

Variability in the atmosphere of the hot giant planet HAT-P-7 b

D. J. Armstrong^{1,2*}, E. de Mooij², J. Barstow³, H. P. Osborn¹, J. Blake¹ and N. Fereshteh Saniee¹

An exoplanet reflects and emits light as it orbits its host star, thus forming a distinctive phase curve^{1,2}. By observing this light, we can study the atmosphere and surface of distant planets. The planets in our Solar System show a wide range of atmospheric phenomena, including stable wind patterns, changing storms and evolving anomalies. Brown dwarfs also exhibit atmospheric variability^{3,4}. Such temporal variability in the atmosphere of a giant exoplanet has not yet been observed. HAT-P-7 b is an exoplanet with a known offset in the peak of its phase curve^{5,6}. Here we present variations in the peak offset ranging between $-0.086^{+0.033}_{-0.033}$ and $0.143^{+0.040}_{-0.037}$ in phase, which implies that the peak brightness repeatedly shifts from one side of the planet's substellar point to the other. The variability occurs on a timescale of tens to hundreds of days. These shifts in brightness are indicative of variability in the planet's atmosphere, and result from a changing balance of thermal emission and reflected flux from the planet's day-side. We suggest that variation in wind speed in the planetary atmosphere, leading to variable cloud coverage on the day-side and a changing energy balance, is capable of explaining the observed variation.

HAT-P-7 b is a hot Jupiter of radius $1.4 R_J$ that transits its host star with a period of 2.20 days⁷. It is extremely hot, with a dayside brightness temperature of 2,860 K and an equilibrium temperature of 2,200 K⁶. It was continuously observed for four years by the Kepler satellite⁸ at optical wavelengths. HAT-P-7 b has also been intensively observed at infrared wavelengths with the Spitzer satellite^{9,10}.

Both the optical and infrared phase curves of HAT-P-7 b have been studied previously^{5,6,9,11–13}. The optical phase curve exhibits a significant fraction of thermal emission (potentially as high as 77%⁶), due to the high temperature of the planet. The optical phase curve also presents a known eastward shift on average^{6,14}, which, in a thermal-emission-dominated phase curve, could be caused by the underlying global circulation pattern created by planetary-scale Rossby waves arising from day-side–night-side temperature differences¹⁵. Large-scale weather in the atmosphere of hot Jupiters is theoretically expected^{16–18}, however no previous searches for variations in the optical phase curve parameters with time have been made. Infrared variability is predicted to be of the order of 1%¹⁹ and has likewise not been observed. Previous work⁹ found a marginally significant increase in the secondary eclipse depth at $3.6 \mu\text{m}$ of 59% between measurements, but this was attributed to differences in the analysis technique. Spitzer observations of multiple eclipses spaced at significant intervals have been used to put an upper limit on the variability of HD 189733 b at $<2.7\%$ ^{20,21}, and have detected eclipse depth changes in the super-Earth 55 Cancri²².

However, the nature of those observations prevented continuous tracking of the variability.

We use four years of public Kepler data of the HAT-P-7 system. We detrend the lightcurve for instrumental effects (see Methods), and then combine the observations using a sliding bin of ten orbits using published ephemeris⁶. Next we fit each of the combined phase curves using a model that describes both planetary and stellar effects (see Methods). The model combines ellipsoidal variation and Doppler beaming on the star, planetary occultation, planetary brightness modelled as a Lambertian sphere with an offset, and a previously detected cosine third harmonic⁶. Binning over ten orbits allows us to build up significance for each fit by gaining more data-points, but averages out variations on timescales shorter than ten times the planetary period. This also removes stellar or instrumental variability on timescales significantly shorter than the bin size. Each successive bin starts one orbit later than the previous bin. As such, successive fit parameters are not fully independent, and hence care must be taken when assessing significance. Independent fit parameters produced from discrete bins are shown in Supplementary Fig. 1. We fix the planet mass and associated parameters to published values⁶, and fit for the planetary brightness amplitude A_p , secondary eclipse depth F_{ec} and peak brightness offset θ for each bin of ten planetary orbits. To assess the reliability of the measurements and determine the uncertainties, we use four separate methods, including a 'prayer bead' residual permutation method²³ and Markov chain Monte Carlo (MCMC), as described in the Methods.

Clear variation is seen in θ (Fig. 1). The variation spans from $-0.086^{+0.033}_{-0.033}$ to $0.143^{+0.040}_{-0.037}$ in phase, with a standard deviation (σ) of 0.033, and an average error on the fits of 0.022 in phase. There is also marginal variation in A_p , with a standard deviation of 14.6 ppm, but this is probably the result of systematic noise biasing the fits (see injection tests in Methods). The derived value of A_p is susceptible to systematic noise occurring near the transit; θ is more robust, however, as changing the peak location requires the noise to have an amplitude larger than the peak. The occultation depth F_{ec} does not vary within our errors (Supplementary Fig. 1).

Individual fit values are given in the Supplementary Information. Example phase curves and models are shown in Fig. 2, and further examples can be found in Supplementary Fig. 2. Since the Kepler data must be detrended to correct for instrument effects, we use different detrending techniques to assess the reliability of our observations. We also test for systematic biases by injecting a non-variable HAT-P-7 b phase curve into the lightcurves of two other stars. Each of these tests confirm the variations in θ and are described in the Methods.

Weather in HAT-P-7 b's atmosphere would be expected to contribute to both the observed shifts in θ , and to potential shifts

¹Department of Physics, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, UK. ²ARC, School of Mathematics & Physics, Queen's University Belfast, University Road, Belfast BT7 1NN, UK. ³Astrophysics Group, Department of Physics and Astronomy, University College London, London, NW1 2PS, UK. *e-mail: d.j.armstrong@warwick.ac.uk




Figure 1 | Variation in peak offset of HAT-P-7b phase curve with time. Derived from fitting to a sliding bin of ten planetary orbits. The horizontal black line shows the value obtained by fitting the entire dataset simultaneously. Negative offsets correspond to a movement towards the morning side of the planet, and are associated with an increased proportion of reflected light in the Kepler bandpass. Positive offsets represent movement to the afternoon side, and thus increased thermal emission in the bandpass. Error bars are 1σ errors on the fit parameters. BJD, Barycentric Julian date.

in A_p and F_{ecl} . Although we do not observe variability in A_p or F_{ecl} , we can limit the variations in A_p to $\pm 59\%$ at 3σ , and variations in F_{ecl} to $\pm 51\%$. This leaves a large potential for variability. Significant temperature differences could be expected if, for example, the circulation efficiency transporting heat to the planet's nightside was particularly variable, thereby causing energy to build up on the planet's dayside.

It is important to consider what physical scenarios can lead to such significant changes in peak offset. Here we implement a published semi-analytical model²⁴ to gain an understanding of the origins of the observed variation (see reference for full description). The model depends primarily on the planet's Bond albedo A_B , heat redistribution efficiency ε , cloud reflectivity boosting factor κ and condensation temperature T_C . These parameters and their effects are discussed in the Supplementary Information. Both thermal emission and reflected light are included. We assume positive values for ε , which corresponds to superrotating winds on the planet. This assumption is supported by current hot Jupiter atmosphere models¹⁵. The model allows the calculation of brightness maps on the planetary surface, and typical maps for both positive and negative Θ are shown in Fig. 3. We discuss the conditions that could lead to these cases below.

We calculate a grid of models in A_B [0–0.6], ε [0–20] and κ [1–40] space. For each model, we fit the offset Lambertian phase curve that we applied to the data above, thus extracting the corresponding A_p , Θ and F_{ecl} . We find that within the ranges tested, condensation temperatures between 1,600 K and 2,200 K can cause the observed Θ variations. Despite the brightness temperature of the planet being hotter than this range, we find that at the morning terminator the temperature can fall below these values, leading to increased cloud coverage. Outside this temperature range, clouds are either not

present or occur over most of the planet's dayside, resulting in a symmetric optical component and no negative Θ values in either case. An example grid is shown in Fig. 4. In the case found that requires the smallest albedo change ($T_C = 2,000$ K, $\varepsilon = 8$, $\kappa = 40$), the entire range of Θ we observe, excluding the three outlier points in Fig. 1 with $\Theta > 0.1$, can be explained by albedo shifts of only 0.05, although changes of the order of 0.1, combined with variation in ε , are more typical. This albedo shift produces variability in F_{ecl} in the model of at most 25 ppm, or 35% of the global fit value of 71.2 ppm⁶, which is smaller than the constraint of 51% we place above. At the above values of T_C , ε and κ , an albedo change of 0.1 corresponds to a peak temperature change of 100 K on the planet. As the sharp edge of the Planck function is near the edge of the Kepler bandpass for HAT-P-7b, this significantly changes the thermal flux emitted in the bandpass, and hence magnifies the change in observed peak offset. However, a wide range of parameter combinations can produce the observed values. Several model grids for a range of temperatures are given in Supplementary Figs 6 and 7.

Here we discuss what physical mechanisms could lead to the above parameter changes, and hence the observed variations in Θ . Circulation models of tidally locked hot Jupiters predict the presence of strong superrotating winds¹⁵, which are responsible for the observed eastward shifts of thermal hotspots for hot Jupiters with infrared phase curves¹⁴. Although dayside temperatures of HAT-P-7b as inferred from measured Spitzer transit depths ($\sim 2,800$ K) are likely to be too warm for possible condensates to form, the extremely fast winds can transport aerosols from the cooler nightside. Due to the short horizontal advection timescales involved, aerosols may persist in regions between the morning terminator and the substellar point—despite the high dayside temperatures—before eventually evaporating towards the evening side of the planet.




Figure 2 | Individual phase curves, with best-fitting model. The best-fitting peak offsets are, from left to right, 0.021, 0.076, 0.108 and -0.086 . For clarity, the data is binned to 0.01 in phase before plotting, however the model was fit to the unbinned data. The lower panels show the binned residuals after the best-fitting model is removed. Error bars are the standard error on the mean of the datapoints contributing to each bin.

Circulation models for planets like HAT-P-7b with equilibrium temperatures of 2,000–2,200 K have been published²⁵, and suggest that condensate species Al_2O_3 (corundum) and CaTiO_3 (perovskite) would condense out at around 100 mbar on the nightside of these planets. According to the same set of models, these condensates could persist in the atmosphere at the morning terminator, and would then be advected onto the dayside. The cloud-covered proportion of the planet between the dayside and morning terminator will be dictated by the relative timescales of advection and evaporation. Increased wind speeds would reduce advection timescales and allow the aerosols to travel further into the dayside atmosphere prior to evaporation; increased wind speeds would also move the thermal hotspot further towards the planet's evening terminator. A plausible mechanism for the phase offset variation in the Kepler lightcurves is therefore variation in the speed of the superrotating jet, which aperiodically transports a greater proportion of reflective aerosols onto the dayside and thus causes reflection from the cloud top to dominate the phase curve shape instead of the thermal contribution. The mechanism for such aperiodic variation in wind speed is unknown, but a possible driver may be the extreme tidal forces experienced by HAT-P-7b owing to its close orbit.

An increase in albedo on the dayside, due to a larger concentration of advected aerosols, will lead to a cooling of the upper atmosphere, which in turn will reduce the day–night temperature contrast and is therefore likely to weaken the superrotating flow²⁶. This will act to reduce the transport of aerosols to the dayside and thus create a feedback mechanism, resulting in oscillations in wind strength, dayside albedo and dayside temperature, which explains the observed phenomenon of the shifting phase curve peak.

Previous studies of the phase curve of HAT-P-7b have ignored the time dimension and therefore show a wide range of often disagreeing values in measured amplitudes and derived albedos and temperatures, especially in the infrared^{5,6,9–13,27}. This has so far been attributed to differing wavelengths, datasets and analysis methods. We show here that the planet itself is variable, and so care should be taken when simultaneously analysing measurements of

hot Jupiters at disparate times. This detection of variation in an exoplanet's atmosphere implies that the temporal dimension of planetary weather can now be explored outside the Solar System. Given the wide variety of exoplanet types observed so far, including many not yet seen in the Solar System, new models will be required to explain dynamical atmospheric changes. Future space missions such as CHEOPS, JWST, PLATO and ARIEL will be able to expand on this data and create a sample of temporally resolved exoplanet atmospheres.

Methods

Data source. We obtained publicly available data from the NASA Kepler satellite, which operated from 2 May 2009 to 11 May 2013⁸. The spacecraft reoriented itself approximately every 90 days, thereby separating data into 'quarters'. We downloaded the complete dataset for HAT-P-7 (Kepler-2, KIC 10666592), which comprised quarters 0–17 in both long (29.5 minute, data release 24 and 25) and short (58.9 second, data release 24) cadences. There is a known problem with the short cadence data smear correction in data release 24, and so, with short cadence data release 25 data not yet available, we focus our analysis on the long cadence data. Each data release and cadence gives consistent results. Data quarters were stitched together by normalizing each quarter by its own median flux to produce a single lightcurve.

Data detrending. We use several differing methods of detrending to reduce the chance that our chosen method affects our results. We start with the raw SAP_FLUX data. Long-term trends in this data are removed by fitting a third-degree polynomial function. First points within 0.05 in phase of the transit centre are masked using published ephemeris⁸. The data is then split into sections of width 0.3 d. A running window around each of these sections is used to fit a third-degree polynomial. We also trial a second-degree polynomial; although the resulting lightcurve contained increased systematic noise, there was no significant effect on the fit results. The polynomial is fit iteratively, with datapoints discrepant from the previous fit by more than 8σ masked from the succeeding fit. This is repeated for ten iterations. We trial several different lengths of window for fitting, comprising 10d, five planetary orbital periods (11.02d), and three planetary orbital periods (6.61d). Each produces consistent results. Shorter windows give decreased systematic error in the lightcurve, at increased risk of overfitting. As we found no significant difference in the resulting fit parameters between window sizes, we use three planetary periods for the final analysis. The window is not allowed to cross gaps of greater than 1d in the data. In cases where such gaps are found, the window is extended before or after the gap to meet the required window size.




Figure 3 | Model outputs giving positive and negative Θ values.

a,b, Example flux maps of the planetary dayside. **c,d**, Thermal emission (red), reflected flux (blue) and total (magenta) phase curve components, calculated using a published model²⁴, normalized to the total flux. Increased cloud coverage results in an increased reflected flux and lower temperatures, leading to $\Theta = -0.1$. (**a,c**). Reduced cloud coverage results in lower reflected flux and increased temperatures, leading to $\Theta = 0.1$. (**b,d**). Variability in the wind speed, which brings clouds from the morning terminator to the dayside, could cause the varying cloud coverage. The temperature change at the substellar point between each case is of the order of 100 K. Limb darkening is not shown here for clarity.

The resulting polynomial is then divided out from the section under consideration, and the process repeated for each section. At this stage, data within 0.7d of any gap larger than 1d is removed, as data ramps and other systematics are prevalent near gaps in the Kepler data. We then manually remove any regions that show clear instrumental systematics. The Kepler data before and after detrending are shown in Supplementary Fig. 8.

As an alternative detrending method, we employ the covariance basis vectors (CBVs) provided at the Michulski Archive for Space Telescopes (MAST). These are used to remove instrumental trends that are common to nearby stars on the CCD. We use the first five CBVs for detrending, enacted through the kepcotrend tool in the PyKE software package²⁸. Planetary transits and occultations are masked from the CBV fit. Following the CBV step, the polynomial fit is repeated as above in order to flatten the lightcurve, as using CBVs alone leaves significant long-term trends present. The results from the CBV-detrended data are consistent. We also tried fitting to the lightcurve produced by the CBVs alone, with no polynomial step. The resulting lightcurve contained significant instrumental noise that was visible by eye, although we found fitting results to be consistent with the polynomial lightcurves. This strengthens our confidence in the analysis and detrending applied.

Finally we compared the results between the long- and short-cadence data. These were again consistent, aside from a 27-day region centred on 543 (BJD-2454833), where a small but systematic offset was found. We conservatively removed data from this region before analysis.

Phase curve fitting. The full optical curve of HAT-P-7b is shown in Supplementary Fig. 9. We model the planetary phase curve using a previously published model⁶. This has four key components: the planetary transits/occultations, the tidal bulge caused by the planet's mass on the star, Doppler beaming from the star's reflex motion, and variation from the planetary brightness itself. Each is significant for this system. Also included is a cosine third-harmonic term empirically found to be present⁶. A full list of fit parameters and a detailed description of the model is given in the Supplementary Information. We hold all parameters associated with ellipsoidal variation, Doppler beaming and the planetary transit constant at the values of ref. 6, which are listed in Supplementary




Figure 4 | Contour map of Θ variations. $T_c = 2,000$ K, $\kappa = 40$, as a function of heat redistribution efficiency ϵ and Bond albedo A_B . Solid contours are positive, dashed are negative, red is zero. To produce the variations in Θ , the parameters must cross the red line repeatedly. The discontinuity seen at higher ϵ values marks the boundary of cloud formation on the planet's dayside. The greyed out regions mark areas excluded by the $\pm 59\%$ limit on amplitude variations.

Table 1. The planetary brightness amplitude, peak offset and secondary eclipse depth are allowed to vary, along with a constant offset in flux, leading to a four free-parameter fit.

To test the variation of the phase curve with time, the lightcurve is split into predefined segments of planetary orbits. We test segment sizes in the range of 5–20 orbits. Segments are not allowed to cross quarter boundaries. Segments with less than half the expected number of datapoints (calculated from the segment size and the data cadence) are not included. Each segment is then phase-folded using the planetary ephemeris and fit independently. We remove data during planetary transit to avoid this affecting the fit. We also trial a sliding window with the same size, shifting the window by one orbit and refitting. This allows higher time resolution but does not produce independent fits, as some of the same data is included in successive fits. As such the discrete segment fits should be used for measuring variation significance, but the sliding window can be used to explore the variation in greater detail.

We derive errors on these fits through a combination of three separate methods. These are the covariance matrix, the 'prayer bead' residual permutation technique (for example ref. 23), which is designed to incorporate systematic noise into the errors, and error resampling. These are explained in more detail in the below section. The method that gives the largest error is used for each fit and parameter. For final values we use a segment size of ten planetary orbits. This allows robust detection of the planetary brightness in each segment, at 9σ confidence on average. Decreasing the segment size to five planetary orbits shows no clear shorter timescale variation, while reducing the significance of each fit. Values for all fits are given in a supplementary file, and the individual fits and segments for discrete fits are shown in Supplementary Fig. 2.

It would be ideal to fit for the planetary mass-dependent processes simultaneously with the phase curve parameters. However, we note that the planetary mass must stay constant in time. As such, fitting for it would require a combined fit across the whole dataset, whereas variations were allowed independently in each measured segment (that is, a simultaneous fit of at least 250 parameters). This is computationally expensive to perform, and so we test the impact of the planetary mass by using the radial velocity (RV)-derived mass, as it is independent of the mass derived from the phase curve. We test fitting to the phase curve segments using different fixed planetary masses, at the -1σ , best-fitting and $+1\sigma$ values of the RV-derived mass (1.725, 1.781 and 1.862 M_J , respectively⁶). The resulting change in phase curve fit parameters was at the level of 5% of the parameter errors, and hence we ignore this effect. We conclude that the effect of an imperfectly known planetary mass, through Doppler beaming or ellipsoidal variation, is insignificant in this case.

We cut fits for which the phase curve amplitude was not detected at 3σ confidence, or where the detected amplitude was lower than 30 ppm. In these cases, the peak offset is not well-defined. This removed 26 fits from the original 592 produced using the sliding bin.

Fitting and error methods. We combine two methods of error derivation (prayer bead and error resampling) when calculating our values and errors. These errors are then compared to the error derived through the covariance matrix, and the largest error used in the final values.

The prayer bead technique (otherwise known as residual permutation) proceeds as follows. The data is initially fit, and initial best-fit parameters are obtained. The residuals of the data from this initial fit are stored. These residuals are then permuted; shifted through the data in time such that correlated noise in the residuals is maintained, but occurs at different phases. At each step, the shifted residuals are applied to the initial best-fit model, then the resulting data is refit. This can be repeated for as many datapoints as are available—beyond this number the fits repeat each other. We combine this method with another well-known method: error resampling. In this technique, multiple fits are repeated, but for every fit each datapoint is resampled from a Gaussian distribution formed from its own errors. We also randomize the initial guesses for each fit parameter, drawn from a uniform distribution within 20% of their initial best-fitting values. To combine both methods simultaneously, we shift the residuals by a random amount at each iteration.

With these variations applied each time, the fit is repeated for 1,000 iterations, for each segment of lightcurve that we consider. In each iteration, the best-fit parameters are found through least-squares minimization, enacted using the 'scipy.optimize.leastsq' routine in Python²⁹. The final values and errors are found from the distribution of resulting fit parameters. We extract the 15.87th, 50th and 84.14th percentiles from this distribution, to form the lower 1σ , best-fit value and upper 1σ errors, respectively. In this way, the errors contain 68.27% of the distribution. We then test the resulting errors against those derived from taking the square root of the diagonal of the covariance matrix. The covariance matrix errors are generally smaller than those derived through the resampling or prayer bead methods; if they are not we adopt the larger error for each data segment.

MCMC fitting. We also test the fitting using an MCMC routine. This is applied to two representative phase curve segments, one with significant positive θ and the other with significant negative θ . We implement an MCMC around the phase curve model using the open-source emcee software³⁰. We fit for the original four parameters, plus an additional factor f that is multiplied to the datapoint errors. This is to account for correlated noise not present in the given errors. We applied the MCMC to two representative phase curve chunks, at times 226 and 1,238 (BJD-2454833). These have significantly high and low θ , respectively. The MCMC was performed with 100 walkers, starting with randomized parameters near the same initial guess as the previous fitting methods. We iterated the chain for 5,000 steps, and found that all parameters had converged by 1,000 steps. The remaining 4,000 values (that is, 400,000 iterations considering all walkers) were used to test for parameter correlations and consistency with the other fitting. Triangle plots and derived values are shown in Supplementary Figs 4 and 5. Values for all parameters matched those of the previous analysis, although we find that the MCMC errors are underestimated compared with the prayer bead analysis. The larger errors are used for final values.

Quarter-to-quarter variations. It has previously been shown that the transit depths of HAT-P-7b vary on a seasonal cycle, as the target falls on different CCD modules each quarter³¹. Such systematic variations could bias our results. However, the variation in transit depth is approximately 1% around the average, which is much less than the variation we detect. As further evidence that our observed variability is not caused by seasonal variations, we note that the observed variability does not repeat on a yearly timescale (that is, when the star falls on the same CCD module), and shows no significant discontinuities between quarters, as seen in the transit depth analysis³¹. We also note that no significant power is detected in the periodogram at multiples of the 90 d quarter timespan (see below).

Orbit-to-orbit variations. We sometimes see apparently significant variations in planetary brightness amplitude, peak offset and secondary eclipse depth, on planetary orbital timescales. These should be treated with caution, as the phase curve is not significantly detected in individual orbits of the planet. We cannot exclude that they are due to either low-level stellar activity or small-scale changes in the data arising from the specific method used, although they may be the signal of short-timescale variations in the planetary atmosphere.

Periodicity analysis. A Lomb–Scargle^{32,33} periodogram analysis of θ (using discrete bins and hence independent fit parameters) returns no significant frequencies. We estimated significance by running periodograms on 10,000 simulated white noise datasets, with the same time sampling and scatter as the observed data. For each simulated dataset, the periodogram peak was extracted. The 1σ and 3σ percentiles of the distribution of maximum peaks were used as significance thresholds. We find that peaks in the real data are all weaker than 1σ under this significance test, and hence concluded that no periodicity is detected.

Exclusion of systematics/stellar variability. At the high level of precision at which we are interested, there is potential for systematic and instrumental noise in the Kepler data to affect our results. Although HAT-P-7 is an inactive star³⁴,

low-level stellar variability may also have an effect, and gravity darkening has been noted in the transit³⁵. First, all phases in transit are cut during the analysis, thereby removing any impact from gravity darkening. The third-order harmonic previously found⁴ has an amplitude of 1.93 ppm on average, and is hence negligible to the overall phase curve, and undetectable in the lightcurve segments we analyse. We trial a large array of detrending methods in an attempt to test for issues caused by our choice of methods. We find the results robust to any combination of parameters that we trial. Furthermore, non-planetary variability would have to occur in a way that mimicked a planetary phase curve, and would also need to occur on a much shorter timescale than the typical quarter-to-quarter systematics presented by Kepler. Variability with a timescale longer than three planetary periods is removed by the polynomial detrending, and variability with a timescale shorter than ten planetary periods is partially averaged out by the binning of ten successive phase curves. The stellar rotation period is >13.23 days³⁶, and hence would be removed by the polynomial detrending. Even before polynomial detrending, no significant starspot activity (which occurs on the stellar rotation period) has been detected by previous studies³⁴. Shorter-timescale stellar variability (such as 'flicker'³⁷) is averaged out by combining multiple planetary orbital periods in each fit. Finally, we note that our 'sliding window' fits present a smooth variation in fit parameters, which would not be the case if single events (such as cosmic rays, detector malfunctions, stellar flares or one-off instrumental effects) were causing the variation. Any such one off event would produce sharp variations in the fit parameters, both as the event entered the window and also as it left.

We further test for systematics by selecting random planetary orbits from throughout the dataset, then performing the fit on the combination of these. If the signals we see are coherent, as they should be if they are connected to the underlying atmospheric dynamics, then randomizing the orbits used should remove the variability. We find that the variability is markedly reduced when using ten planetary orbits, with the standard deviation of the extracted θ values dropping by $\sim 20\%$. We tested that this reduction in variability was reasonable for a real signal by considering a simulated dataset with similar properties to the true variation in θ . The simulated data consisted of two sinusoids with periods of 10 and 100 days (a simple recreation of a coherent variability signal and some correlated noise), and amplitude 0.035, overlaid on a white-noise signal with $\sigma = 0.1$. Considering randomized binned samples from the simulated dataset—as opposed to consecutive bins—led to a drop of $\sim 20\%$ in the standard deviation of θ , which matches the effect on real data. Using 20 planetary orbits rather than 10 removed the variation entirely.

Injection testing. As a further test against systematic noise, we identify two nearby stars with similar magnitude and temperature to HAT-P-7, observed on the same Kepler CCD module. These are KIC 10861893 and 11027406. For each of these targets, we follow the same procedure as HAT-P-7 for getting the lightcurve. We inject a constant phase curve using the HAT-P-7b parameters from ref. 6 into the lightcurves before the polynomial flattening stage. We then run the same detrending and fitting code and attempt to extract phase curve parameters. In both targets, no significant variation is found in θ (Supplementary Fig. 3), although some structure can be seen, probably as a result of underlying starspot activity. For KIC 10861893, however, marginally significant variation is seen in A_p , on the same level as for HAT-P-7b. Because the A_p variation in HAT-P-7b was inconclusive, it seems that this systematic noise source was included in our errors, as we would expect. This test confirms that the variations in A_p do not originate from the planet. The lack of variation in θ supports the physical origin of these variations in HAT-P-7b.

Alternative explanations. It is worth exploring whether other physical effects could give rise to the observed variation. A resonance between stellar rotation and planetary orbital period, combined with variable stellar activity, is excluded in the previous section. To give rise to a signal occurring on the planetary orbital period, other potential origins for the signal are necessarily related to the planet. They could include a variable dark spot on the star, consistently below the planet's position. This would require some form of magnetic connection between the planet and star, and hence a particularly strong planetary magnetic field. Such a spot would be expected to affect the transit depths of the planet, yet does not, thus making this explanation unlikely. Similarly, a variable bright spot on the opposite side of the star to the planet could lead to the observed signal, although it is unclear why such a spot would exist. Both explanations are intrinsically unlikely, as maintaining a coherent interaction between the two would be difficult due to the planet's misalignment with the stellar spin.

Atmospheric model. We implement a published semi-analytical model to allow us to explore the atmosphere of the planet²⁴. This model takes certain atmospheric parameters and calculates the expected phase curve. These parameters are: A_p : Bond albedo. Controls the proportion of light reflected by the planet. This reflected light comes from both cloudy and non-cloudy regions of the atmosphere, and hence the albedo in part controls the cloud coverage. f : Greenhouse factor. A free-scaling factor to represent possible greenhouse effects. As we do not have enough information to constrain this, we set it to unity, which corresponds to no greenhouse effect.

ϵ : Heat redistribution efficiency. Defined as the ratio between the radiative timescale and the advective timescale³⁸. $|\epsilon| \gg 1$ corresponds to efficient heat redistribution, such that temperature is smoothed out around the planet's longitude by strong winds. $|\epsilon| \ll 1$ represents the opposite, where heat is radiated faster than it is transported, and longitudinal differences in temperature will be large. We assume positive values for ϵ , which correspond to superrotating winds on the planet.

κ : Cloud reflectivity boosting factor. The factor by which cloudy regions are brighter than non-cloudy regions in reflected light.

T_C : Condensation temperature. A_p and ϵ , together with the known planetary and stellar parameters, allow for the calculation of a temperature map on the planet. Clouds form where the temperature falls below T_C . For positive ϵ , cloud formation occurs preferentially on the morning side of the planet, as the hottest point is shifted to the afternoon side.

With a temperature map, it is possible to calculate the thermal contribution to the phase curve by integrating the Planck function over the Kepler bandpass then integrating the local emission over the planetary disk. The reflected light contribution can be calculated by considering the cloudy and cloud-free regions. We refer the interested reader to the original publication²⁴ for full details.

Condensates. In the main text we propose corundum or perovskite as plausible condensates in HAT-P-7b's atmosphere. As such it is important to check if these condensates may have sufficient material to form the necessary clouds. Based on the work of ref. ³⁹, following the relative elemental mass calculations of ref. ⁴⁰, we find that if corundum and perovskite are the sole condensables for Al and Ti, respectively, then over one scale height the optical depth in the Kepler bandpass would be ~ 2 –4 for corundum and ~ 0.1 –1 for perovskite. If the cloud extends for greater than one scale height, these depths could increase. As such there is sufficient available material to form an optically thick cloud for corundum, but less so for perovskite, mainly because Al is more abundant than Ti. As there are a number of assumptions here, this does not exclude perovskite, but makes corundum the more likely of the two. We consider only particles of 0.1 μm and 1 μm in size, which are strongly scattering at short wavelengths for both species, meaning any cloud formed would probably be reflective.

Visualization. We visualize the changes on the planet by extracting temperature maps and cloud coverage from the atmospheric model. The model is not fully constrained, so we arbitrarily select a region of parameter space that produces the observed variations with modest parameter changes (0.38–0.17 in A_p , 12.9–5.3 in ϵ , with $\kappa=40$ and $T_C=2,000\text{ K}$). We extract the thermal and reflected flux for each contour shown in Fig. 4 between these values. For each observed θ , we interpolate between these contours to produce a representative flux map, cloud coverage and brightness. These are shown in Supplementary Video 1, which gives a visual representation of the changes observed. We stress that this video is for clarification and explanatory purpose only; the values used to produce it are not best-fits. Nevertheless, the proportion of thermal emission and reflected light seen at each epoch is the necessary combination to produce the directly observed peak offset.

In the video the bright band at the left (morning) side of the planet shows the area of cloud formation and associated reflected light. The more diffuse bright spot towards the right (afternoon) side of the planet represents the thermal emission seen in the Kepler bandpass. Note that although the thermal emission observed changes significantly, this only corresponds to relatively small changes in the underlying temperature. This is due to two effects. First, the shorter-wavelength edge of the thermal emission overlaps with the edge of the Kepler bandpass. As such, a slight increase in temperature markedly increases the amount of flux emitted into the Kepler bandpass. This would be expected to increase the secondary eclipse depth and phase curve amplitude, and indeed this is found to be the case by the model. However, the increase in depth and amplitude is not significant enough to be seen over our (relatively large) errors on these parameters, as described in the main text. Second, we do not show the effects of planetary limb darkening in the video, which would decrease the flux from the planetary limbs. Displaying the planet without limb darkening amplifies the cloud and thermal variations to increase clarity.

Data availability. Individual fit values are provided in the Supplementary Information as a CSV file. Any other data that support the plots within this paper and other findings of this study are available from the corresponding author upon reasonable request.

Received 16 June 2016; accepted 3 November 2016;
published 12 December 2016

References

- Knutson, H. A., Charbonneau, D., Allen, L. E., Burrows, A. & Megeath, S. T. The 3.6–8.0 μm broadband emission spectrum of HD 209458b: Evidence for an atmospheric temperature inversion. *Astrophys. J.* **673**, 526–531 (2008).
- Snellen, I. A. G., de Mooij, E. J. W. & Albrecht, S. The changing phases of extrasolar planet CoRoT-1b. *Nature* **459**, 543–545 (2009).
- Artigau, É., Bouchard, S., Doyon, R. & Lafrenière, D. Photometric variability of the T2.5 brown dwarf SIMP J013656.5+093347: Evidence for evolving weather patterns. *Astrophys. J.* **701**, 1534–1539 (2009).
- Radigan, J. *et al.* Large-amplitude variations of an L/T transition brown dwarf: Multi-wavelength observations of patchy, high-contrast cloud features. *Astrophys. J.* **750**, 105 (2012).
- Esteves, L. J., de Mooij, E. J. W. & Jayawardhana, R. Optical phase curves of Kepler exoplanets. *Astrophys. J.* **772**, 51 (2013).
- Esteves, L. J., de Mooij, E. J. W. & Jayawardhana, R. Changing phases of alien worlds: Probing atmospheres of Kepler planets with high-precision photometry. *Astrophys. J.* **804**, 150 (2015).
- Pál, A. *et al.* HAT-P-7b: An extremely hot massive planet transiting a bright star in the Kepler field. *Astrophys. J.* **680**, 1450–1456 (2008).
- Borucki, W. J. *et al.* Kepler planet-detection mission: Introduction and first results. *Science* **327**, 977–980 (2010).
- Wong, I. *et al.* 3.6 and 4.5 μm Spitzer phase curves of the highly-irradiated hot Jupiters WASP-19b and HAT-P-7b. *Astrophys. J.* **823**, 122 (2016).
- Christiansen, J. L. *et al.* Studying the atmosphere of the exoplanet HAT-P-7b via secondary eclipse measurements with Epoxi, Spitzer, and Kepler. *Astrophys. J.* **710**, 97–104 (2010).
- Borucki, W. J. *et al.* Kepler's optical phase curve of the exoplanet HAT-P-7b. *Science* **325**, 709–709 (2009).
- Welsh, W. F. *et al.* The discovery of ellipsoidal variations in the Kepler light curve of HAT-P-7. *Astrophys. J. Lett.* **713**, L145–L149 (2010).
- von Paris, P., Gratier, P., Bordé, P. & Selsis, F. Inferring heat recirculation and albedo for exoplanetary atmospheres: Comparing optical phase curves and secondary eclipse data. *Astron. Astrophys.* **587**, A149 (2016).
- Faigler, S. & Mazeh, T. BEER analysis of Kepler and CoRoT light curves. II. Evidence for superrotation in the phase curves of three Kepler hot Jupiters. *Astrophys. J.* **800**, 73 (2015).
- Showman, A. P. & Polvani, L. M. Equatorial superrotation on tidally locked exoplanets. *Astrophys. J.* **738**, 71 (2011).
- Fromang, S., Leconte, J. & Heng, K. Shear-driven instabilities and shocks in the atmospheres of hot Jupiters. *Astron. Astrophys.* **591**, A144 (2016).
- Webber, M. W. *et al.* Effect of longitude-dependent cloud coverage on exoplanet visible wavelength reflected-light phase curves. *Astrophys. J.* **804**, 94 (2015).
- Rauscher, E., Menou, K., Cho, J. Y.-K., Seager, S. & Hansen, B. M. S. On signatures of atmospheric features in thermal phase curves of hot Jupiters. *Astrophys. J.* **681**, 1646–1652 (2008).
- Showman, A. P. *et al.* Atmospheric circulation of hot Jupiters: Coupled radiative-dynamical general circulation model simulations of HD 189733b and HD 209458b. *Astrophys. J.* **699**, 564–584 (2009).
- Agol, E. *et al.* The climate of HD 189733b from fourteen transits and eclipses measured by Spitzer. *Astrophys. J.* **721**, 1861–1877 (2010).
- Knutson, H. A. *et al.* 3.6 and 4.5 μm phase curves and evidence for non-equilibrium chemistry in the atmosphere of extrasolar planet HD 189733b. *Astrophys. J.* **754**, 22 (2012).
- Demory, B.-O., Gillon, M., Madhusudhan, N. & Queloz, D. Variability in the super-Earth 55 Cnc e. *Mon. Not. R. Astron. Soc.* **455**, 2018–2027 (2015).
- Gillon, M. *et al.* Accurate Spitzer infrared radius measurement for the hot Neptune GJ 436b. *Astron. Astrophys.* **471**, L51–L54 (2007).
- Hu, R., Demory, B.-O., Seager, S., Lewis, N. & Showman, A. P. A semi-analytical model of visible-wavelength phase curves of exoplanets and applications to Kepler-7 B and Kepler-10 B. *Astrophys. J.* **802**, 51 (2015).
- Parmentier, V., Fortney, J. J., Showman, A. P., Morley, C. V. & Marley, M. S. Transitions in the cloud composition of hot Jupiters. Preprint at <https://arxiv.org/abs/1602.03088> (2016).
- Kataraia, T. *et al.* The atmospheric circulation of a nine-hot-Jupiter sample: Probing circulation and chemistry over a wide phase space. *Astrophys. J.* **821**, 9 (2016).
- Schwartz, J. C. & Cowan, N. B. Balancing the energy budget of short-period giant planets: Evidence for reflective clouds and optical absorbers. *Mon. Not. R. Astron. Soc.* **449**, 4192–4203 (2015).
- Still, M. & Barclay, T. PyKE: Reduction and analysis of Kepler simple aperture photometry data. *Astrophys. Source Code Lib.* 1208.004 (2012).
- Oliphant, T. E. Python for scientific computing. *Comput. Sci. Eng.* **9**, 10–20 (2007).
- Foreman-Mackey, D., Hogg, D. W., Lang, D. & Goodman, J. emcee: The MCMC hammer. *Publ. Astron. Soc. Pacific* **125**, 306–312 (2013).
- Van Eylen, V., Lindholm Nielsen, M., Hinrup, B., Tingley, B. & Kjeldsen, H. Investigation of systematic effects in Kepler data: Seasonal variations in the light curve of HAT-P-7b. *Astrophys. J. Lett.* **774**, L19 (2013).
- Lomb, N. R. Least-squares frequency analysis of unequally spaced data. *Astrophys. Space Sci.* **39**, 447–462 (1976).
- Scargle, J. D. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *Astrophys. J.* **263**, 835 (1982).

34. McQuillan, A., Mazeh, T. & Aigrain, S. Stellar rotation periods of the Kepler objects of interest: A dearth of close-in planets around fast rotators. *Astrophys. J. Lett.* **775**, L11 (2013).
35. Morris, B. M., Mandell, A. M. & Deming, D. Kepler's optical secondary eclipse of HAT-P-7b and probable detection of planet-induced stellar gravity darkening. *Astrophys. J. Lett.* **764**, L22 (2013).
36. Lund, M. N. *et al.* Asteroseismic inference on the spin-orbit misalignment and stellar parameters of HAT-P-7. *Astron. Astrophys.* **570**, A54 (2014).
37. Cranmer, S. R., Bastien, F. A., Stassun, K. G. & Saar, S. H. Stellar granulation as the source of high-frequency flicker in Kepler light curves. *Astrophys. J.* **781**, 124 (2014).
38. Cowan, N. B. & Agol, E. A model for thermal phase variations of circular and eccentric exoplanets. *Astrophys. J.* **726**, 82 (2010).
39. Wakeford, H. R. *et al.* High temperature condensate clouds in super-hot Jupiter atmospheres. *Mon. Not. R. Astron. Soc.* Preprint available at <https://arxiv.org/abs/1610.03325> (2016).
40. Lodders, K. in *Formation and Evolution of Exoplanets* Ch. 8 (Wiley, 2010).

Acknowledgements

D.J.A. acknowledges funding from the European Union Seventh Framework programme (FP7/2007–2013) under grant agreement No. 313014 (ETA-EARTH). E.d.M. acknowledges support from the Michael West Fellowship. This paper includes data collected by the Kepler mission. Funding for the Kepler mission is provided by the NASA Science Mission directorate. All of the data presented in this paper were obtained from

the Mikulski Archive for Space Telescopes (MAST). STScI is operated by the Association of Universities for Research in Astronomy, under NASA contract NAS5-26555. Support for MAST for non-HST data is provided by the NASA Office of Space Science via grant NNX09AF08G and by other grants and contracts.

Author contributions

D.J.A. obtained and detrended the data, developed and fit the phase curve models, implemented the atmospheric model, produced the figures and wrote the manuscript. E.d.M. developed the discussion, contributed to the tests performed to check the results, and tested the results with his own models. H.P.O. contributed to the phase curve model, and produced visual interpretations of the results. J.Ba. developed the discussion of the atmospheric processes behind the peak offset variations. J.Bl. provided the initial development of the phase curve model. N.F.S. contributed to development of the figures. All authors commented on the manuscript.

Additional information

Supplementary information is available for this paper.

Reprints and permissions information is available at www.nature.com/reprints.

Correspondence and requests for materials should be addressed to D.J.A.

How to cite this article: Armstrong, D. J. *et al.* Variability in the atmosphere of the hot giant planet HAT-P-7 b. *Nat. Astron.* **1**, 0004 (2016).

Competing interests

The authors declare no competing financial interests.