# LETTER

# Detection of titanium oxide in the atmosphere of a hot Jupiter

Elyar Sedaghati[1,2,3], Henri M. J. Boffin[1], Ryan J. MacDonald[4], Siddharth Gandhi[4], Nikku Madhusudhan[4], Neale P. Gibson[5], Mahmoudreza Oshagh[6,7], Antonio Claret[8] & Heike Rauer[2,3]

As an exoplanet transits its host star, some of the light from the star is absorbed by the atoms and molecules in the planet's atmosphere, causing the planet to seem bigger; plotting the planet's observed size as a function of the wavelength of the light produces a transmission spectrum[1]. Measuring the tiny variations in the transmission spectrum, together with atmospheric modelling, then gives clues to the properties of the exoplanet's atmosphere. Chemical species composed of light elements—such as hydrogen, oxygen, carbon, sodium and potassium—have in this way been detected in the atmospheres of several hot giant exoplanets[2–5], but molecules composed of heavier elements have thus far proved elusive. Nonetheless, it has been predicted that metal oxides such as titanium oxide (TiO) and vanadium oxide occur in the observable regions of the very hottest exoplanetary atmospheres, causing thermal inversions on the dayside[6,7]. Here we report the detection of TiO in the atmosphere of the hot-Jupiter planet WASP-19b. Our combined spectrum, with its wide spectral coverage, reveals the presence of TiO (to a confidence level of 7.7$\sigma$), a strongly scattering haze (7.4$\sigma$) and sodium (3.4$\sigma$), and confirms the presence of water (7.9$\sigma$) in the atmosphere[5,8].

Hot Jupiters are gas-giant exoplanets with sizes like that of Jupiter but much shorter orbital periods. WASP-19b is the shortest-period hot Jupiter to be discovered so far[9], and has an excessively bloated radius, owing to the extreme radiation that it receives from its host star; as a result of this radiation, the planet's effective temperature is more than 2,000 K (obtained via secondary-eclipse measurements[10]). It is thought that high atmospheric temperatures imply the presence of metal oxides such as TiO, but despite extensive searches[11,12] a definitive detection of metal oxides in exoplanetary atmospheres has proved elusive.

We observed three transits of WASP-19b with the 8.2-metre Unit Telescope 1 (UT1) of the European Southern Observatory's Very Large Telescope (VLT), using the low-resolution FORS2 spectrograph. By using three of FORS2's grisms—600B (blue), 600RI (green) and 600z (red), thereby covering the entire visible-wavelength domain (0.43–1.04 μm)—together with the multi-object spectroscopy configuration, we were able to obtain relatively high-resolution, precise, broadband transmission spectra. Such results were made possible through optimized observing strategies[13] and careful design of the observing mask used for the multi-object observations: this has slits about 30″ wide, which minimized differential losses owing to variations in telescope guiding and seeing conditions. The observations presented here were made between 11 November 2014 and 29 February 2016.

For each set of observations, we obtained a series of spectra for the main target (WASP-19), as well as for several comparison stars. After standard data-reduction steps, we integrated those spectra for the largest common wavelength domain and 10-nm bins, to produce the 'white' and 'spectrophotometric' light curves, respectively. To correct for the imprint of telluric variations on the detected light, we divided the light curve of the target star by that of the one comparison star providing the best result. This corrected light curve was chosen through a detailed, statistical analysis of the differential light curves of all comparison stars, as well as all their possible combinations[14,15].

We initially fitted the broadband light curves by using an analytical solution for a transit[16], as well as a Gaussian process model for the estimation of correlated noise in the data[17]. This approach provides a model-independent stochastic method for including the systematic component of noise into our model. We derived the maximum posterior probability distributions for the fitted parameters by solving the Bayesian relation and running multiple, very long Monte Carlo Markov Chain (MCMC) simulations of our multivariate model[18]. The Gaussian process systematic model is trained by using carefully selected inputs, which are those physical variants that play a role in introducing correlated noise to the time-series data[19]. This method gives us a robust estimate of the uncertainties in determining the planetary physical and orbital parameters, by fully accounting for the contribution of systematic noise[17]. Once the optimal broadband solutions are found for each data set, we use the wavelength-independent parameter values as strict, informative priors in our Bayesian analysis of the spectral light curves, and solve for variations of the planetary radius as a function of wavelength bin. Broadband light curves for the three sets of observations, exemplar spectral light curves and the fitted systematic models are shown in Fig. 1.

Figure 2 shows the wavelength-dependent radius variations—the transmission spectrum—where the combination of results from the three observing campaigns constructs a broadband spectrum with high precision. It is this precision that facilitates the detection of features originating from the exoplanetary atmosphere. Such signals are imprinted upon the observed spectrum owing to the discrete absorption of the host star's light by gases as the light traverses the day–night boundary of the exoplanetary atmosphere[1]. Interpreting this transmission spectrum requires fitting theoretical atmospheric models to the data, a process that involves exploring a wide range of the atmospheric parameter space, described shortly. However, before performing this task, an important factor to consider is the role that the host star's activity plays in introducing spurious signals into our final spectrum[20]. Especially important is the presence of stellar active regions, be it spots (dark) or faculae (bright). If the transiting planet happens to traverse one of these regions from the point of a view of an observer on Earth, then the light curve shows an anomaly in the form of a small bump or a dip. For instance, we see an evidence for such a spot-crossing event (a bump) in the blue light curve, where the anomaly cannot be attributed to changes in observing conditions. To circumvent this, we include an additional complexity in our analytical model for a circular-shaped spot on the stellar surface with a temperature below that of the surrounding stellar photosphere. This ensures that our

[1]European Southern Observatory, Alonso de Córdova 3107, Santiago, Chile. [2]Deutsches Zentrum für Luft- und Raumfahrt, Rutherfordstrasse 2, 12489 Berlin, Germany. [3]Zentrum für Astronomie und Astrophysik, TU Berlin, Hardenbergstrasse 36, 10623 Berlin, Germany. [4]Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK. [5]Astrophysics Research Centre, School of Mathematics and Physics, Queens University Belfast, Belfast BT7 1NN, UK. [6]Institut für Astrophysik, Georg-August Universität Göttingen, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany. [7]Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto, CAUP, Rua das Estrelas, 4150-762 Porto, Portugal. [8]Instituto de Astrofísica de Andalucía, CSIC, Apartado 3004, 18080 Granada, Spain.
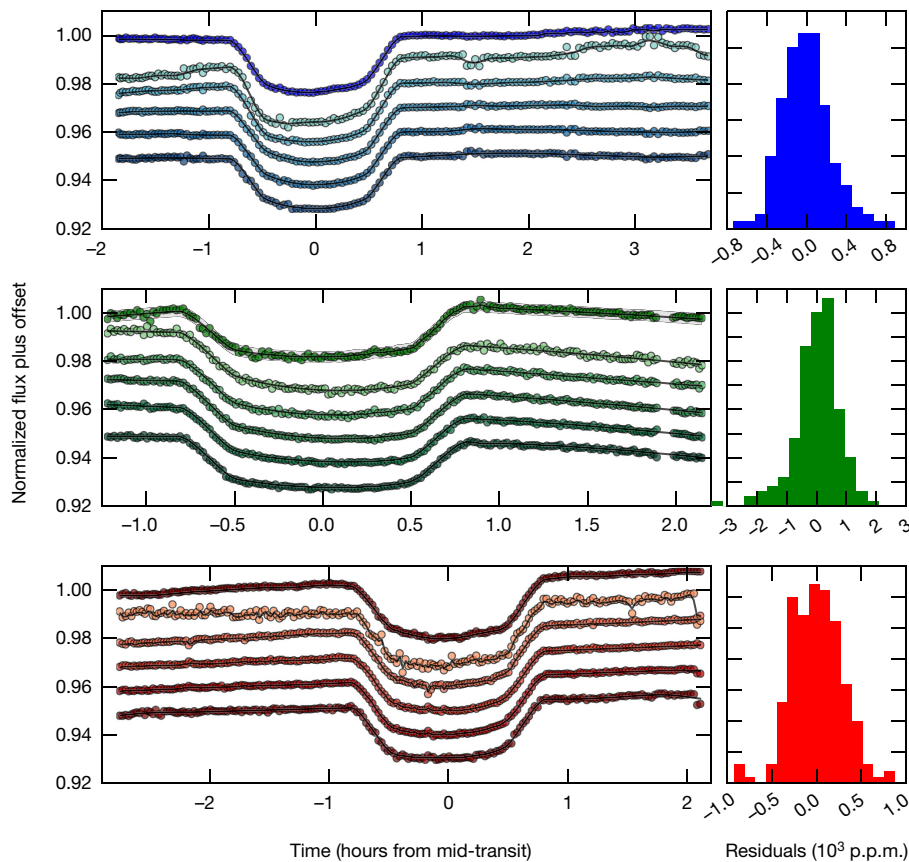
**Figure 1 | Light curves and models.** Broadband and spectrophotometric light curves of WASP-19b from our three transit observation campaigns. The colour of each row corresponds to the data set that it represents, obtained using the 600B (blue), 600RI (green) and 600z (red) grisms. In the left panels, for each set, the top plot is the broadband transit light curve, used for determining wavelength-independent parameters, and the others are demonstrative spectral light curves produced using 50-nm integration bins. The grey-shaded regions in the broadband plots highlight the $3\sigma$ confidence level of the Gaussian process model. The panels on the right show the distribution of the residuals of the analytical models fit to the broadband light curves.

results are not biased by the spot-crossing event. We also account for the possible presence of active regions on parts of the stellar disk that are not covered by the transit chord, better known as unocculted spots or faculae. Their presence introduces a bias in determining the out-of-transit baseline flux, which in turn leads to errors in the inferred value of the transit depth and hence to errors in the estimated planetary radius. This is particularly important for transmission spectroscopy, as the effect is chromatic[21]. Our analysis reveals that the enhanced planetary radius seen towards the ultraviolet end of our spectrum can be attributed only partially to the possibility of unocculted spots, even when assuming a large spot-filling factor and high temperature differences.

We infer the atmospheric properties of WASP-19b at the day–night boundary, probed by our transmission spectrum, by using the atmospheric retrieval algorithm POSEIDON[22]. This algorithm explores many millions of transmission spectra—spanning a wide range of chemical compositions, temperatures, and cloud/haze properties—to identify the range of atmospheres that is consistent with the observations. In addition to standard absorbers expected in hot Jupiter atmospheres[4] ($H_2$, He, Na, K and $H_2O$), we consider a wide range of metal oxides and hydrides—TiO, VO, AlO, TiH, NaH, MgH, CrH, CaH, ScH and FeH. Our models span the continuum from clear to cloudy atmospheres, both with and without scattering hazes, and include two-dimensional models with inhomogeneous cloud coverage.

We report detections of $H_2O$ (confidence limit $7.9\sigma$), TiO ($7.7\sigma$) and Na ($3.4\sigma$), as well as a strongly scattering haze ($7.4\sigma$) that envelopes the planet. Constraints on the volume mixing ratios and haze properties are presented in Fig. 3. The derived abundances of $H_2O$ (18–1,300 p.p.m.) and Na (0.028–140 p.p.m.) span a wide range, at one end being consistent with expectations of an atmosphere of solar composition, and at the other being substantially sub-solar[23]; the TiO abundances (the $1\sigma$ range is 0.015–1.1 p.p.b.) are sub-solar to a
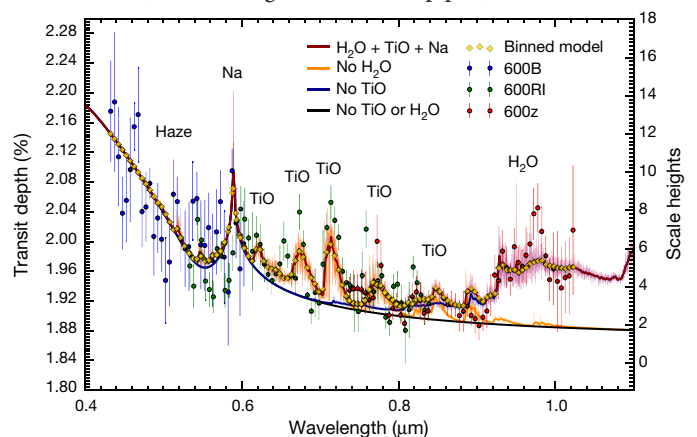


**Figure 2 | WASP-19b transmission spectrum.** Blue, green and red data points correspond to observations made using the 600B, 600RI and 600z grisms, respectively. The associated $1\sigma$ error bars were derived from posterior probability distributions of the planetary radius parameter in a joint analysis of the MCMC chains, where the mean of each distribution is plotted as the solution. The overall best-fitting spectrum—which includes opacity resulting from the presence of $H_2O$, TiO, Na, and a global haze—is shown as a red curve and yellow point at a representative resolution of about 3,000. Other curves (orange, blue and black) represent models with specific species removed. We applied a corrective wavelength shift between the model and the data of 73.6 Å.
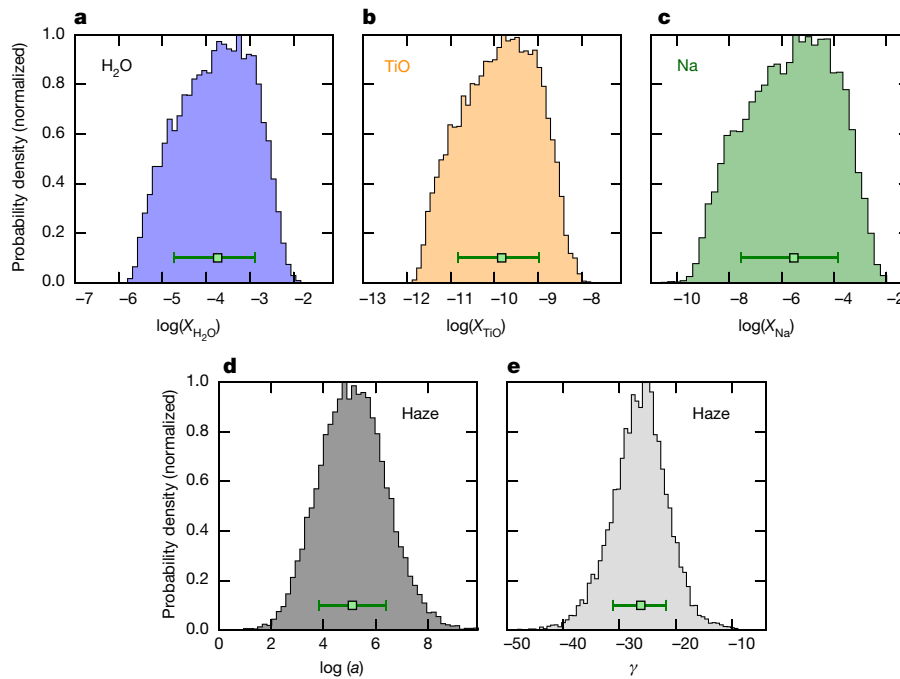
**Figure 3 | Constituents detected in WASP-19b's atmosphere.**
**a–c**, Histograms showing the marginalized posterior probability densities of the $H_2O$ (**a**), TiO (**b**) and Na (**c**) volume mixing ratios at WASP-19b's day–night terminator. **d, e**, Histograms showing the marginalized posterior probability densities of the haze Rayleigh enhancement

factor, $a$, and scattering slope[25], $\Upsilon$, as defined by the power law $\sigma(\lambda) = a\sigma_0(\lambda/\lambda_0)^{\Upsilon}$, where $\lambda_0$ is a reference wavelength (350 nm) and $\sigma_0$ is the $H_2$–Rayleigh scattering cross-section at the reference wavelength ($5.31 \times 10^{-31}$ m$^2$).

confidence limit of more than $5\sigma$ (the upper limit of which is 44 p.p.b.). Despite suggestions of substructure in the low-wavelength region of the spectrum (at around 0.5 μm or slightly higher), which could be explained by the presence of metal hydrides, our Bayesian model comparison leads to the conclusion that a haze optimally explains these observations, without the need to invoke other chemical species. The haze we detect is about 100,000 times stronger than Rayleigh scattering from $H_2$ alone, follows a power-law with an exponent of $-26^{+4}_{-5}$, and is consistent with 100% coverage across the terminator. We do not detect an opaque cloud deck. Finally, the planet's atmospheric temperature in the line-of-sight at 1 mbar pressure is constrained to $2,350^{+168}_{-314}$ K.

Our detection of TiO in WASP-19b's atmosphere is consistent with expectations given the atmosphere's high temperature (more than 2,000 K)[7]. Moreover, our spectroscopic detection of a refractory metal oxide in an exoplanetary atmosphere demonstrates the importance of visible-wavelength molecular opacity in transmission spectra. The strong visible opacity that results from the presence of TiO could have substantial effects on the temperature structure and circulation of the planet's atmosphere. If present in large enough quantities on the dayside, the TiO might cause a thermal inversion in the dayside atmosphere[6,7], which could in the future be observable in high-precision thermal emission spectra obtained with the Hubble Space Telescope and James Webb Space Telescope. The increased opacity resulting from the presence of TiO could also lead to strong day–night temperature gradients in the atmosphere[24]. On a more general note, our results demonstrate the usefulness of ground-based transmission spectroscopy for pursuing a detailed molecular analysis of exoplanetary atmospheres at optical wavelengths.

1. Seager, S. & Sasselov, D. Theoretical transmission spectra during extrasolar giant planet transits. *Astrophys. J.* **537**, 916–921 (2000).
2. Deming, D. *et al.* Infrared transmission spectroscopy of the exoplanets HD 209458b and XO-1b using the Wide Field Camera-3 on the Hubble Space Telescope. *Astrophys. J.* **774**, 95 (2013).
3. Kreidberg, L. *et al.* A precise water abundance measurement for the hot Jupiter WASP-43b. *Astrophys. J.* **793**, L27 (2014).
4. Madhusudhan, N., Agúndez, M., Moses, J. I. & Hu, Y. Exoplanetary atmospheres: chemistry, formation conditions, and habitability. *Space Sci. Rev.* **205**, 285–348 (2016).
5. Sing, D. K. *et al.* A continuum from clear to cloudy hot-jupiter exoplanets without primordial water depletion. *Nature* **529**, 59–62 (2016).
6. Hubeny, I., Burrows, A. & Sudarsky, D. A possible bifurcation in atmospheres of strongly irradiated stars and planets. *Astrophys. J.* **594**, 1011–1018 (2003).
7. Fortney, J. J., Lodders, K., Marley, M. S. & Freedman, R. S. A unified theory for the atmospheres of the hot and very hot Jupiters: two classes of irradiated atmospheres. *Astrophys. J.* **678**, 1419–1435 (2008).
8. Huitson, C. M. *et al.* An HST optical-to-near-IR transmission spectrum of the hot Jupiter WASP-19b: detection of atmospheric water and likely absence of TiO. *Mon. Not. R. Astron. Soc.* **434**, 3252–3274 (2013).
9. Hebb, L. *et al.* WASP-19b: the shortest period transiting exoplanet yet discovered. *Astrophys. J.* **708**, 224–231 (2010).
10. 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).
11. Haynes, K., Mandell, A. M., Madhusudhan, N., Deming, D. & Knutson, H. Spectroscopic evidence for a temperature inversion in the dayside atmosphere of hot Jupiter WASP-33b. *Astrophys. J.* **806**, 146 (2015).
12. Evans, T. M. *et al.* Detection of $H_2O$ and evidence for TiO/VO in an ultra-hot exoplanet atmosphere. *Astrophys. J.* **822**, L4 (2016).
13. Boffin, H. M. J. *et al.* Regaining the FORS: making optical ground-based transmission spectroscopy of exoplanets with VLT+ FORS2 possible again. *Proc. SPIE* **99082**, http://dx.doi.org/10.1117/12.2232094 (2016).
14. Sedaghati, E. *et al.* Potassium detection in the clear atmosphere of a hot-Jupiter-FORS2 transmission spectroscopy of WASP-17b. *Astron. Astrophys.* **596**, A47 (2016).
15. Sedaghati, E. *et al.* Probing the atmosphere of a sub-Jovian planet orbiting a cool dwarf. *Mon. Not. R. Astron. Soc.* **468**, 3123–3134 (2017).
16. Mandel, K. & Agol, E. Analytic light curves for planetary transit searches. *Astrophys. J.* **580**, L171 (2002).
17. Gibson, N. *et al.* A Gaussian process framework for modelling instrumental systematics: application to transmission spectroscopy. *Mon. Not. R. Astron. Soc.* **419**, 2683–2694 (2012).
18. Collier Cameron, A. *et al.* Efficient identification of exoplanetary transit candidates from Super-WASP light curves. *Mon. Not. R. Astron. Soc.* **380**, 1230–1244 (2007).
19. Gibson, N. *et al.* Probing the haze in the atmosphere of HD 189733b with Hubble Space Telescope/WFC3 transmission spectroscopy. *Mon. Not. R. Astron. Soc.* **422**, 753–760 (2012).

20. Oshagh, M. *et al.* Effect of stellar spots on high-precision transit light-curve. *Astron. Astrophys.* **556,** A19 (2013).
21. Oshagh, M. *et al.* Impact of occultations of stellar active regions on transmission spectra: can occultation of a plage mimic the signature of a blue sky? *Astron. Astrophys.* **568,** A99 (2014).
22. MacDonald, R. J. & Madhusudhan, N. HD 209458b in new light: evidence of nitrogen chemistry, patchy clouds and sub-solar water. *Mon. Not. R. Astron. Soc.* **469,** 1979–1996 (2017).
23. Madhusudhan, N. C/O ratio as a dimension for characterizing exoplanetary atmospheres. *Astrophys. J.* **758,** 36 (2012).
24. 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).
25. Lecavelier des Etangs, A., Pont, F., Vidal-Madjar, A. & Sing, D. Rayleigh scattering in the transit spectrum of HD 189733b. *Astron. Astrophys.* **481,** L83–L86 (2008).

## METHODS

We here present, in greater detail, the observational, data-reduction and retrieval methods used in our analysis of WASP-19b's transmission spectrum.

**Observations.** We have observed multiple transits of WASP-19b (ref. 9), a hot Jupiter that is in an extremely tight orbit around its host star, WASP-19: it revolves around this magnitude-12.3 G8V star on a 0.789-day orbit, bringing it very close to its host's Roche limit. All observations were performed with ESO's FOcal Reducer and low dispersion Spectrograph (FORS2)[26], mounted at the Cassegrain focus of UT1 of the VLT. The data were obtained over a two-year period. Details of the three observing campaigns are given in Extended Data Table 1. The instrument has a $6.8' \times 6.8'$ field of view and comprises a mosaic of two $2k \times 4k$ pixel detectors, which are either red optimized (MIT; green and red data) or blue optimized (E2V; blue data). The observations were prepared and mostly performed by members of our team.

For all sets of observations, the instrument's atmospheric dispersion corrector (the LADC, which was upgraded in 2014; ref. 13) was left in 'park' position throughout the entire observing sequence, with its two prisms kept at a minimal separation of 30 mm. For the two most recent sets of observations, we chose to use the non-standard 200 kHz readout. This strategy allowed us to spend more time on the target for a given cadence owing to the reduced readout time, subsequently increasing the signal-to-noise ratio of the spectra. For observations performed with the Mask eXchange Unit (MXU) mode, we created custom-designed masks with typically $30''$-wide slitlets placed on top of the target star, and also on top of several reference stars and the sky in order to monitor variations in the background contamination level. As part of a routine calibration sequence, we also took bias, flat-field and arc images. The arc frames were taken with the same slit configuration as the scientific sequences, but with the width of the slitlets set to $1''$ for increased resolution. Best effort was made in placing all of the observed stars in the middle of the detector set as much as possible, in order to ensure maximal wavelength overlap between the targets. An example of a mask created for observations in MXU mode is shown in Extended Data Fig. 1a. This is in keeping with the optimal observation strategy when using FORS2 for multi-object spectroscopy[13]. The observational conditions for all sets of observations were generally clear but not photometric. Some sporadic cirrus clouds affected the raw light curves in the 'blue' campaign towards the end of the observing sequence; the influence of these clouds is mostly corrected for by differential photometry, although some residual effects remain. These effects are however well modelled by the Gaussian process model. Finally, the results of the 'green' campaign are from the reanalysis of a data set that has been studied previously[27].

**Data reduction.** For the purpose of data reduction, we wrote a specialized Pyraf pipeline[14,15], which included: first, overscan and bias-shape subtraction; second, spectroscopic flat-fielding; third, spectral extraction; and fourth, wavelength calibration. For the extraction of one-dimensional spectra from the two-dimensional science frames, we implemented an optimal extraction algorithm[28]. The size of the extraction box used for each star was varied, and chosen by analysing dispersion in the subsequent light curves (see Extended Data Fig. 1b for an example). Specifically, the width of the extraction box was selected once the gradient of the light curve dispersion relation to the extraction width was below $10^{-4}$ pix$^{-1}$, for the frame with the largest point spread function in the spatial direction. This value was then used for all of the frames in the corresponding data set. In the example shown, 30 binned pixels were used for both the target and the comparison star. However, this extraction size varied from star to star, as well as from night to night, owing to variations in brightness and seeing conditions. Examples of such extracted spectra of the target and the comparison stars, for each data set, are shown in the top row of Extended Data Fig. 2. To obtain the dispersion solutions we used a Chebyshev polynomial function fit of order 4, using the one-dimensional extracted arc frames, which provided results with a root mean square of 0.06 Å or better after deleting a few outlying features.

Once extracted, the series of the stellar spectra were integrated, initially within the largest common wavelength domain of all the targets. Their variations as a function of time result in the raw light curves. We then produced differential transit light curves by dividing the target light curve (that of WASP-19) by the light curves of all of the observed comparison stars, as well as all their possible combinations. We chose the final comparison star as the one that produced the cleanest transit light curve (that is, the light curve with the lowest residuals relative to a transit model). Once chosen, we reintegrated the target and the chosen comparison star for their largest common wavelength overlap. Examples of these final raw and differential light curves are given in the middle and bottom rows of Extended Data Fig. 2, respectively, for the observing campaigns.

The spectrophotometric or narrow-band transit light curves were produced in an identical manner, but using integration bins of 10 nm. We chose this value through a statistical analysis of all three data sets, where the final decision was dictated by the least-precise set[14]. All of these light curves, used to produce the

transmission spectrum, are shown in the left panels of Extended Data Figs 3–5 for the blue, green and red data sets. We further corrected these light curves by division through the residuals of the analytical model to the broadband light curve. This is known as common-mode systematic correction[29–31], and it removes the wavelength-independent component of the systematic noise from the spectral light curves. The corrected light curves are shown in the right panels of Extended Data Figs 3–5, together with the Gaussian process models (described below).

**Data analysis.** We modelled all transit light curves using a Gaussian process model, which provides a model-independent stochastic approach to including systematic signals[32]. Their application to modelling exoplanet transit data and transmission spectroscopy has been demonstrated for the analysis of Hubble Space Telescope (HST)/Wide-Field Camera 3 (WFC3) data for probing the haze in the atmosphere of HD 189733b (ref. 19). We find maximum posterior solutions, $\mathcal{P}$, by optimizing the Bayesian relation:

$$\log \mathcal{P}(\theta, \phi \,|\, \boldsymbol{f}, X) = \log \mathcal{L}(\boldsymbol{r}|X, \theta, \phi) + \log P(\theta, \phi) \qquad (1)$$

that is true up to a constant term, where $\theta$ and $\phi$ are collections of noise and transit model parameters respectively; $\boldsymbol{f}$ is the vector of flux measurements; $X$ is an $N \times K$ matrix ($N$ being the number of observations and $K$ the number of inputs needed to train the model); and $\boldsymbol{r}$ is the residual vector for the model, elements of which are calculated from flux measurements relative to an analytical transit model[16]. The two terms on the right-hand side are the likelihood function, $\mathcal{L}$, and prior probability, $P$, respectively. In order to introduce correlated noise to our covariance matrix, we write the likelihood function, $\mathcal{L}$, as a matrix equation:

$$\log \mathcal{L}(\boldsymbol{r}|X, \theta, \phi) = -\frac{1}{2}(\boldsymbol{r}^T \Sigma \boldsymbol{r} + \log|\Sigma| + N \log 2\pi)$$

where the superscript $T$ is the transpose and $\Sigma$ represents the covariance matrix. Ideally one would like to know the full covariance matrix, and we simplify this process by 'modelling' the elements of this $N \times N$ matrix via a covariance function, also more commonly known as the kernel ($k$). In this framework, the covariance matrix is written as $\Sigma_{ij} = k(x_i, x_j, \theta) + \delta_{ij}\sigma^2$, $x$ being an input of the kernel, $\delta$ the Kronecker delta and $\sigma^2$ the variance term. The last term ensures the addition of Poisson or white noise to the diagonal of the covariance matrix. Our choice for the kernel is the squared exponential (SE), which for a multi-dimensional parameter space ($K$; defined above), in its additive form, is written as:

$$k_{SE}(x_i, x_j, \theta) = \zeta \exp\left[-\sum_{\alpha=1}^{K} \eta_\alpha (x_{\alpha,i} - x_{\alpha,j})^2\right]$$

We chose this form of the kernel as it is the *de facto* default form for Gaussian processes, being infinitely differentiable and easy to integrate against most functions. In this definition, $\zeta$ is the maximum covariance and $\eta_\alpha$ are the inverse scale parameters of the input vectors $\boldsymbol{x}$ which are essentially the columns of the $X$ matrix. In the description of our analytical transit model, we use the quadratic limb-darkening law[33] to describe the centre-to-limb intensity variations across the stellar disk. This selection was made through a comparison of various alternative laws, using the $\Delta$BIC formalism[34]. Finally, we generally assume uninformative, flat priors for the searched parameters, represented as $P(\theta, \phi)$ in equation (1), with the following exceptions:

$$\ln P(\theta, \phi) = \begin{cases} \ln \mathcal{N}(0.78884, 10^{-10}) & \text{for } \phi \in [P_0] \\ \ln \mathcal{N}(0, 10^{-10}) & \text{for } \phi \in [e] \\ -\infty & \text{if } c_1 + c_2 > 1 \\ \ln \mathcal{N}(\mu_i, 9\sigma^2) & \text{for } \theta \in [c_1, c_2] \\ \ln \Gamma(1, 1) & \text{for } \phi \in [\zeta, \eta_\alpha] \\ -\infty & \text{if } \sigma < 0 \end{cases}$$

What these prior assumptions essentially mean is that we fix the period, $P_0$, and eccentricity, $e$, of the transit to previously determined values, and set prior Gaussian distributions upon the values of the two limb-darkening coefficients, $c_1$ and $c_2$, on the basis of theoretically calculated values from PHOENIX models[35], with a width three times larger than the uncertainty in the theoretical model values for a greater flexibility in the search space. A gamma function, $\Gamma$, is chosen as the prior for the parameters of the kernel in order to encourage their values towards 0 if they are truly irrelevant in describing the correlated noise in the data[17].

We initially obtain values for the free parameters by using the Nelder–Mead simplex algorithm[36], and then find the maximum posterior solution by optimizing the log posterior given in equation (1). We obtain the posterior probability distribution functions by using the adaptive MCMC method, which explores the joint probability for our multivariate models. For each light curve analysed here, we

ran four independent MCMC simulations of length 100,000 iterations each, and checked their mutual convergence using the Gelman–Rubin[37] diagnostic. Example sets of samples drawn from such chains are given in Extended Data Fig. 6 for a broadband and a spectral light curve fit. We note that in addition to our analytical transit model, we also include a baseline function that is a quadratic polynomial of the parallactic angle throughout the observations that accounts for modulations introduced to the light curves by the rotation of the telescope, as well as the second-order, colour-dependent extinction of the stars. This last fact is the reason why we fit for the parameters of the baseline model of each spectrophotometric light curve, independently.

We infer best-fit parameter solutions from the analysis of those derived posterior distributions from all the MCMC simulations. Given that all fitted parameters have a Gaussian posterior distribution, our results for the mean, mode and median of each parameter are statistically identical, with the exception of those noise parameters with gamma distribution prior functions. In those cases, we quote the median of the distribution as the solution (Extended Data Table 2).

**Activity analysis.** It has been shown that transit light curves of WASP-19b towards the blue edge of the visible domain suffered substantially from the presence of stellar spots[8]. As a consequence, transmission spectra could not be produced from HST observations performed with the G430L grism of the Space Telescope Imaging Spectrograph (STIS) owing to the planet crossing such spot anomalies, whose temperature difference compared with the surrounding stellar photosphere is larger towards shorter wavelengths. Therefore, we assume that our observations could possibly be affected by stellar activity, although not to the same extent as the HST results, because our wavelength coverage from the blue data set starts at $0.44\,\mu m$ (as compared with $0.29\,\mu m$ from STIS). There are two avenues through which presence of stellar spots can influence the transmission spectrum, namely occulted and unocculted spots.

*Occulted spots.* Visual inspection shows that there is marginal evidence for a spot-crossing event in the blue data set after mid-transit. To calculate the impact of this anomaly, we modelled the light curves from the blue data set with a single spot-crossing event included as part of the analytical model. This approach meant the addition of four fitted parameters to the light curves to describe the physical properties of the spot anomaly: two positional variables (parallel and perpendicular to the transit chord), angular size ($R_\bullet$) in units of stellar radius ($R_\star$), and the contrast ratio ($r_\bullet$), which is wavelength dependent. We first modelled the broadband light curve from the blue data set as before, with the inclusion of four additional free parameters for the spot. This transit model including a spot-crossing event is shown in Extended Data Fig. 7a, quoting the derived spot parameters and the relative planetary radius and the limb-darkening coefficients. We detect a large spot anomaly overlapping with the transit chord, with a spectral radiance value very close to that of the surrounding photosphere, that is, a contrast ratio of 95%. In order to obtain the spot temperature, we modelled all of the spectrophotometric light curves of the blue data set with a spot model, only fitting for the contrast ratio, as this is the only wavelength-dependent aspect of the spot, the relation for which is given by a differential application of Planck's law:

$$r_\bullet \equiv \frac{B_\bullet(\lambda, T_\bullet)}{B_\star(\lambda, T_\star)} = \frac{e^{hc/(\lambda k_B T_\star)} - 1}{e^{hc/(\lambda k_B T_\bullet)} - 1} \qquad (2)$$

where $B$ is the spectral radiance of the spot ($\bullet$) or the star ($\star$); $h$ is Planck's constant; $c$ is the speed of light; $k_B$ is Boltzmann's constant; $\lambda$ is the wavelength at which the flux is measured; and $T$ is the temperature of the spot or stellar photosphere.

The stellar photospheric temperature is $5,568 \pm 71\,K$, obtained through precise spectroscopic measurements[38]. From fitting the spectral channels for a spot model, we obtained the wavelength dependence of the contrast ratio (Extended Data Fig. 7b). This relation is governed by equation (2), and, through a least-squares minimization approach, we obtained a value of $5,530 \pm 10\,K$ for the spot temperature—38 K cooler than the surrounding photosphere. Furthermore, we compared the initial transmission spectrum in the blue data set with that obtained by modelling the light curves with a spot model. These two sets of results are highly consistent, without any change to the obtained slope towards the shorter wavelengths (Extended Data Fig. 7c). This consistency highlights the strength of the Gaussian process red noise model in accounting for any systematic deviation from a transit function. We note that the error bars for the set that include a spot-crossing component are underestimated, because here no systematic noise model has been included. This is done to fully capture the impact of spot crossing on the differential light curves. We also note that there are degeneracies between spot latitude, angular size and contrast. Thus our estimation of the spot contrast should be used with caution.

*Unocculted spots.* WASP-19 is an active late G-star with spot anomalies on its surface that were detected by numerous previous studies[39–41]. Consequently, we have to assume that, in addition to the occulted spot detected here, there might be

spots on the stellar surface that are not occulted by the planetary transit chord. The presence of such unocculted spots can lead to an overestimation of the integrated flux from the stellar disk, and therefore to an overestimation of the planetary relative radius, which is a chromatic effect. To estimate the impact of such unocculted spots on the transmission spectrum in the blue end of the spectrum and therefore their contribution to the observed slope, we simulated transit light curves for WASP-19b over a rotating spotted star, using the SOAP-T tool[42]. For information minimization, we combined possible spots into a single spot with a 20% filling factor, on the basis of previous spot detections in transit light curves of this planet.

To estimate the maximum impact of this unocculted spot on the transit depth measurements, we selected the longitude and latitude of the combined spot to be at the centre of the stellar disk, because this set-up maximizes the rotational modulations introduced into the photometry. The spot contrast is calculated via Planck's law, with the temperature of the spot set to 200 K, 600 K and 1,000 K below the surrounding photospheric temperature of 5,568 K. We then simulated transit light curves for each of the wavelength bins in the blue data set, at each given spot temperature, adapting those limb-darkening coefficient values that were used as priors in modelling the spectral channels. These light curves were then modelled in a similar manner to the spectrophotometric time series of the observed data, and the inferred relative radii as a function of wavelength were compared with the observational results for each given spot temperature (Extended Data Fig. 7c). The observed slope can be explained partially by the possible presence of unocculted spots on the stellar disk, given large spot temperature differences of more than 1,000 K. A spot temperature difference of about 500 K has been reported for an occulted spot[39]; this finding, together with our much lower temperature measurement (about 40 K below the photosphere for the occulted spot), suggests that the observed slope towards the blue end of the spectrum cannot be explained simply by the presence of unocculted spots on the surface of the star. However, any atmospheric conclusions made from measurements of this slope have to be used with this fact in mind. A similar analysis has been performed to determine the impact of unocculted spots on the transmission spectroscopy of HD 189733b, ruling out such an impact as the cause of the observed transmission spectral signals[43] (that is, confirming the atmospheric nature of these signals).

**Retrieval analysis.** *Atmospheric models.* We modelled WASP-19b's atmosphere at the day–night terminator via 100 axially symmetric layers, uniformly spaced in log-pressure between $10^{-6}$ and $10^2$ bar. We take $R_p$ to be $1.31\,R_J$, $R_\star$ to be $0.993\,R_\odot$, and a surface gravity of $g = 14.3\,m\,s^{-2}$ (where $R_p$ is the planetary radius, given in units of Jupiter's equatorial radius, $R_J$, and $R_\star$ is the radius of the host star, given in units of the Sun's radius, $R_\odot$). We assume hydrostatic equilibrium, along with terminator-averaged temperature structure[44] and chemistry, with each species distributed uniformly with altitude. Clouds are parameterized by an opaque cloud deck, below which no electromagnetic radiation may pass. Hazes are included via a two-parameter power law[25], with two free parameters $a$ and $\Upsilon$: $\sigma(\lambda) = a\sigma_0(\lambda/\lambda_0)^\Upsilon$, where $\sigma(\lambda)$ is the wavelength-dependent haze cross-section, $\lambda_0$ is a reference wavelength (350 nm) and $\sigma_0$ is the $H_2$–Rayleigh scattering cross-section at the reference wavelength ($5.31 \times 10^{-31}\,m^2$). Our usage of the terms 'cloud' and 'haze' refers to the spectral features they cause, and makes no distinction as to the formation mechanisms and microphysics involved. Inhomogeneous cloud and haze distributions across the terminator are considered, parameterized by a cloud-coverage factor[22].

The chemical composition of the atmosphere is modelled as $H_2/He$ dominated, with a fixed $H_2/He$ ratio of 0.17. This background gas is considered to contain a mixture of trace molecular and atomic species, each with parameterized volume mixing ratios ranging between $10^{-16}$ and $10^{-1}$ of the total atmospheric content. We consider a wide range of chemical species that have prominent absorption features at visible wavelengths, namely: $H_2O$, Na, K, TiO, VO, AlO, TiH, NaH, MgH, CrH, CaH, ScH and FeH. We obtained the molecular line list for $H_2O$ from the HITEMP[45] database, while those for the metal oxides and hydrides originated from the ExoMol project[46]. For each molecular species, cross-sections were pre-computed line-by-line[47] and binned to a resolution of $1\,cm^{-1}$ on a grid of temperatures and pressures spanning the range $10^{-4}$ bar to $10^2$ bar, and 300 K to 3,500 K. The Na and K cross-sections were based on semi-analytic Lorentzian line profiles[48]. $H_2$–$H_2$ and $H_2$–He collision-induced absorption is included from the HITRAN database[49].

We computed the transmission spectrum of a given model atmosphere by integrating the stellar intensity, exponentially attenuated by the slant optical depth, over successive annuli for both cloud-free and uniformly cloudy terminator regions[22]. The two resulting spectra were then linearly superimposed, weighted by the terminator cloud-coverage factor. The stellar intensities during and outside transit were integrated over the solid angle subtended at the observer to obtain their respective fluxes, with the transit depth given by the fractional stellar flux difference induced as the planet transits its host star. For the visible wavelength spectrum of WASP-19b, we computed the transit depth at 2,000 wavelength points, spaced uniformly between $0.4\,\mu m$ and $1.1\,\mu m$. Model spectra were convolved with the instrument

point spread functions and integrated over grism response curves to produce binned model points corresponding to the blue, green and red data sets.

*Retrieval methodology.* We explored the range of atmospheric properties consistent with WASP19b's transmission spectrum by using the POSEIDON[22] atmospheric retrieval algorithm. This algorithm generates millions of model transmission spectra, mapping the high-dimensional parameter space via the MultiNest[50–52] multimodal nested sampling algorithm, implemented by the python wrapper PyMultiNest[53]. This allows statistical estimation of the underlying atmospheric parameters, in addition to establishing detection significances for model features (for example, chemical species, clouds or hazes) via nested Bayesian model comparison[54].

Each model atmosphere was parameterized by a six-parameter pressure–temperature profile[44], a four-parameter inhomogeneous cloud/haze prescription[22], the *a priori* unknown pressure at $R_p$, $P_{ref}$, and independent parameters for the volume mixing ratios of a subset of chemical species from Na, K, TiO, VO, AlO, TiH, NaH, MgH, CrH, CaH, ScH and FeH. Given the high resolution of the observed spectrum, we also allowed for the possibility of a systematic relative wavelength shift between the model and data as a nuisance parameter (often used in brown-dwarf retrievals of similar resolution[55,56]). This resulted in a maximum parameter space dimensionality of 25. Priors for each parameter were taken as either uniform (for example, top-of-atmosphere temperature) or uniform-in-the logarithm (for example, mixing ratios), depending on whether the prior range was less than or more than two orders of magnitude, respectively.

We initially assessed plausibility of each chemical species by computing a full 25-dimensional retrieval with 2,000 MultiNest live points. We observed flat posteriors (no evidence, beyond an upper bound) for all chemical species except Na, K, $H_2O$, TiO and MgH. The terminator cloud fraction sloped sharply towards the upper edge of the prior (100% cloud/haze coverage). With respect to this reference model, we ran two identical retrievals—one with the cloud deck removed and one with the haze removed. Removing the haze incurred a substantial penalty to the Bayesian evidence ($\ln Z = 987.2 \rightarrow 970.2$), corresponding to a Bayes factor of $2.0 \times 10^7$ and equivalent to a $6.2\sigma$ detection of a uniform haze across the terminator (with respect to the full chemistry model). The Bayesian evidence was unchanged within the error bars (about 0.1) when the cloud deck was removed, so the data do not support the presence of a cloud deck. We found a relative wavelength shift of $-72.8^{+10.0}_{-11.3}$ Å between the data and the model ($\ln Z = 987.0 \rightarrow 970.6$, compared with a model in which the shift is fixed at 0, equivalent to a $6.1\sigma$ detection of a wavelength shift).

A second round of retrievals was conducted to assess the potential presence of Na, K, $H_2O$, TiO and MgH. Given the strong evidence for uniform hazes from the first round, we fixed the cloud fraction to 100% for these runs, leading to a dimensionality of 16 for the reference model. We conducted a nested Bayesian model comparison by removing individual chemical species from the reference model, running a new retrieval for each, and computing the change in the Bayesian evidence. We observed little change, or only a very slight increase, when we removed K and MgH, indicating that these species were not detected and that their presence is suggested only marginally by the data. However, we observed changes in $\ln Z > 4$ when Na, $H_2O$ or TiO was removed, suggesting that they are present.

We conducted a final round of retrievals by using a minimal reference model containing just Na, $H_2O$, TiO, haze and a cloud deck—corresponding to 14 free parameters. This simple reference model has a Bayesian evidence of $\ln Z = 990.71$, a notable increase over the initial reference model, reaffirming that the additional model complexity incurred by the large collection of other metal oxides/hydrides is not justified in light of the data. The results of the Bayesian model comparison, minimum best-fit reduced chi-squares, and detection significances for Na, TiO, $H_2O$ and the terminator-spanning haze, with respect to this reference model, are given in Extended Data Table 3. The overall best-fitting transmission spectrum, drawn from the posterior of the reference model (shown in red in Fig. 2), has Na, $H_2O$ and TiO volume mixing ratios of 6.4 p.p.m., 180 p.p.m. and 0.12 p.p.b., respectively; haze parameters $a = 470,000$ and $\Upsilon = -30$; a wavelength shift of $-73.6$ Å; and a reference pressure of 0.93 bar.

It has been claimed that accurate $H_2O$ abundances cannot be extracted from HST WFC3 transmission spectra[57], unless a reference radius and pressure are assumed. This degeneracy may be lifted by observations of cloud-free atmospheres at visible wavelength, as the $H_2$–Rayleigh slope fixes the reference pressure[57]. Here, we have demonstrated that constraints of less than 1 dex can be placed on individual molecular abundances when retrieving the visible wavelength transmission spectrum of a hazy atmosphere, despite making no *a priori* assumptions as to the value of $P_{ref}$. Our ability to accomplish this stems from the data points that are less than about 0.5 μm, which sample the continuum slope owing to our parametric haze model. This slope, in turn, provides a normalization to the transmission spectrum and enables a loose constraint to be placed on $P_{ref}$, allowing
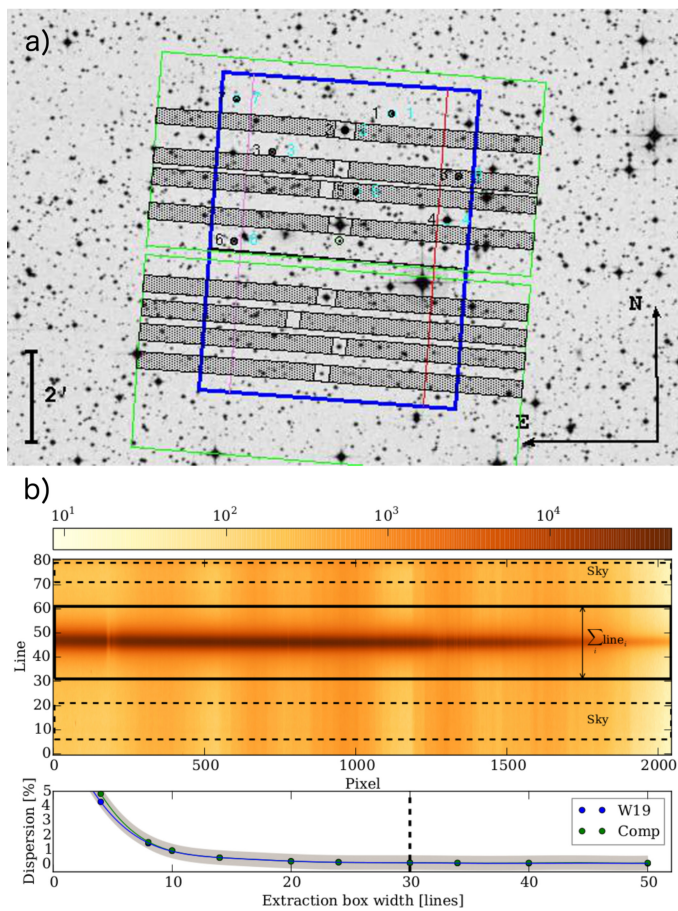
the degeneracy with the molecular abundances to be partially collapsed. We stress that observations at short wavelengths, away from spectral features, are essential for constraining the absolute abundances of individual species.

**Data availability.** The data used in this work can be accessed at the ESO science archive (archive.eso.org), using identification numbers 60.A-9203(F) for the green data set, and 96.C-0465(B,C) for the blue and red data. Reduced one-dimensional and two-dimensional frames are available from E.S. upon reasonable request. In addition, the broadband (Fig. 1) and spectrophotometric light curves (Extended Data Figs 3–5), as well as the transmission spectrum data and models (Fig. 2), are included as source data online.
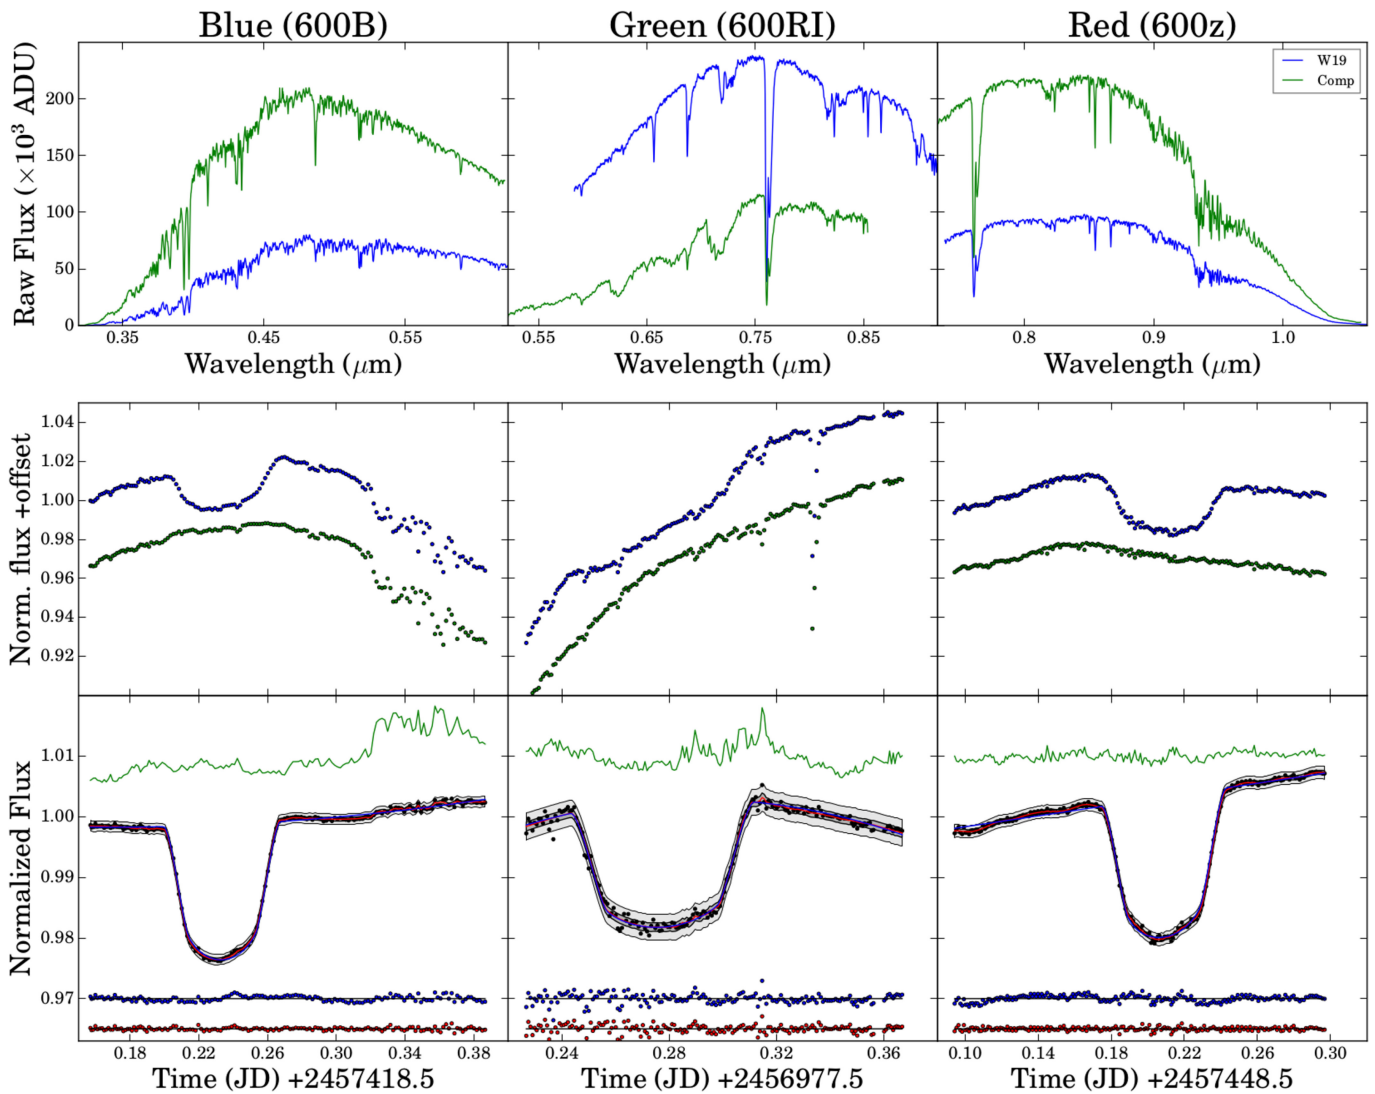
**Code availability.** The GP (GeePea) and the MCMC inference (Infer) modules are written in python programming language and are freely available from https://github.com/nealegibson. The multi-modal nested sampling algorithm MultiNest and its python wrapper PyMultiNest are freely available from https://ccpforge.cse.rl.ac.uk/gf/project/multinest/ and https://github.com/JohannesBuchner/PyMultiNest. The SOAP-T tool is also freely available at http://astro.up.pt/resources/soap-t/.

26. Appenzeller, I. *et al.* Successful commissioning of FORS1—the first optical instrument on the VLT. *ESO Messenger* **94**, 1–6 (1998).
27. Sedaghati, E. *et al.* Regaining the FORS: optical ground-based transmission spectroscopy of the exoplanet WASP-19b with VLT + FORS2. *Astron. Astrophys.* **576**, L11 (2015).
28. Horne, K. An optimal extraction algorithm for CCD spectroscopy. *Publ. Astron. Soc. Pacif.* **98**, 609–617 (1986).
29. Lendl, M. *et al.* FORS2 observes a multi-epoch transmission spectrum of the hot Saturn-mass exoplanet WASP-49b. *Astron. Astrophys.* **587**, A67 (2016).
30. Nikolov, N. *et al.* VLT FORS2 comparative transmission spectroscopy: detection of Na in the atmosphere of WASP-39b from the ground. *Astrophys. J.* **832**, 191 (2016).
31. Gibson, N. P. *et al.* VLT/FORS2 comparative transmission spectroscopy II: confirmation of a cloud deck and Rayleigh scattering in WASP-31b, but no potassium? *Mon. Not. R. Astron. Soc.* **467**, 4591–4605 (2017).
32. Rasmussen, C. & Williams, C. *Gaussian Processes for Machine Learning* (MIT Press, 2006).
33. Kopal, Z. Detailed effects of limb darkening upon light and velocity curves of close binary systems. *Harvard College Observ. Circular* **454**, 1–12 (1950).
34. Schwarz, G. *et al.* Estimating the dimension of a model. *Ann. Stat.* **6**, 461–464 (1978).
35. Claret, A., Hauschildt, P. & Witte, S. New limb-darkening coefficients for Phoenix/1d model atmospheres-II. Calculations for 5000 K ≲ $T_{eff}$ ≲ 10,000 K Kepler, CoRoT, Spitzer, uvby, UBVRIJHK, Sloan, and 2MASS photometric systems. *Astron. Astrophys.* **552**, A16 (2013).
36. Nelder, J. A. & Mead, R. A simplex method for function minimization. *Comput. J.* **7**, 308–313 (1965).
37. Gelman, A., Carlin, J. B., Stern, H. S. & Rubin, D. B. *Bayesian Data Analysis* vol. 2 (Chapman and Hall, 2014).
38. Torres, G. *et al.* Improved spectroscopic parameters for transiting planet hosts. *Astrophys. J.* **757**, 161 (2012).
39. Tregloan-Reed, J., Southworth, J. & Tappert, C. Transits and starspots in the WASP-19 planetary system. *Mon. Not. R. Astron. Soc.* **428**, 3671–3679 (2013).
40. Mancini, L. *et al.* Physical properties, transmission and emission spectra of the WASP-19 planetary system from multi-colour photometry. *Mon. Not. R. Astron. Soc.* **436**, 2–18 (2013).
41. Mandell, A. M. *et al.* Exoplanet transit spectroscopy using WFC3: WASP-12b, WASP-17b, and WASP-19b. *Astrophys. J.* **779**, 128 (2013).
42. Oshagh, M. *et al.* SOAP-T: a tool to study the light curve and radial velocity of a system with a transiting planet and a rotating spotted star. *Astron. Astrophys.* **549**, A35 (2013).
43. McCullough, P., Crouzet, N., Deming, D. & Madhusudhan, N. Water vapor in the spectrum of the extrasolar planet HD 189733b. I. the transit. *Astrophys. J.* **791**, 55 (2014).
44. Madhusudhan, N. & Seager, S. A temperature and abundance retrieval method for exoplanet atmospheres. *Astrophys. J.* **707**, 24–39 (2009).
45. Rothman, L. S. *et al.* HITEMP, the high-temperature molecular spectroscopic database. *J. Quant. Spectrosc. Radiat. Transf.* **111**, 2139–2150 (2010).
46. Tennyson, J. & Yurchenko, S. N. ExoMol: molecular line lists for exoplanet and other atmospheres. *Mon. Not. R. Astron. Soc.* **425**, 21–33 (2012).
47. Hedges, C. & Madhusudhan, N. Effect of pressure broadening on molecular absorption cross sections in exoplanetary atmospheres. *Mon. Not. R. Astron. Soc.* **458**, 1427–1449 (2016).
48. 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).
49. Richard, C. *et al.* New section of the HITRAN database: collision-induced absorption (CIA). *J. Quant. Spectrosc. Radiat. Transf.* **113**, 1276–1285 (2012).
50. Feroz, F. & Hobson, M. P. Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses. *Mon. Not. R. Astron. Soc.* **384**, 449–463 (2008).
51. Feroz, F., Hobson, M. P. & Bridges, M. MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics. *Mon. Not. R. Astron. Soc.* **398**, 1601–1614 (2009).

52. Feroz, F., Hobson, M. P., Cameron, E. & Pettitt, A. N. Importance nested sampling and the MultiNest algorithm. Preprint at https://arxiv.org/abs/1306.2144 (2013).

53. Buchner, J. *et al.* X-ray spectral modelling of the AGN obscuring region in the CDFS: Bayesian model selection and catalogue. *Astron. Astrophys.* **564,** A125 (2014).

54. Trotta, R. Bayes in the sky: Bayesian inference and model selection in cosmology. *Contemp. Phys.* **49,** 71–104 (2008).

55. Line, M. R., Teske, J., Burningham, B., Fortney, J. J. & Marley, M. S. Uniform atmospheric retrieval analysis of ultracool dwarfs. I. Characterizing benchmarks, Gl 570D and HD 3651B. *Astrophys. J.* **807,** 183 (2015).

56. Burningham, B. *et al.* Retrieval of atmospheric properties of cloudy L dwarfs. *Mon. Not. R. Astron. Soc.* **470,** 1177–1197 (2017).

57. Heng, K. & Kitzmann, D. The theory of transmission spectra revisited: a semi-analytical method for analyzing WFC3 data and an unresolved challenge. *Mon. Not. R. Astron. Soc.* **470,** 2972–2981 (2017).
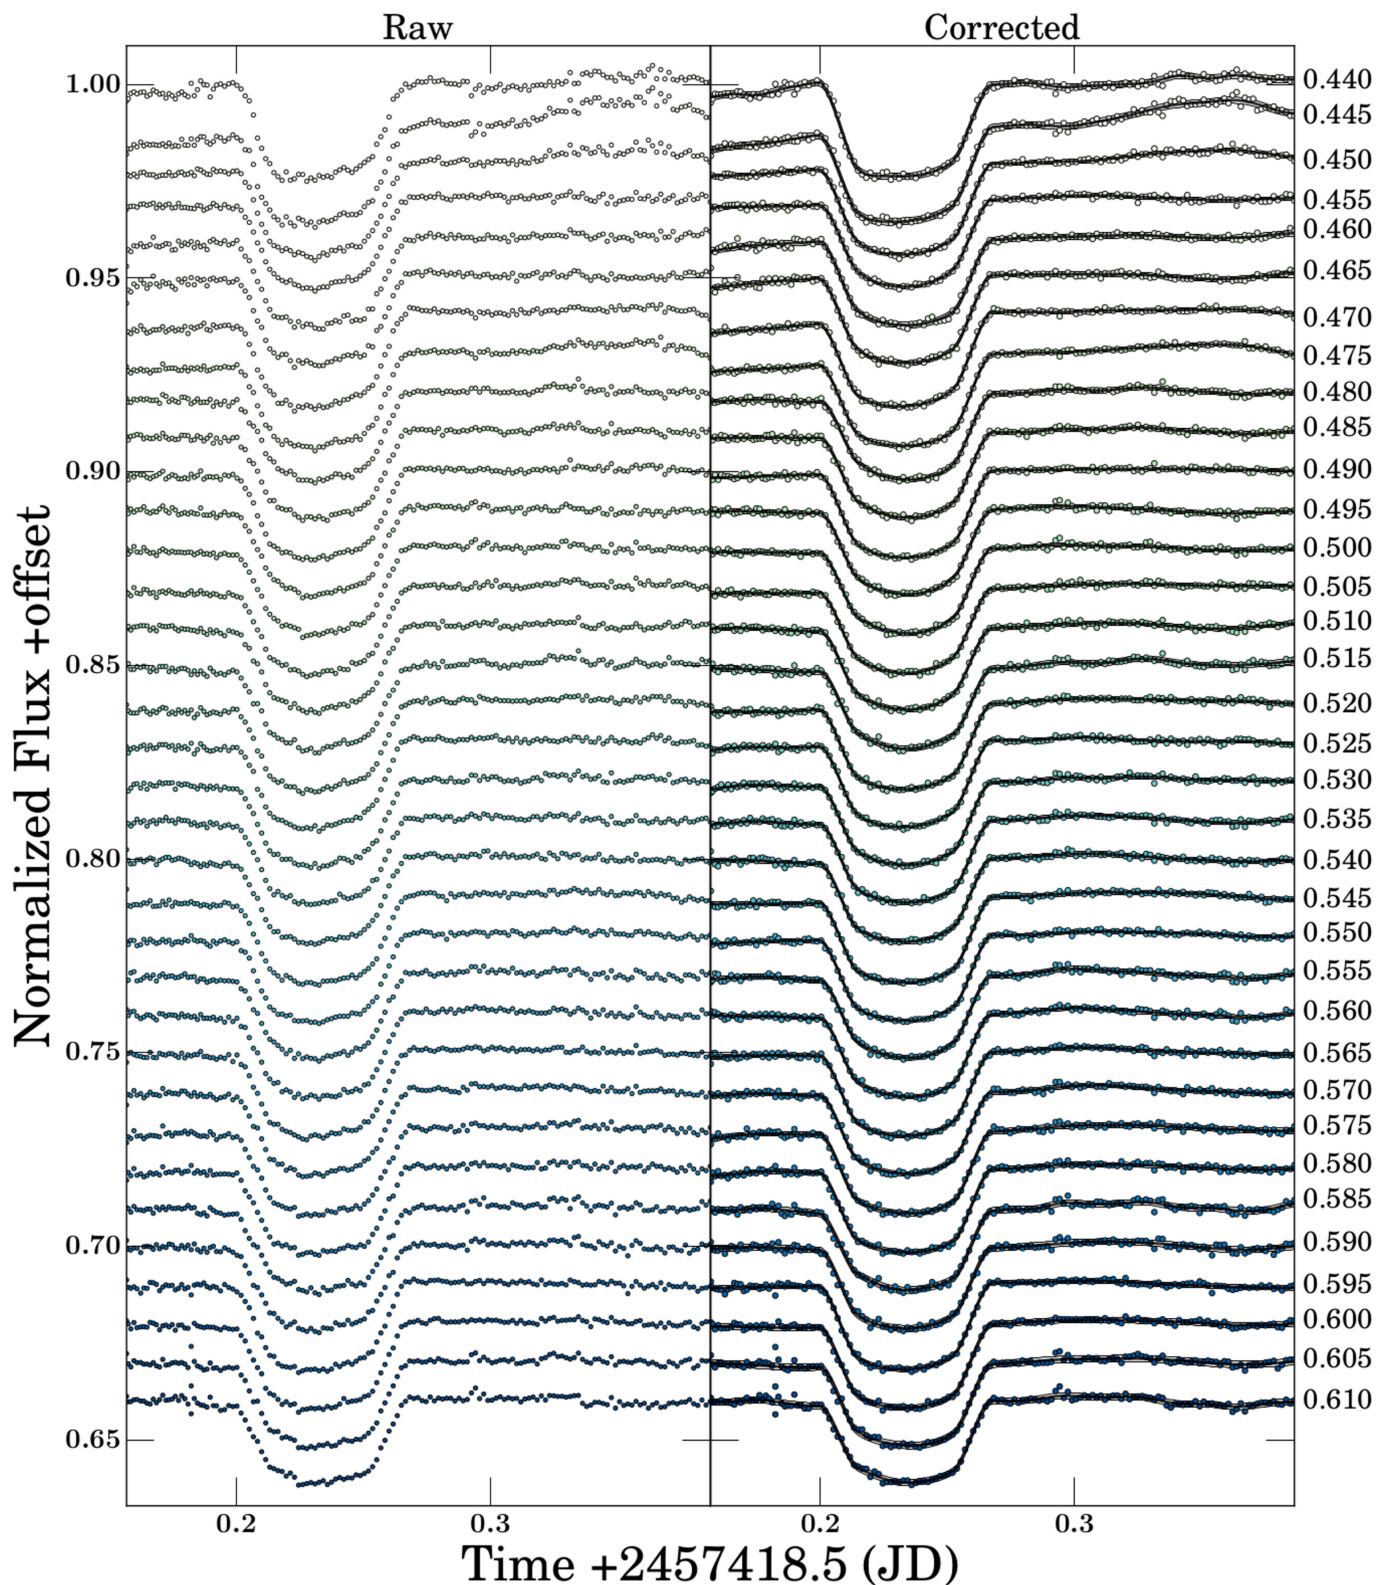
**Extended Data Figure 1 | Observations and reduction. a**, An example of a mask design used for MXU observations. The field of view of the FORS2 spectrograph is shown in blue, with the green lines indicating the two-chip detector mosaic. The grey shaded regions show the areas of the detectors used for recording the stellar spectra. In the instance shown, WASP-19 is the star in the upper-most slit. **b**, Top, an example of a two-dimensional spectrum, from the red data set, extracted from a frame taken with the FORS2 instrument. The final size of the extraction box and the regions used for sky subtraction are indicated. Bottom, the process of choosing the width of the extraction box, where the final value is shown as a dashed line and the grey shading represents the $1\sigma$ confidence limits. The exemplar frame used to produce these plots is selected at large seeing. W19, WASP-19; comp., comparison star.
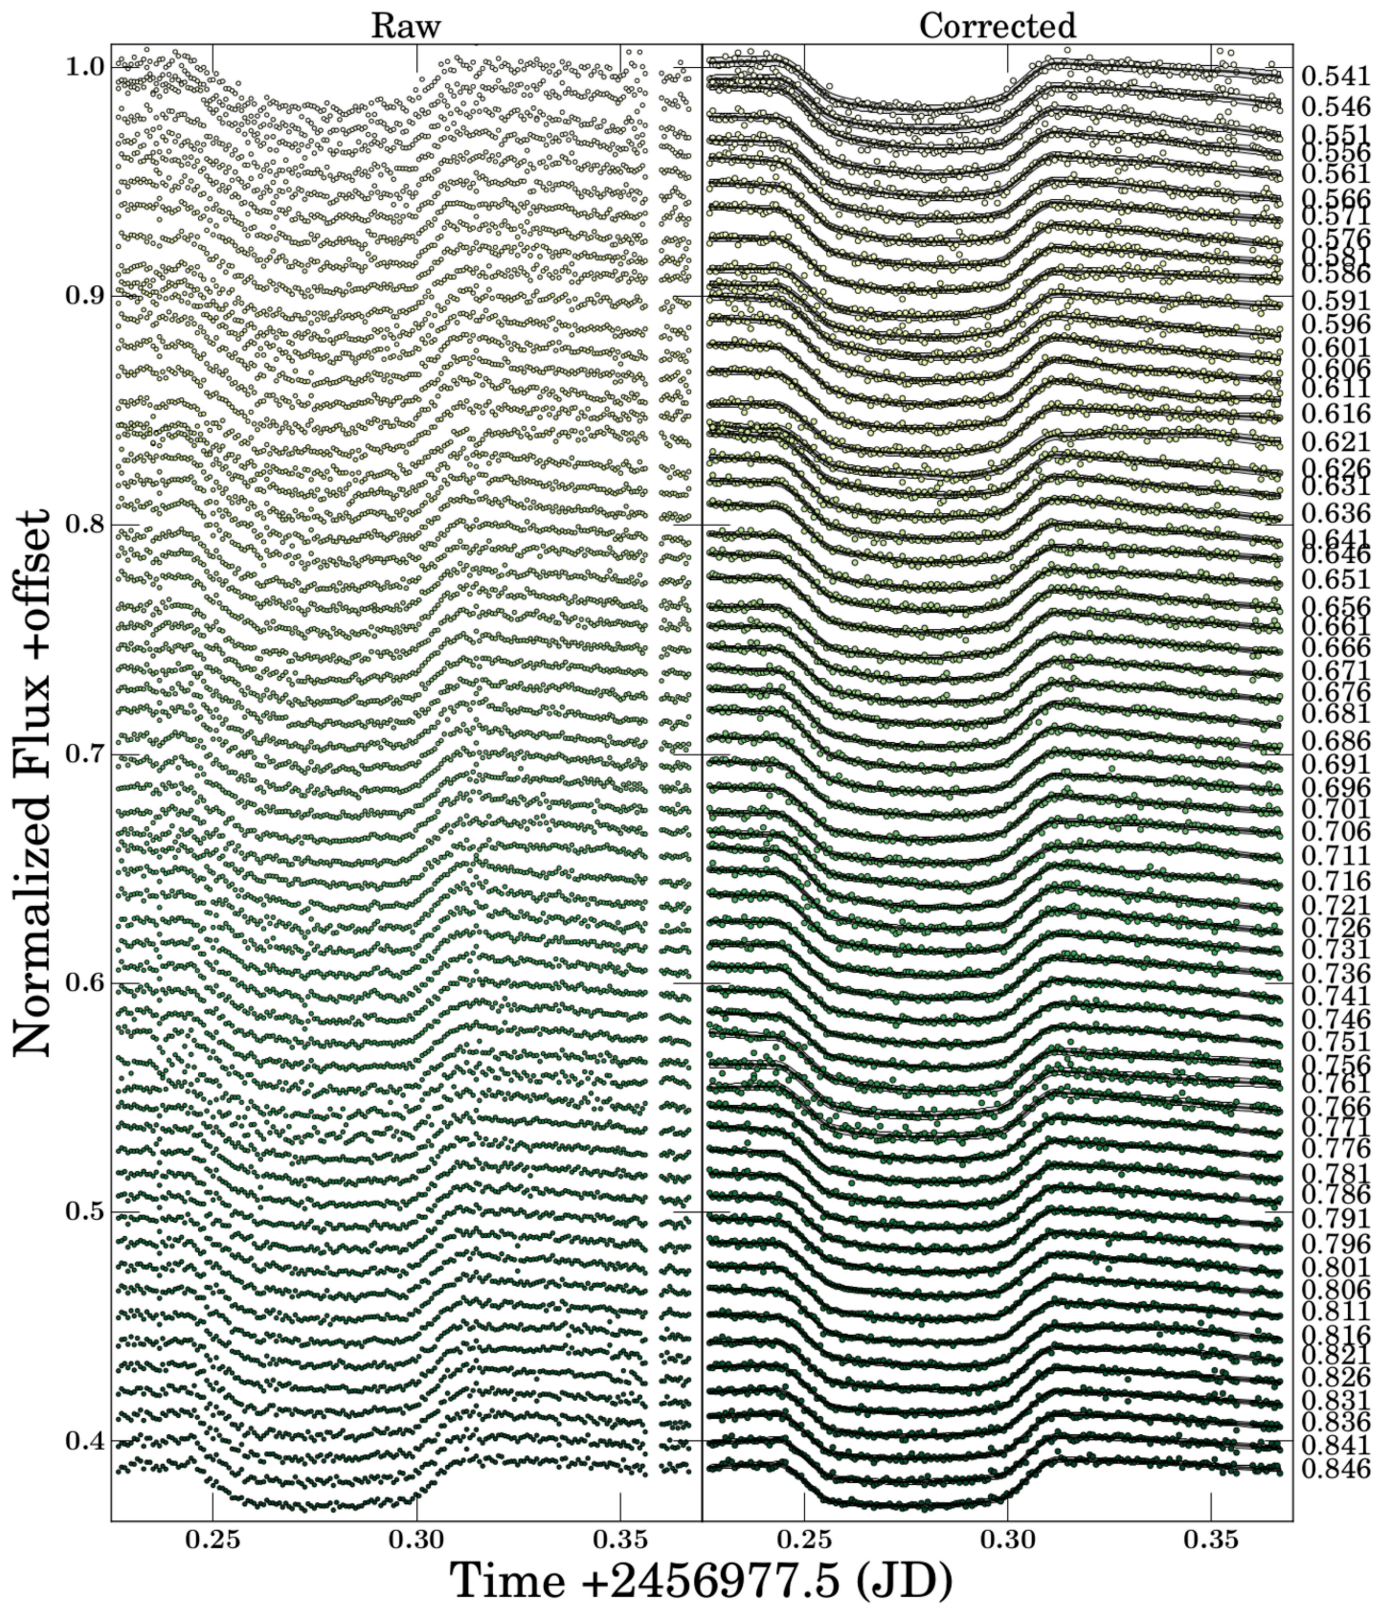
**Extended Data Figure 2 | Spectra and light curves.** Top row, an example set of spectra for the target (WASP-19) and the chosen comparison star, for each observing run (blue, green and red). Counts are given in analogue-to-digital units (ADUs), read directly from the sum of values of charge-coupled-device (CCD) pixels. Middle row, light curves for the target and comparison stars for each data set, obtained through broadband integration of the series of spectra. Colours match those in the top row; values are normalized to the mean of the out-of-transit fluxes and shifted for clarity. The transit imprint from WASP-19b is clearly evident even in these raw light curves. Bottom row, differential broadband light curves (black data point) obtained simply by dividing the two light curves in the middle row. We also show our fitted transit model (blue curve) and the Gaussian process systematic noise model (red curve) with its $1\sigma$ (dark grey shading) and $3\sigma$ (light grey shading) confidence levels. The points below are the residuals of the two models, where the colours correspond to the fit that they represent. The green line shows the flux variations resulting from changes in seeing conditions, used as an input for our Gaussian process model. JD, Julian day.
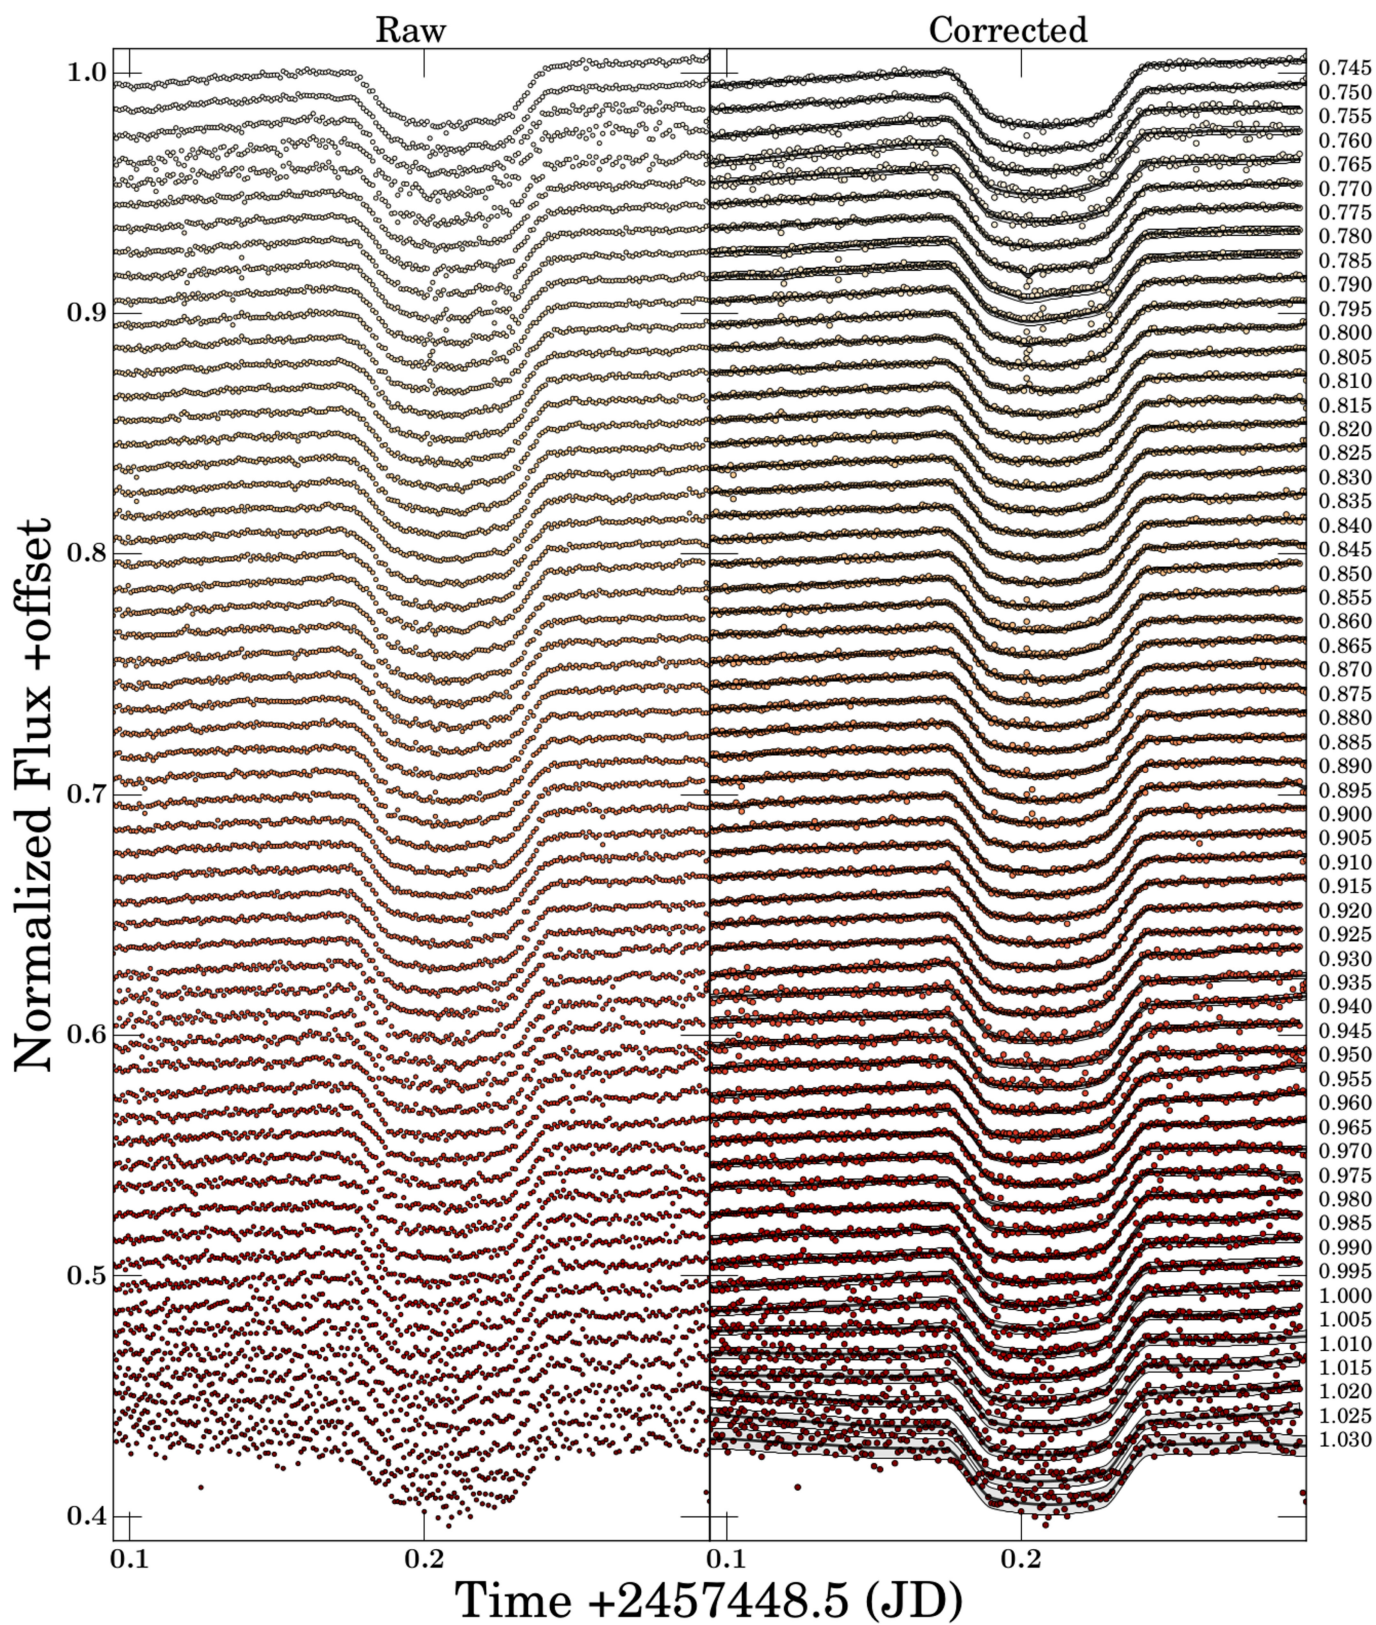
**Extended Data Figure 3 | Spectrophotometric light curves for the blue data set.** Left, raw light curves produced from each of the narrow-band channels in the blue data set. Right, those light curves that have been corrected for the common-mode systematics. Our best-fit Gaussian process systematic noise models are shown as solid black lines, with the centre of the integration bin for each light curve given to the right of it in micrometres. All light curves have been shifted vertically for clarity.

**Extended Data Figure 4 | Spectrophotometric light curves for the green data set.** As for Extended Data Fig. 3 but for the green data set.

**Extended Data Figure 5 | Spectrophotometric light curves for the red data set.** As for Extended Data Figs 3 and 4 but for the red data set.

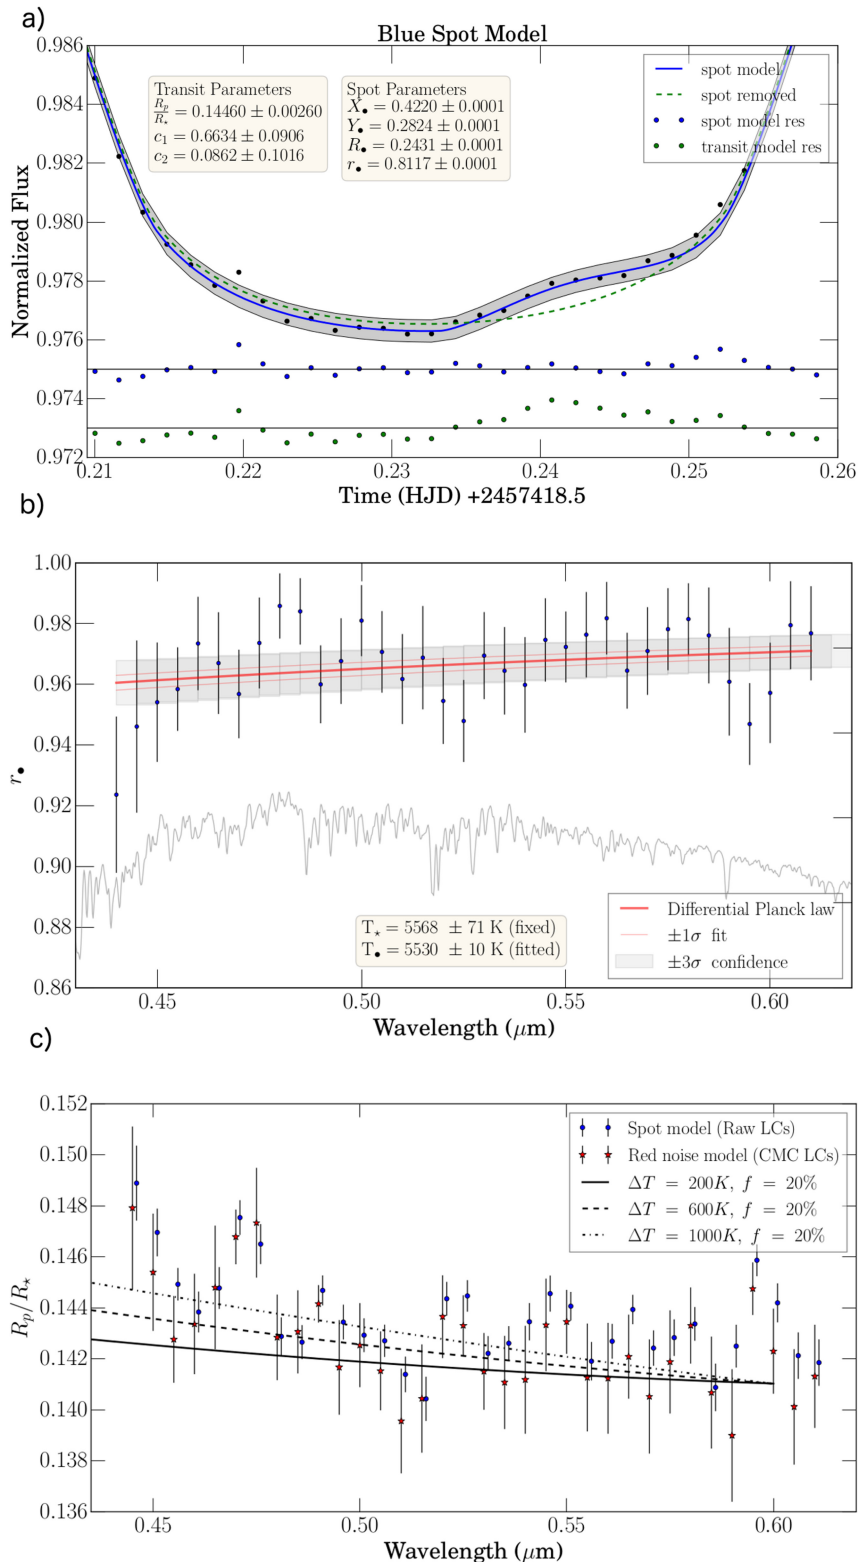**Extended Data Figure 6 | Correlations.** Random samples drawn from the four MCMC simulations, for all the fitted parameters (see Methods for definitions), in modelling a broadband light curve (lower-left triangle) and a spectroscopic light curve (upper-right triangle). Both examples are from the blue data set. Mutual convergence of all independent chains is evident, as are the well documented degeneracies between the impact parameter ($b$) and the scaled semi-major axis ($a/R_\star$), and between the two coefficients of the limb-darkening law ($c_1$ and $c_2$). $\eta_{fwhm}$ is the Gaussian process inverse length scale for 'seeing'.

**Extended Data Figure 7 | Stellar activity impact. a**, Broadband light curve from the blue data set, modelled using an analytical model that includes a spot-crossing event by the planet. The new inferred planetary radius and the limb-darkening coefficients (the 'transit parameters') are shown, for which the offset to our previous results is substantially lower than the derived precision. This is because our systematic model accounts well for this anomaly. The inferred spot parameters are also shown. HJD, heliocentric Julian day; res, residuals. **b**, Dependence of spot contrast ratio on the observation wavelength, from which the spot temperature is determined using Planck's law. The $1\sigma$ error bars are derived from a joint analysis of posterior probability distributions of the relative planetary radius, from the MCMC simulations. For reference, a spectrum of WASP-19 is plotted in light grey. The prior stellar photospheric temperature ($T_\star$) and the fitted spot temperature ($T_\bullet$) are also given. **c**, Comparison of transmission spectra in the blue data set, from red noise analysis and spot modelling; $1\sigma$ error bars were derived as above. The spot-analysis results (blue points) have been shifted by $+0.01\,\mu$m to better distinguish between the two sets of results. Wavelength-dependent radius variations induced by the presence of spots with 20% filling factor (f) and temperature differences of 200 K, 600 K or 1,000 K are also plotted. LC, light curve; CMC, common mode corrected.

**Extended Data Table 1 | Observational information on the three sets of observational campaigns**

| Dataset | Date | Airmass | Median seeing ($''$) | Grism 600 | $\lambda$ ($\mu$m) | Readout mode | Exposures (in-transit) |
|---------|------|---------|-------------|-----------|-------------|--------------|------------------------|
| Blue | 30-01-16 | 1.20→1.07→1.42 | 1.49 | B | 0.330-0.620 | 200kHz | 150 (43) |
| Green | 15-11-14 | 2.57→1.19 | 1.13 | RI | 0.536-0.853 | 100kHz | 180 (88) |
| Red | 29-02-16 | 1.15→1.07→1.47 | 1.18 | z | 0.740-1.051 | 200kHz | 212 (86) |

**Extended Data Table 2 | Transit parameters from broadband analysis of all three data sets**

| Parameter (inferred) | Blue | Green | Red |
| --- | --- | --- | --- |
| | 4,865 (2,530) Å | 6,941.5 (3,172) Å | 8,951.5 (3,111) Å |
| Mid-Transit, $T_c$ (JD) +2,400,000 | 57,418.73368±0.00010 | 56,977.77754±0.00019 | 57,448.70885±0.00020 |
| $T_c$, Barycentric corrected (BJD$_{TDB}$) | 57,418.73688 | 56,977.77653 | 57,448.71294 |
| Period, P (days) [fixed] | | 0.78884 | |
| Eccentricity, $e$ [fixed] | | 0.0 | |
| Scaled semi-major axis, $a/R_\star$ | | 3.5875±0.0574 | |
| Impact parameter, $b$ | | 0.6525±0.0233 | |
| Relative planetary radius, $R_p/R_\star$ | 0.14440±0.00151 | 0.14366±0.00181 | 0.14060±0.00148 |
| Linear LD coefficient, $c_1$ | 0.6709±0.0616 | 0.5045±0.1037 | 0.3152±0.1195 |
| Quadratic LD coefficient, $c_2$ | 0.0599±0.0866 | 0.1361±0.0934 | 0.1673±0.1074 |
| GP output scale, $\zeta$ (ppm) | $748^{+418}_{-220}$ | $650^{+1,723}_{-218}$ | $644^{+308}_{-181}$ |
| GP inverse length scale for $t$, $\eta_t$ | $0.028^{+0.010}_{-0.009}$ | $0.027^{+0.093}_{-0.015}$ | $0.008^{+0.002}_{-0.001}$ |
| GP inverse length scale for *seeing*, $\eta_{fwhm}$ | $0.006^{+0.004}_{-0.002}$ | $0.010^{+0.018}_{-0.003}$ | $0.006^{+0.013}_{-0.002}$ |
| Poisson noise, $\sigma_w$ (ppm) | 276±19 | 729±71 | 318±20 |
| Parameter (derived) | | | |
| Semi-major axis, $a$ (AU) | | 0.01651±0.00064 | |
| Orbital inclination, $i$ (°) | | $79.52^{+0.54}_{-0.56}$ | |
| Planet radius, $R_p$ ($R_{jup}$) | 1.3907±0.0426 | 1.3836±0.0454 | 1.3541±0.0416 |

For each data set, the band centre and (in parentheses) the bandwidths are given at the top of the column. The data that are shown only in the 'green' column are common to all columns. GP, Gaussian process; LD, limb darkening; $R_{jup}$, radius of Jupiter.

**Extended Data Table 3 | Bayesian model comparison detections of WASP-19b's terminator chemistry and cloud properties**

| Model | Evidence $\ln \mathcal{Z}_i$ | Best-fit $\chi^2_{r,\min}$ | Bayes factor $\mathcal{B}_{0i}$ | Detection of Ref. |
|---|---|---|---|---|
| Reference | 990.71 | 1.63 | Ref. | Ref. |
| No TiO | 963.17 | 2.06 | $9.1 \times 10^{11}$ | $7.7\sigma$ |
| No $H_2O$ | 961.83 | 2.05 | $3.5 \times 10^{12}$ | $7.9\sigma$ |
| No Na | 986.36 | 1.65 | 77 | $3.4\sigma$ |
| No Haze | 965.33 | 1.95 | $1.0 \times 10^{11}$ | $7.4\sigma$ |

The 'reference' (ref.) model includes opacity resulting from $H_2$, He, Na, $H_2O$ and TiO, along with a parameterized cloud and haze prescription. $\mathcal{Z}_i$ is the Bayesian evidence of the $i$th model; $B_{0i}$ is the Bayes factor corresponding to the preference for the reference model over model $i$; and $\chi^2_{r,min}$ is the minimum reduced chi-squared value for the best-fitting spectrum within each model. An $n\sigma$ detection (where $n \geq 3$) indicates the degree of preference for the reference model over the alternative model.