ and 0.60 (larger than published values), and geometrical factors of $a/R_* = 2.9$ and 2.7 (smaller than published values) at 3.6 and 4.5 μm, respectively.
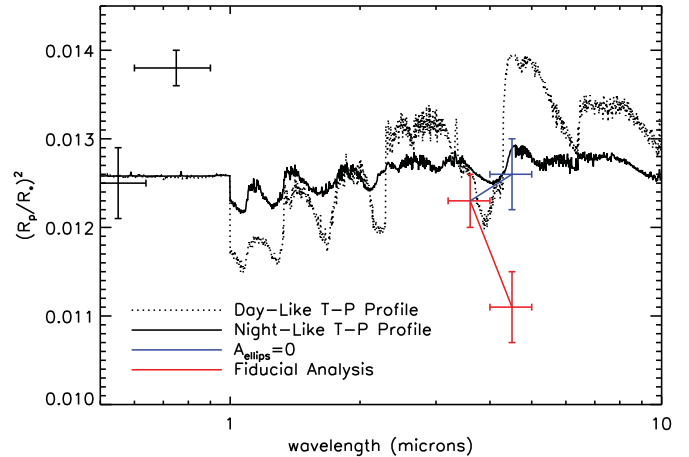

Figure 6. Predicted wavelength-dependent transit depth of WASP-12b based on a solar composition (Burrows et al. 2007, 2008). The dotted black line shows a model with a day-like temperature–pressure (T-P) profile (large scale height); the solid line shows a model with a night-like T-P profile (short scale height). The red points correspond to a model with ellipsoidal variations; the blue point corresponds to a model without 4.5 μm ellipsoidal variations (see first Section 5.3.2). The two black points with error bars on the left show the optical transit depth from Hebb et al. (2009) and Maciejewski et al. (2011) (top) and Chan et al. (2011) (bottom); we normalize the model transit spectrum to the latter's observation.

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

of hazes or optical absorbers in the atmosphere of WASP-12b, as has been inferred for HD 189733b (Pont et al. 2008; Sing et al. 2011). The larger radius at 3.6 μm as compared to 4.5 μm indicates a higher atmospheric opacity at the shorter wave band, which is difficult to reconcile with current models.

In an attempt to explain its peculiar eclipse spectrum, Deming et al. (2011) hypothesized that CoRoT-2b might have equal opacity at 3.6 and 4.5 μm (and lower opacity at 8 μm) due to a haze of—as yet unknown—μm-sized particles. One may similarly explain the unusual WASP-12b transit spectrum in terms of a haze of slightly smaller particles, such that the opacity drops from 3.6 to 4.5 μm. Given the large difference in transit depth between the two wave bands, this hypothesis also requires a large atmospheric scale height at the planet's terminator.

5.2. Eclipses

At 3.6 μm, the eclipse depth is $\sim 1\sigma$ lower using the polynomial fit as compared to the decorrelation. If we fit the two eclipses individually using the polynomial IPSV fit, we obtain depths of 0.0030 (highly correlated residuals) and 0.0038 (incomplete egress), respectively. Given the prior measurement of 0.0038(1) by Campo et al. (2011), we adopt the larger value from the Gaussian decorrelation for our analysis (this choice does not significantly affect any of our conclusions).

At 4.5 μm, we obtain comparable χ^2 values and eclipse depths regardless of our treatment of systematics. In both cases they are consistent with the Campo et al. (2011) value: 0.0038(2).

In all cases our error estimates are somewhat larger than the Campo et al. (2011) estimates. We observed the same planetary system with the same instrument, so one expects the same eclipse depths and uncertainties. The only significant change in the detector is that our observations occurred after the fall 2010 change in heater cycling, as mentioned in Section 2. This means that the short-term telescope jitter for our observations has a period of 30 minutes rather than 1 hr, and the amplitude of

the centroid excursions—and hence flux variations—is reduced by a factor of two (for reference, the ingress/egress time for WASP-12b is 21 minutes, and the total transit/eclipse duration is 3 hr).

It is conceivable that the near coincidence between the centroid jitter half-period and the ingress/egress timescale leads to greater residual systematics in our data than in the Campo et al. (2011) time series. However, our eclipse depths are based on two occultations, and we have a much longer baseline of observations to help us correct for detector systematics, characterize noise properties, and estimate uncertainties (see Section 4.2). Our MCMC error estimates are similar to those of Campo et al. (2011), but residual-binning and prayer-bead analyses produce eclipse depth uncertainties more than $2\times$ larger than the MCMC. This means that there is still red noise present in our residuals, and the prayer-bead analysis is a more realistic estimate of our parameter uncertainties.

5.2.1. Dayside Emergent Spectrum

If one assumes solar composition, the relative eclipse depths at 3.6 and $4.5\,\mu$m can probe the temperature versus pressure (T-P) profile of the planet. Water vapor absorbs less at $3.6\,\mu$m than at $4.5\,\mu$m, so the shorter wave band will have a higher brightness temperature if the temperature is locally dropping with height or vice versa (e.g., Burrows & Orton 2010).

If the eclipse depth is measured at sufficiently many wavelengths, one may hope to simultaneously constrain a planet's atmospheric composition and T-P profile (e.g., Madhusudhan & Seager 2009). The high C/O chemistry invoked by Madhusudhan et al. (2011) was the result of an abnormally low eclipse depth at $4.5\,\mu$m in the Campo et al. (2011) data, which was interpreted as being due to CO absorption in the planet's relatively cool upper atmosphere. More recently, Kopparapu et al. (2011) used a photochemical model to study the disequilibrium chemistry of WASP-12b, confirming that CO would be enhanced in a high C/O composition atmosphere.

In Figure 7, we compare the near- to mid-IR broadband spectrum of WASP-12b to various one-dimensional (1D) radiative transfer models (Burrows et al. 2007, 2008). We vary the abundance of CO as a proxy for varying the C/O ratio. Our eclipse depths are consistent with those of Campo et al. (2011), so we still favor models with enhanced CO ($10 \times$ solar) and a weak inversion for this planet, in agreement with Madhusudhan et al. (2011). We also find that we can obtain an equally good fit to the data by reducing H_2O to 1% solar abundance and partitioning carbon evenly between CO and CH_4. Given the caveat that these models are in radiative—but not chemical—equilibrium, the crux of fitting WASP-12b's unique dayside spectrum is to suppress H_2O with respect to CO.

Our larger error bars, however, make these composition statements marginal: the solar composition model with modest inversion (the green line in Figure 7) is only worse by $\Delta\chi^2 = 8$ for 8 degrees of freedom. Furthermore, neither the standard composition nor the enhanced CO scenario is consistent with the relative transit depths at 3.6 and $4.5\,\mu$m (see Figure 6 and previous section). It may be possible to reconcile these two measurements if the atmospheric composition is grossly different at the day–night terminator than near the sub-stellar point.

5.3. Ellipsoidal Variations

The best-fit ellipsoidal variations at $3.6\,\mu$m are consistent with the predicted amplitude of 2×10^{-4}, but as one can see
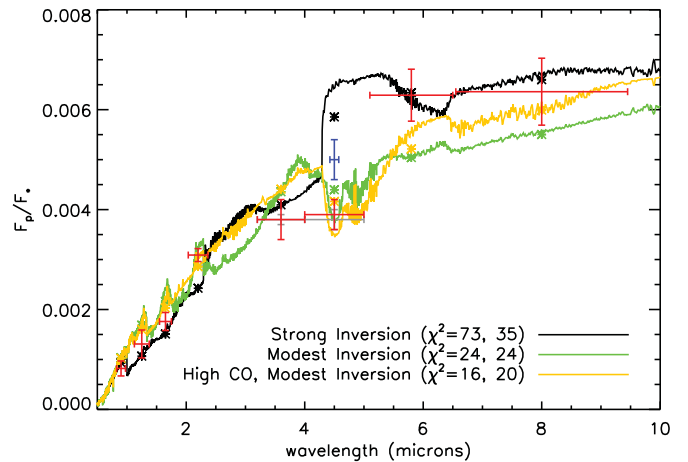

Figure 7. Dayside emergent spectrum of WASP-12b. The colored lines show various 1D atmospheric models (Burrows et al. 2007, 2008), while the red points show the measured secondary eclipse depths. From left to right, the data are z band (López-Morales et al. 2010); J, H, and K_s band (Croll et al. 2011), IRAC channels 1 and 2 (this study); and IRAC channels 3 and 4 (Campo et al. 2011). Note that the K_s-band eclipse depth and K–H eclipse color have been confirmed by Zhao et al. (2011) and Crossfield et al. (2012), respectively. The Campo et al. (2011) eclipse depths at 3.6 and $4.5\,\mu$m are shown in gray. The blue point shows the $4.5\,\mu$m eclipse depth if ellipsoidal variations are set to zero, the "null hypothesis." In the legend, the first χ^2 value for each model is for the fiducial analysis (including ellipsoidal variations), the second value is for the null hypothesis (setting ellipsoidal variations to zero). The enhanced CO model (yellow line) offers the best fit, but the solar composition model (green line) is not significantly worse.

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

from the relative uncertainty (or from glancing at Figure 4), they are not robustly detected: the χ^2, residuals and remaining astrophysical parameters do not change significantly if A_{ellips} is set to zero.

As shown in Figure 5, however, we clearly detect power at $4.5\,\mu$m in the second cosine harmonic, $\cos(2\alpha)$, consistent with the prediction of ellipsoidal variations due to the prolate shape of the planet (Li et al. 2010; Leconte et al. 2011; Budaj 2011). The semi-amplitude of the variations, however, is $1.2(2) \times 10^{-3}$ using either decorrelation or polynomial IPSV-removal, approximately $6\times$ the predicted value.

Since ellipsoidal variations are primarily a geometrical effect, it is difficult to understand how the measured amplitude could be so different at 3.6 and $4.5\,\mu$m. The upper layers of the atmosphere should be more distorted; the deeper layers more spherical. The relative strengths of ellipsoidal variations at the two wave bands imply that the $4.5\,\mu$m flux is originating from much higher up in the planet's atmosphere than the $3.6\,\mu$m flux. The simplest way to do this is for the atmosphere to have a greater opacity at $4.5\,\mu$m than $3.6\,\mu$m. But the relative transit depths indicate exactly the opposite, as described above and shown in Figure 6.

Alternatively, it is possible that detector systematics still present after IPSV-removal attenuate the ellipsoidal signal at $3.6\,\mu$m or enhance it at $4.5\,\mu$m. The raw photometry shown in Figure 1 implies that the $4.5\,\mu$m light curve is more trustworthy of the two. We therefore begin by assuming that the ellipsoidal signal at $4.5\,\mu$m is entirely astrophysical in nature, then consider the opposite scenario.

5.3.1. Interpreting the Ellipsoidal Variations at 4.5 μm

If we take the $4.5\,\mu$m ellipsoidal variations at face value, they have surprising implications for the planet's shape. The
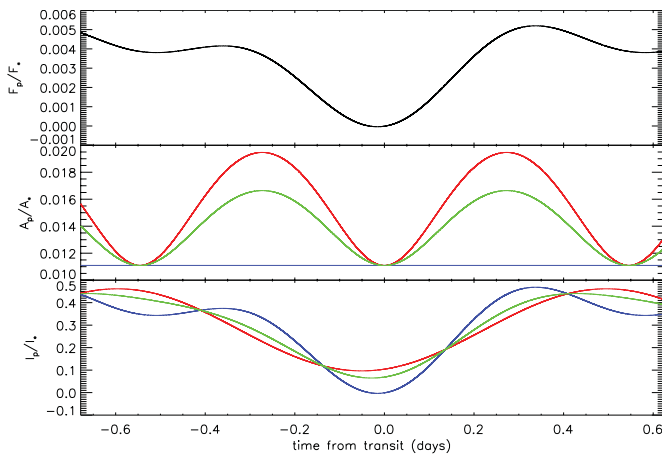
Figure 8. Our best-fit 4.5 μm phase variations (top panel) can be modeled as a geometrical component due to the planet's changing projected area (middle panel) and a thermal component due to longitudinal variations in the planet's brightness (bottom panel). For the red lines all of the 2α power is attributed to the planet's ellipsoidal shape; this hypothesis can be ruled out by the transit morphology (Figure 9). The blue lines result from assuming that the planet is spherical in shape; this hypothesis leads to unphysical brightness variations on the planet. The green lines show the middle path: most of the 2α power is due to the planet's shape, but the thermal component also contributes a bit. Note that this a posteriori analysis has limitations: in our astrophysical model, the thermal and ellipsoidal components of phase variations were added, while strictly speaking the planet's intensity and cross-sectional area should be multiplied.

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

dimensions of a prolate planet may be described by its short and long radii, denoted by R_p and R_{long}, respectively. The third dimension (parallel to the system's angular momentum vector) is assumed to be R_p because we are neglecting rotational effects, which tend to produce oblate planets. Note that WASP-12b is on a very short period orbit, so—if tidally locked—it has a rotation rate only a factor of two slower than Jupiter or Saturn. As a result of this rotation, Budaj (2011) estimates the planet's polar radius to be 2.5% shorter than its lateral equatorial radius,[19] so strictly speaking WASP-12b is a triaxial ellipsoid, but this does not affect the analysis below because of the planet's edge-on orbit, and our relatively short baseline of observations is insensitive to the spin precession of the planet (Carter & Winn 2010b).

At conjunction—either transit or eclipse—we are seeing the planet's smallest projected area (πR_p^2, in the case of a perfectly edge-on orbit). The projected area of an ellipsoid on an edge-on orbit varies as (e.g., Vickers 1996)

$$A_P(\alpha) = \pi R_p^2 \sqrt{\cos^2 \alpha + \left(\frac{R_{\mathrm{long}}}{R_p}\right)^2 \sin^2 \alpha}. \qquad (5)$$

For $R_{\mathrm{long}}/R_p \approx 1$, the changes in projected area follow a $A_P \propto \cos(2\alpha)$ shape and this is in fact how we modeled them; for more severe elongations, the peaks become broader and the troughs narrower, until a limiting case of $A_P \propto |\sin \alpha|$.

If we interpret all of the power in the second cosine harmonic as being due to the changing cross-sectional area of the planet (the red curves in Figure 8), we may estimate the planet's aspect ratio as

$$R_{\mathrm{long}}/R_p \approx 1 + 2A_{\mathrm{ellips}} \langle F_*/F_{\mathrm{day}} \rangle = 1.8(1). \qquad (6)$$
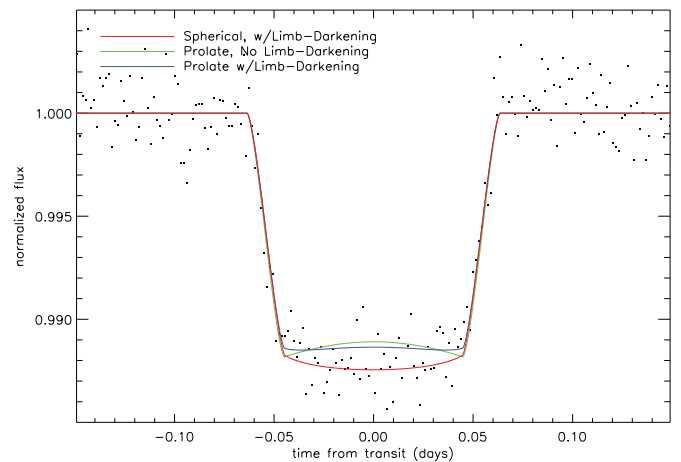
Figure 9. WASP-12b transit light curve at 4.5 μm; the data have been binned for plotting. The red line shows our fiducial model: a spherical planet with nonlinear limb darkening of the star. The green line shows the effect of the planet's changing cross-sectional area (but not its changing shape). The blue line shows how limb darkening partially washes out this signal. For the two prolate planet models, we use $R_{\mathrm{long}}/R_p = 1.8$, the most extreme scenario supported by the phase variations.

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

The predicted aspect ratio for the planet is $R_{\mathrm{long}}/R_p = 1.1$, while the aspect ratio for a Roche lobe is $3/2 = 1.5$.

It is worth noting that a planet with an aspect ratio of 1.8 would have a 13% larger projected area at the start and end of transit as compared to transit center, resulting in a w-shaped transit, in the absence of stellar limb darkening (note that this differs from changes in transit morphology due to an oblate planet discussed in Carter & Winn 2010a). Since the transit depth of WASP-12b is approximately 1%, this shape-induced transit effect would come in at the 1.3×10^{-3} level and might be detectable in our current data. In Figure 9, we estimate the expected transit morphology by treating the planet as a sphere with variable cross-sectional area. The red line shows our fiducial model: a spherical planet with nonlinear limb darkening of the star. The green line shows the effect of the planet's changing cross-sectional area (but not its changing shape). The blue line shows how limb darkening partially washes out this signal. Figure 9 implies that our data rule out the most extreme prolate toy model, but clearly the treatment of limb darkening is important here.

If we instead interpret the phase curve as being caused entirely by longitudinal brightness variations on a spherical planet (the blue curves in Figure 8), it can be inverted into a longitudinal intensity map of the planet, following Cowan & Agol (2008). Because of the low-pass filtering that occurs in the map → light curve convolution, the deconvolution will enhance the highest frequency terms present in the light curve (e.g., Figure 2 of Cowan & Agol 2009). In the current case, the resulting map has two prominent temperature peaks: one at the dawn terminator and one at the dusk terminator. More importantly, the only way to simultaneously fit the bright terminators and dark night side is by having negative intensity at the anti-stellar point, clearly an unphysical solution.

Another argument against a spherical planet is that while the thermal phase variations of a spherical planet will in general contain power in the second harmonic, there is no reason for it to all appear in the cosine rather than sine term.[20] When we run fits

---

[19] Note, furthermore, that the sub-stellar and anti-stellar planetary radii are not equal. Nevertheless, the dominant effect is the planet's prolate shape.

[20] This is in stark contrast to the first harmonic, where one expects most of the power to be in the cos α term due to the extreme day–night forcing.

allowing for the phase of the $\cos(2\alpha)$ term to vary, the offset is consistent with zero at the 2σ level, with the largest offset being less than $6°$. This indicates that the flux is peaking at quadrature, as expected for ellipsoidal variations. To our knowledge, there is no reason to expect such a temperature profile if the planet is spherical.

If the planet is prolate, however, one might expect something akin to the gravity darkening/brightening seen in some binary stars: the sub-stellar and anti-stellar points on the planet are farther from the center of the planet than the terminator, leading to lower surface gravity. The lower surface gravity at the sub-stellar and anti-stellar regions might lead to cooler temperatures, all things being equal (von Zeipel 1924). But this would only affect the intrinsic component of the planet's power budget, expected to be insignificant for hot Jupiters.

Finally, it is possible that the planet's shape affects the circulation of its atmosphere in more subtle ways, leading to relatively hot regions at the dawn and dusk terminators. This brings us to our favored astrophysical interpretation of the ellipsoidal variations. The planet may have an aspect ratio of $R_{\text{long}}/R_p \approx 1.5$ (somewhat easing the transit morphology constraints described above), with an additional enhancement of the $\cos(2\alpha)$ power because of a relatively hot day–night terminator (the green curve in Figure 8). Since the atmospheric dynamics on severely non-spherical planets has not yet been addressed in the literature, it is difficult to say whether this scenario is reasonable.

5.3.2. The Null Hypothesis at 4.5 μm

If the amplitude of ellipsoidal variations, A_{ellips}, is set to zero at 4.5 μm, the residuals become more correlated (as one would guess from Figure 5), increasing the red noise and hence uncertainties in the other astrophysical parameters. Furthermore, including terms up to sixth order in the polynomial IPSV-removal does not obviate the need for the ellipsoidal term, and the same signal is detected in the Gaussian decorrelation version of the analysis. The χ^2 is worse when ellipsoidal variations are ignored ($\Delta\chi_R^2 = 0.007$ for either the decorrelation or polynomial fit).

However, it is conceivable that the 4.5 μm ellipsoidal signal is in fact uncorrected detector systematics. It is therefore worth briefly considering the astrophysical implications of this scenario.

The most obvious implication of the null hypothesis is that the Roche-filling upper atmosphere of the planet need not be optically thick. Since the predicted amplitude of ellipsoidal variations were only at the 2σ level, the null hypothesis is consistent with the predictions of Li et al. (2010), Leconte et al. (2011), and Budaj (2011).

Furthermore, the ellipsoidal variations have minima at inferior and superior conjunction, so the null hypothesis causes the transit depth and eclipse depths to significantly increase. Notably, the 4.5 μm transit depth becomes 0.0126(4), commensurate with that at 3.6 μm. The null hypothesis transit depths are consistent with a short (night-like) scale height and solar composition (solid black line in Figure 6).

The deeper 4.5 μm eclipse depth, $F_{\text{day}}/F_* = 0.0050(4)$, no longer favors the enhanced CO scenario (yellow line in Figure 6) invoked by Madhusudhan et al. (2011). The fit to the solar composition model with modest inversion (green line in Figure 7) is only worse by $\Delta\chi^2 \approx 4$, for 8 degrees of freedom.

5.4. Thermal Phase Variations

We model thermal phases using only first harmonics ($\cos\alpha$ and $\sin\alpha$), while our parameterization of ellipsoidal variations is a second harmonic ($\cos 2\alpha$). These functions are by definition orthogonal, so it is not surprising that our conclusions about thermal phase variations (A_{therm}, α_{max}) described below are not significantly affected by the presence or absence of ellipsoidal variations discussed above.

That said, the amplitude and offset we obtain for the thermal phase variations depends on which IPSV-removal scheme we use. Unfortunately, it is not clear how to perform a direct model comparison between decorrelation and polynomial fits using the BIC. The Gaussian decorrelation can be thought of as having a large number of free parameters and a somewhat smaller number of additional constraining equations, but it is not obvious how to estimate its degrees of freedom. As discussed by Ballard et al. (2010), the point-by-point decorrelation does a great job of removing the short-term jitter and long-term detector drift, but may also remove any longer-term astrophysical signal. (This did not interfere with their goal of searching for transits.) Insofar as diurnal phase variations are the most gradual astrophysical signal in our study, it is not surprising that it is the most dependent on IPSV-removal.

The polynomial fit leads to χ^2 values at least as good as—and sometimes significantly better than—the Gaussian decorrelation. More importantly, the scatter in the residuals on the critical 21 minute timescale is lower for the polynomial fits, indicating that this method is doing a better job of removing correlated noise.

Our 4.5 μm pixel sensitivity map obtained from Gaussian decorrelation (top right panel of Figure 3) shows a valley at $y \approx 14.75$; this does not follow the usual pattern of sensitivity decreasing toward pixel edges. (Note that we observe the same ripples in sensitivity as Ballard et al. 2010, but those occur on a much smaller spatial scale—and exhibit a much smaller amplitude—than the $y \approx 14.75$ valley.)

Finally, the middle left panel of Figure 3 suggests that the decorrelation method has overcorrected the 3.6 μm system flux near superior conjunction: the eclipse bottoms—which should be flat since the planet is hidden from view—slope upward toward the central transit.

We therefore argue that the decorrelation has filtered out much of the phase variations along with the systematics and adopt the polynomial-fitted thermal phase parameters in what follows.[21]

Following Cowan & Agol (2011b), we estimate the hemispheric effective temperatures to be $T_{\text{day}} = 2928(97)$ K and $T_{\text{night}} = 983(201)$ K, where the uncertainties include an estimate of systematic errors in going from brightness temperatures to effective temperatures.[22] This implies a Bond albedo of $A_B = 0.25$ (presumably due to Rayleigh scattering and/or reflective clouds) and very low heat recirculation efficiency, $\varepsilon < 0.1$, at 1σ (see Figure 10).[23] Note that our 1D radiative transfer models used for interpreting transit and eclipse

---

[21] For completeness, we include the thermal phase parameters resulting from the decorrelation in Tables 2 and 3. That analysis leads to smaller phase amplitudes than the polynomial fit. If taken at face value, this implies lower albedo and higher heat transport efficiency.

[22] WASP-12b has a sub-stellar equilibrium temperature of $T_0 = 3555(132)$ K, a no-albedo, no-recirculation dayside temperature of $T_{\varepsilon=0} = (2/3)^{1/4}T_0 = 3213(119)$ K, and a no-albedo full-recirculation global temperature of $T_{\text{uni}} = (1/4)^{1/4}T_0 = 2514(92)$ K (e.g., Cowan & Agol 2011b).

[23] If one presumes that advection is the dominant mode of heat transport, then $\varepsilon \approx \tau_{\text{rad}}/(\tau_{\text{adv}} + \tau_{\text{rad}})$, where τ_{rad} and τ_{adv} are the characteristic radiative and advective timescales at the mid-IR photosphere.

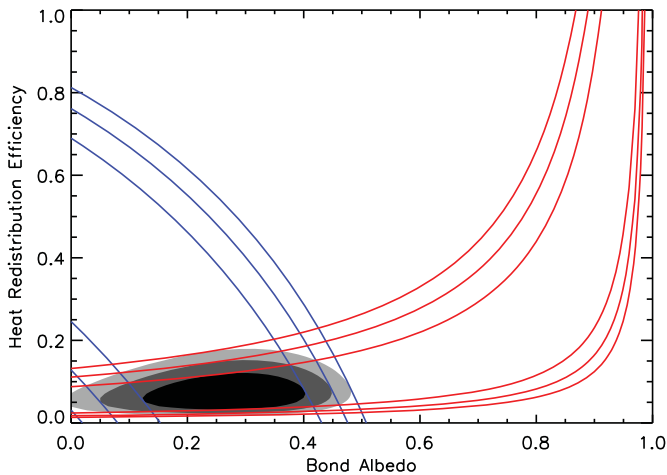**Figure 10.** 1σ, 2σ, and 3σ constraints on the Bond albedo and recirculation efficiency of WASP-12b from thermal eclipse (blue) and phase variation (red) observations, using the parameterization of Cowan & Agol (2011b). The gray scale shows the confidence intervals for the combined constraints.

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

spectra are cloud-free and have an albedo lower than this inferred value.

Space-based optical secondary eclipse depths have been measured for a handful of hot Jupiters. If the planet's equilibrium temperature is sufficiently low, such measurements provide an unambiguous estimate of geometric albedo, A_g: 0.04(5) for HD 209458b (Rowe et al. 2008), 0.32(3) for Kepler-7b (Demory et al. 2011), 0.10(2) for Kepler-17b (Désert et al. 2011a), 0.30(8) for KOI-196b (Santerne et al. 2011), and 0.025(7) for TrES-2b (Kipping & Spiegel 2011). Acknowledging that WASP-12b is more than 1000 K hotter than any of those planets, a Bond albedo of 0.25 is well within the observed range for hot Jupiters.

Assuming gray albedo and a Lambertian scattering phase function ($A_B = 3/2A_g$; Hanel et al. 1992), the reflected-light secondary eclipse of WASP-12b should have a depth of 2.4×10^{-4}, comparable to current ground-based precision for this target. In practice, most planets exhibit an opposition surge (making them disproportionately bright at superior conjunction) and Rayleigh scattering, so the actual contrast ratio in blue optical wave bands is likely more favorable than this estimate.

Assuming solar atmospheric composition and hence opacity, the $3.6\,\mu$m thermal flux should originate from deeper in the atmosphere than any other mid-IR wave band. Insofar as radiative times increase monotonically with pressure, we may therefore expect the $3.6\,\mu$m phase variations to be muted compared to the $4.5\,\mu$m phase variations, and the phase offset should be greater at the shorter wavelength (e.g., Figure 9 of Burrows et al. 2010).

The hot spot offset is 53(7)° east of the sub-stellar point at $3.6\,\mu$m. While not as extreme as υ-Andromeda b (Crossfield et al. 2010), it is difficult to reconcile our large phase offset and large amplitude at $3.6\,\mu$m.[24] It is worth noting, however, that there are highly correlated residuals near the purported peak ($\sim$0.35 day after transit) which may be partially responsible for the large offset.

On the other hand, we find that the $4.5\,\mu$m phase amplitude and offset, 16(4)°E, are consistent with a Cowan & Agol (2011a) model with heat transport efficiency of $\epsilon \equiv \tau_{\rm rad}/\tau_{\rm adv} \approx 0.1$. To put this in context, phase variations and eclipse timing offset of HD 189733b at $8\,\mu$m indicate $\epsilon \approx 0.7$ (Agol et al. 2010).

5.4.1. Implications of Thermal Phase Variations

Cowan & Agol (2011b) noted that the dayside temperatures of the hottest short-period giant planets are very close to the theoretical upper limit of no albedo and no recirculation. This was in contrast to run-of-the-mill hot Jupiters (e.g., HD 189733b, HD 209458b), which exhibit a variety of albedos/recirculation efficiencies, albeit consistent with generally low albedos ($A_B < 0.3$). In a statistical study of *Kepler* planetary candidates, Coughlin & Lopez-Morales (2012) also found generally low albedos for hot Jupiters based on optical secondary eclipses.

The amplitude of the phase variations for WASP-12b depends on the details of the systematics correction, but—for reasons stated at the start of Section 5.4—we favor the polynomial fit, which implies a large day–night temperature contrast, and a non-zero Bond albedo. This suggests that the difference between the hottest short-period giant planets and other hot Jupiters is not albedo, but recirculation efficiency. (Differences in albedo may very well explain the differences in dayside effective temperature among the remaining hot Jupiters, however.)

What could make the hottest of hot Jupiters poor heat recirculators? There are two classes of solutions: decreasing either the planet's characteristic advective frequency or radiative timescale.[25]

More magnetic drag. Assuming these planets have magnetic fields, the movement of ionized alkali metals through the field produces drag that is collisionaly imparted on the dominant neutral species (presumably H and He). Hotter planets should have more ionized species, more drag and therefore a harder time advecting heat to their night side (Perna et al. 2010). Because of the nonlinear dependence of ionization on temperature, this effect could lead to sudden changes in dynamical regime as one considers increasingly hot planets. The scaling relations of Menou (2011) indicate that the temperature above which magnetic drag severely curtails heat transport is inversely related to the planet's magnetic field strength. Even the weakest field they considered in their study, 3 Gauss, would result in very low recirculation efficiency for a planet as hot as WASP-12b.

Shorter radiative times. Following the argument of Cowan & Agol (2011b), the radiative relaxation time of a parcel of gas scales as $\tau_{\rm rad} \propto T^{-3}$ (Iro et al. 2005; Seager et al. 2005). But zonal wind speeds may also increase with the amplitude of the diurnal forcing. If one assumes that the wind speeds have a fixed Mach number, they should scale as $\upsilon_{\rm wind} \propto T^{1/2}$, and therefore the advective time should scale as $\tau_{\rm adv} \propto T^{-1/2}$ (a more detailed scaling analysis leads to the same temperature dependence; Menou 2011). The stronger dependence on temperature of radiative time compared to advective time implies that—all things being equal—hotter planets should be less efficient at balancing their day–night temperature contrast. This effect should cause the heat transport efficiency to gradually decrease as one considers increasingly hot planets.

---

[24] It is tempting to attribute the early peak of the $3.6\,\mu$m phase curve to Doppler beaming, which would produce a peak in flux when the star is moving toward us, a quarter period before superior conjunction (Loeb & Gaudi 2003). The expected amplitude of this signal, however, is only 7.5×10^{-7} on the Rayleigh–Jeans tail of the star.

[25] One can imagine more exotic means of transporting energy (e.g., gravity waves; Watkins & Cho 2010) but most hydrodynamical simulations suggest that horizontal energy transport on hot Jupiters is primarily a matter of advection.

Weaker greenhouse. Atmospheric opacity is typically greater at the thermal wavelengths of emergent radiation than at the visible wavelengths of incident radiation. For the hottest planets, the blackbody peak of thermal emission approaches the peak of their host star, so the opacities of the incoming and outgoing streams converge. In that limit, one expects the thermal photosphere and the optical deposition depth to be one and the same: thermal radiation should escape the atmosphere just as easily as the incoming stellar radiation came in. As with the scalings above, this effect should lead to gradually decreasing recirculation efficiency as a function of increasing planet temperature.

More observations of thermal eclipses and phase variations for hot Jupiters—especially those near the $T_0 \approx 2700$ K transition—will be necessary to distinguish between the magnetohydrodynamic and the radiative timescale arguments.

6. CONCLUSIONS

We obtained *Warm Spitzer* full-orbit phase observations of WASP-12b at 3.6 and 4.5 μm, allowing us to measure the transit depths, eclipse depths, thermal and ellipsoidal phase variations at both wavelengths. We are able to push *Warm Spitzer* photometry to within 10%–20% of the Poisson limit, but there are two important caveats.

1. Removing IPSVs from the data is inherently a model-dependent endeavor. This means that we must specify not only an IPSV model, but also an astrophysical model before getting close to the quoted precision. The simultaneous fit to astrophysical and systematic effects makes it difficult to produce a "clean" light curve independent of astrophysical assumptions. For example, we obtain very different thermal phase variation parameters depending on how we correct for systematics, and it is difficult to distinguish between these scenarios based solely on goodness of fit. Instead, we must resort to a number of indirect clues as to which IPSV-removal scheme is more trustworthy.

2. There is still red noise in our residuals, no matter how we remove IPSVs. This remaining red noise is the dominant source of uncertainty for all of our astrophysical parameters.

We find that WASP-12b exhibits large-amplitude thermal phases—indicative of poor day–night heat transport and a moderate Bond albedo—but also an unexpectedly large phase offset at 3.6 μm. We do not detect ellipsoidal variations at 3.6 μm, while we detect an unexpectedly strong signal at 4.5 μm. This leads us to two possible hypotheses.

1. If we take the 4.5 μm ellipsoidal variations at face value, we find: deeper transits at 3.6 μm as compared to 4.5 μm, inconsistent with either solar or enhanced CO models; eclipse depths consistent with previous studies. If the 4.5 μm ellipsoidal variations are astrophysical in nature, it indicates that the planet is far more distorted than predicted, and exhibits a bright terminator. In this scenario, the 3.6 μm ellipsoidal variations are attenuated due to detector systematics, possibly throwing off the 3.6 μm transit depth as well.

2. If instead we presume that the 4.5 μm ellipsoidal variations are caused by detector systematics and set them to zero—the null hypothesis—we find: transit depths consistent with a solar composition and short atmospheric scale height at the planet's terminator; eclipse depths consistent with a solar composition and a modest temperature inversion; ellipsoidal variations in line with predictions.

The null hypothesis is attractive in its simplicity, but requires that we were very unlucky; follow-up *Warm Spitzer* observations would have different systematics (the PSF would fall on different regions of the pixels) and could settle the question of ellipsoidal variations. It is likely that near-infrared transit spectroscopy could break the composition degeneracy, or at least determine the atmospheric structure of WASP-12b; if the planet has a short scale height at the terminator it will lend credence to the null hypothesis. Further optical transit photometry will be useful in pinning down the transmission spectrum and refining geometrical parameters; if $a/R_* < 3$, then the planet could very well be more distorted than predicted, making the large ellipsoidal variations more plausible. Optical eclipse measurements from the ground or from space might confirm the moderate albedo of the planet.

The planet is hypothesized to be losing mass to its host star. If this is indeed the case, the presence of an accretion disk, accretion stream, and impact hot spot may necessitate a more holistic model to properly interpret observations.

REFERENCES

Agol, E., Cowan, N. B., Knutson, H. A., et al. 2010, ApJ, 721, 1861
Ballard, S., Charbonneau, D., Deming, D., et al. 2010, PASP, 122, 1341
Budaj, J. 2011, AJ, 141, 59
Burrows, A., Budaj, J., & Hubeny, I. 2008, ApJ, 678, 1436
Burrows, A., Hubeny, I., Budaj, J., Knutson, H. A., & Charbonneau, D. 2007, ApJ, 668, L171
Burrows, A., & Orton, G. 2010, in Exoplanets, ed. S. Seager (Tucson, AZ: Univ. Arizona Press), 419
Burrows, A., Rauscher, E., Spiegel, D. S., & Menou, K. 2010, ApJ, 719, 341
Campo, C. J., Harrington, J., Hardy, R. A., et al. 2011, ApJ, 727, 125
Carter, J. A., & Winn, J. N. 2009, ApJ, 704, 51
Carter, J. A., & Winn, J. N. 2010a, ApJ, 709, 1219
Carter, J. A., & Winn, J. N. 2010b, ApJ, 716, 850
Chan, T., Ingemyr, M., Winn, J. N., et al. 2011, AJ, 141, 179
Charbonneau, D., Allen, L. E., Megeath, S. T., et al. 2005, ApJ, 626, 523
Claret, A. 2000, A&A, 363, 1081
Coughlin, J. L., & Lopez-Morales, M. 2012, AJ, 143, 39
Cowan, N. B., & Agol, E. 2008, ApJ, 678, L129
Cowan, N. B., & Agol, E. 2009, in IAU Symp. 253, Transiting Planets (Cambridge: Cambridge Univ. Press), 544
Cowan, N. B., & Agol, E. 2011a, ApJ, 726, 82
Cowan, N. B., & Agol, E. 2011b, ApJ, 729, 54
Cowan, N. B., Agol, E., & Charbonneau, D. 2007, MNRAS, 379, 641
Croll, B., Lafreniere, D., Albert, L., et al. 2011, AJ, 141, 30
Crossfield, I. J. M., Hansen, B. M. S., & Barman, T. 2012, arXiv:1201.1023
Crossfield, I. J. M., Hansen, B. M. S., Harrington, J., et al. 2010, ApJ, 723, 1436
Deming, D., Knutson, H., Agol, E., et al. 2011, ApJ, 726, 95
Demory, B.-O., Seager, S., Madhusudhan, N., et al. 2011, ApJ, 735, L12
Désert, J.-M., Charbonneau, D., Demory, B.-O., et al. 2011a, ApJS, 197, 14
Désert, J.-M., Sing, D., Vidal-Madjar, A., et al. 2011b, A&A, 526, A12
Faigler, S., & Mazeh, T. 2011, arXiv e-prints
Fazio, G. G., Hora, J. L., Allen, L. E., et al. 2004, ApJS, 154, 10
Fortney, J. J., Shabram, M., Showman, A. P., et al. 2010, ApJ, 709, 1396
Fossati, L., Haswell, C. A., Froning, C. S., et al. 2010, ApJ, 714, L222
Hanel, R. A., Conrath, B. J., Jennings, D. E., & Samuelson, R. E. (ed.) 1992, Exploration of the Solar System by Infrared Remote Sensing (Cambridge Planetary Science Series, No. 7; Cambridge: Cambridge Univ. Press), 484

Hebb, L., Collier-Cameron, A., Loeillet, B., et al. 2009, ApJ, 693, 1920
Husnoo, N., Pont, F., Hébrard, G., et al. 2011, MNRAS, 413, 2500
Iro, N., Bézard, B., & Guillot, T. 2005, A&A, 436, 719
Kipping, D. M., & Spiegel, D. S. 2011, arXiv e-prints
Knutson, H. A., Charbonneau, D., Allen, L. E., et al. 2007, Nature, 447, 183
Knutson, H. A., Charbonneau, D., Burrows, A., O'Donovan, F. T., & Mandushev, G. 2009a, ApJ, 691, 866
Knutson, H. A., Charbonneau, D., Cowan, N. B., et al. 2009b, ApJ, 690, 822
Knutson, H. A., Madhusudhan, N., Cowan, N. B., et al. 2011, ApJ, 735, 27
Kopparapu, R. K., Kasting, J. F., & Zahnle, K. J. 2011, arXiv e-prints
Kurucz, R. L. 1979, ApJS, 40, 1
Kurucz, R. L. 2005, Mem. Soc. Astron. Ital. Suppl., 8, 14
Lai, D., Helling, C., & van den Heuvel, E. P. J. 2010, ApJ, 721, 923
Leconte, J., Lai, D., & Chabrier, G. 2011, A&A, 528, A41
Li, S.-L., Miller, N., Lin, D. N. C., & Fortney, J. J. 2010, Nature, 463, 1054
Llama, J., Wood, K., Jardine, M., et al. 2011, MNRAS, 416, L41
Loeb, A., & Gaudi, B. S. 2003, ApJ, 588, L117
López-Morales, M., Coughlin, J. L., Sing, D. K., et al. 2010, ApJ, 716, L36
Maciejewski, G., Errmann, R., Raetz, S., et al. 2011, A&A, 528, A65
Madhusudhan, N., Harrington, J., Stevenson, K. B., et al. 2011, Nature, 469, 64
Madhusudhan, N., & Seager, S. 2009, ApJ, 707, 24
Mandel, K., & Agol, E. 2002, ApJ, 580, L171
Menou, K. 2011, arXiv e-prints
Morales-Calderón, M., Stauffer, J. R., Kirkpatrick, J. D., et al. 2006, ApJ, 653, 1454
Perna, R., Menou, K., & Rauscher, E. 2010, ApJ, 719, 1421
Pont, F., Knutson, H., Gilliland, R. L., Moutou, C., & Charbonneau, D. 2008, MNRAS, 385, 109
Pont, F., Zucker, S., & Queloz, D. 2006, MNRAS, 373, 231
Ragozzine, D., & Wolf, A. S. 2009, ApJ, 698, 1778
Rauscher, E., Menou, K., Seager, S., et al. 2007, ApJ, 664, 1199
Rowe, J. F., Matthews, J. M., Seager, S., et al. 2008, ApJ, 689, 1345
Santerne, A., Bonomo, A. S., Hébrard, G., et al. 2011, arXiv e-prints
Schwarz, G. 1978, Ann. Stat., 6, 461
Seager, S., Richardson, L. J., Hansen, B. M. S., et al. 2005, ApJ, 632, 1122
Sing, D. K., Pont, F., Aigrain, S., et al. 2011, MNRAS, 416, 1443
Stevenson, K. B., Harrington, J., Fortney, J., et al. 2011, arXiv e-prints
Vickers, G. T. 1996, Powder Technol., 86, 195
Vidotto, A. A., Jardine, M., & Helling, C. 2010, ApJ, 722, L168
von Zeipel, H. 1924, MNRAS, 84, 665
Watkins, C., & Cho, J. Y.-K. 2010, ApJ, 714, 904
Welsh, W. F., Orosz, J. A., Seager, S., et al. 2010, ApJ, 713, L145
Werner, M. W., Roellig, T. L., Low, F. J., et al. 2004, ApJS, 154, 1
Williams, P. K. G., Charbonneau, D., Cooper, C. S., Showman, A. P., & Fortney, J. J. 2006, ApJ, 649, 1020
Winn, J. N. 2010, in Exoplanets, ed. S. Seager (Tucson, AZ: Univ. Arizona Press), 55
Winn, J. N., Henry, G. W., Torres, G., & Holman, M. J. 2008, ApJ, 675, 1531
Winn, J. N., Holman, M. J., & Roussanova, A. 2007, ApJ, 657, 1098
Zhao, M., Monnier, J. D., Swain, M. R., Barman, T., & Hinkley, S. 2011, arXiv e-prints