

2.5D retrieval of atmospheric properties from exoplanet phase curves: application to WASP-43b observations

Patrick G. J. Irwin¹,¹★ Vivien Parmentier,¹ Jake Taylor¹,¹ Jo Barstow¹,²
Suzanne Aigrain,¹ Graham K. H. Lee¹ and Ryan Garland¹

¹Department of Physics, University of Oxford, Parks Road, Oxford OX1 3PU, UK

²Department of Physics and Astronomy, University College London, Gower St, London WC1E 6BT, UK

Accepted 2020 January 12. Received 2019 December 9; in original form 2019 September 7

ABSTRACT

We present a novel retrieval technique that attempts to model phase curve observations of exoplanets more realistically and reliably, which we call the 2.5-dimensional (2.5D) approach. In our 2.5D approach we retrieve the vertical temperature profile and mean gaseous abundance of a planet at all longitudes and latitudes simultaneously, assuming that the temperature or composition, x , at a particular longitude and latitude (Λ, Φ) is given by $x(\Lambda, \Phi) = \bar{x} + (x(\Lambda, 0) - \bar{x}) \cos^n \Phi$, where \bar{x} is the mean of the morning and evening terminator values of $x(\Lambda, 0)$, and n is an assumed coefficient. We compare our new 2.5D scheme with the more traditional 1D approach, which assumes the same temperature profile and gaseous abundances at all points on the visible disc of a planet for each individual phase observation, using a set of synthetic phase curves generated from a GCM-based simulation. We find that our 2.5D model fits these data more realistically than the 1D approach, confining the hotter regions of the planet more closely to the dayside. We then apply both models to WASP-43b phase curve observations of *HST*/WFC3 and *Spitzer*/IRAC. We find that the dayside of WASP-43b is apparently much hotter than the nightside and show that this could be explained by the presence of a thick cloud on the nightside with a cloud top at pressure <0.2 bar. We further show that while the mole fraction of water vapour is reasonably well constrained to $(1-10) \times 10^{-4}$, the abundance of CO is very difficult to constrain with these data since it is degenerate with temperature and prone to possible systematic radiometric differences between the *HST*/WFC3 and *Spitzer*/IRAC observations. Hence, it is difficult to reliably constrain C/O.

Key words: radiative transfer – methods: numerical – planets and satellites: atmospheres – planets and satellites: individual: WASP-43b.

1 INTRODUCTION

Studying the atmospheres of exoplanets is a field that has expanded rapidly in the past decade. Transiting exoplanets have been studied in both primary transit (e.g. Barstow et al. 2013, 2017; Evans et al. 2018; Krissansen-Totton et al. 2018) and secondary eclipse (e.g. Barstow et al. 2014; Lee et al. 2014; Gandhi & Madhusudhan 2018), with each technique providing a unique insight of the atmospheres of these alien worlds. The primary transit spectrum allows us to probe composition at high altitudes at the day/night terminator of these planets and a number of studies have been made comparing the observed properties and inferred atmospheres of a representative range of planets (e.g. Sing et al. 2016; Barstow et al. 2017; Pinhas et al. 2019), finding a wide range of apparent H₂O abundances and

obscuration by clouds, with in one study the presence or absence of clouds apparently dependent on the effective temperature of the planet (Barstow et al. 2017). Meanwhile the secondary eclipse spectrum allows us to probe the temperature and composition of the dayside. Recently, using the *Hubble*/WFC3 and *Spitzer*/IRAC instruments, it has become possible to obtain spectroscopic observations of the full phase curve of three exoplanets: WASP-43b (e.g. Stevenson et al. 2014; Mendonça et al. 2018), discovered by Hellier et al. (2011), WASP-103b (Kreidberg et al. 2018), discovered by Gillon et al. (2014), and WASP-18b (e.g. Maxted et al. 2013; Arcangeli et al. 2019), discovered by Hellier et al. (2009). The *HST*/WFC3 observations cover wavelengths from ~ 1 to $2 \mu\text{m}$ at a spectral resolution of $\Delta\lambda \sim 0.035 \mu\text{m}$, while the broad *Spitzer*/IRAC channel observations at 3.6 and $4.5 \mu\text{m}$, which have higher signal-to-noise ratio (SNR) but lower spectral resolution ($\Delta\lambda \sim 1 \mu\text{m}$), extend the coverage to higher altitudes with different gaseous sensitivities. Techniques have been developed to invert such

* E-mail: patrick.irwin@physics.ox.ac.uk

phase curves to derive longitudinal brightness temperature maps at individual wavelengths (e.g. Burrows, Budaj & Hubeny 2008; Cowan & Agol 2008), but to perform full atmospheric ‘retrievals’ of such data, where we attempt to find a global distribution of temperature and abundance that fits the measured phase curves at all wavelengths simultaneously, the prevalent current technique has been to fit the observed spectra from each phase individually assuming a model that has the same temperature profile and gaseous abundances at all points on the visible disc; this process is then repeated for all the other phases of the planet’s orbit. Such a ‘1D’ model approach ignores the fact that atmospheric conditions will vary across the face of the observable disc and also fails to utilize the constraints on the atmospheric structure arising from observations at nearby phases. A study of former effect has been made by Blečić, Dobbs-Dixon & Greene (2017) who used a retrieval model to retrieve a single atmospheric temperature profile from a secondary eclipse spectra generated by a general circulation model (GCM) simulation of HD189733b. Meanwhile, Feng et al. (2016) explored how the assumption of a single 1D thermal profile can bias the interpretation of synthetic thermal emission spectra generated from a model atmosphere composed of two different thermal profiles (to approximate the appearance of a Hot Jupiter with a sub-stellar region much hotter than the surrounding background). In some cases it was found that considerable biases could arise.

In this paper, we present a novel technique that performs a more consistent retrieval of the complete phase curve to extract atmospheric structure and composition at all latitudes and longitudes simultaneously, which we believe represents a considerable improvement over the traditional 1D approach. In Section 2, we review the numerical process of simulating disc-averaged spectra and propose an efficient scheme for computing this numerically and in Section 3 we introduce our new 2.5D retrieval scheme. Section 4 describes the radiative transfer model we use in our retrieval model and Section 5 describes how we validated our retrieval model with synthetic phase curve spectra generated from a model based on a GCM simulation. In Section 6, we apply our new scheme to observed phase curve spectra of Stevenson et al. (2017) of WASP-43b and we discuss our findings and present our conclusions in Sections 7 and 8.

2 CALCULATING DISC-AVERAGED SPECTRA FROM INHOMOGENEOUS ATMOSPHERES

When calculating the disc-averaged spectra of planets there are several numerical techniques that can be used, which have various levels of precision. In this work, where we attempt to retrieve observations from disc-averaged spectra of inhomogeneous atmospheres, it was vital that the disc-averaging was treated as accurately as possible, but also that we used a scheme that was not too computationally expensive; this is important since in a retrieval we have to iterate over many possible solutions before we arrive at our best estimate.

At the most basic level, the contribution to the total spectral irradiance, $I(\lambda)$ (units of $\text{W m}^{-2} \mu\text{m}^{-1}$), at wavelength λ , seen from an area on the ‘surface’ of a planet is given by

$$dI(\lambda) = R(\lambda) d\Omega, \quad (1)$$

where $R(\lambda)$ is the spectral radiance ($\text{W m}^{-2} \text{sr}^{-1} \mu\text{m}^{-1}$) calculated for a particular area on the disc and $d\Omega$ is the solid angle it projects. An element of area dA on the ‘surface’ of a planet at a distance D from the Earth, viewed at a local zenith angle θ , subtends a solid angle $d\Omega = dA \cos \theta / D^2$ and since for a planet with radius a , $dA =$

$a^2 \sin \theta d\theta d\phi$ (where θ is the local zenith angle, and ϕ is the azimuth angle), the total planetary spectral irradiance or ‘flux’ seen is

$$F_{\text{plan}} = \frac{a^2}{D^2} \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi/2} R(\lambda, \theta, \phi) \sin \theta \cos \theta d\theta d\phi \quad (2)$$

or

$$F_{\text{plan}} = \frac{a^2}{2D^2} \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi/2} R(\lambda, \theta, \phi) \sin 2\theta d\theta d\phi. \quad (3)$$

By substituting $\mu = \cos \theta$, equation (2) can be simplified to

$$F_{\text{plan}} = \frac{a^2}{D^2} \int_{\phi=0}^{2\pi} \int_{\mu=0}^1 R(\lambda, \mu, \phi) \mu d\mu d\phi. \quad (4)$$

The disc-averaged radiance $\bar{R}(\lambda)$ ($\text{W m}^{-2} \text{sr}^{-1} \mu\text{m}^{-1}$) can be calculated from F_{plan} as

$$\bar{R}(\lambda) = \frac{F_{\text{plan}}}{\Omega_{\text{total}}} = \frac{F_{\text{plan}}}{\frac{\pi a^2}{D^2}}. \quad (5)$$

We can see from equation (3) that the planetary flux (and thus disc-averaged radiance) is most strongly weighted by the radiance emitted at a zenith angle of 45° . Hence to a rough first approximation, the disc-averaged radiance can be estimated as

$$\bar{R}(\lambda) = \frac{1}{2\pi} \int_{\phi=0}^{2\pi} R(\lambda, \theta = 45^\circ, \phi) d\phi, \quad (6)$$

and even more simply, if we assume azimuthal symmetry, as $\bar{R}(\lambda) = R(\lambda, \theta = 45^\circ)$. However, the most accurate estimate of the disc-averaged radiance [especially if $R(\lambda, \mu, \phi)$ varies significantly with position on the disc] comes from the full integration:

$$\bar{R}(\lambda) = \frac{1}{\pi} \int_{\phi=0}^{2\pi} \int_{\mu=0}^1 R(\lambda, \mu, \phi) \mu d\mu d\phi. \quad (7)$$

For our integration scheme we integrated equation (7) with respect to $\mu = \cos \theta$ using a Gauss–Lobatto quadrature scheme, while we used a Trapezium rule integration for the azimuth part of the integration, splitting the circle for each zenith angle into equally spaced points on the nightside from disc edge to terminator (the distance between the points on a zenith angle circle being $\sim R/N_\mu$) and equally spaced points on the dayside from terminator to disc edge (Fig. 1) with the same approximate separation. Our computed disc-averaged radiance was thus calculated as

$$\bar{R}(\lambda) = 2 \sum_i^{N_\mu} \sum_j^{N_\phi} R(\lambda, \mu_i, \phi_{ij}) \mu_i \Delta\mu_i w_{ij}, \quad (8)$$

where μ_i are the Gaussian quadrature points, $\Delta\mu_i$ are the Gaussian weights, ϕ_{ij} are the azimuth angles (which are different for different μ_i), and w_{ij} are the azimuth angle trapezium-rule integration weights. To ensure correct normalization, $\sum_j w_{ij} = 1.0$ and $\sum_i^{N_\mu} \Delta\mu_i = 1.0$ (i.e. $\sum_i^{N_\mu} \mu_i \Delta\mu_i = 0.5$). In addition, we assumed that the exoplanet had a transit with low impact parameter such that northern and southern hemispheres were observed at the same emission angle and we also assumed north/south symmetry such that we only had to integrate over half the disc.

An example of the chosen quadrature points for our scheme using $N_\mu = 2, 3, 4$, or 5 zenith angles is shown in Fig. 1 for a phase of 0.125 (i.e. a phase angle $\zeta = 45^\circ$). Here we can see that as we add more zenith angles, the number of sampled points on the disc is greatly increased. In our scheme if a given zenith angle circle intersects the day–night terminator, one azimuth quadrature point is placed on that terminator and the remaining azimuth points are

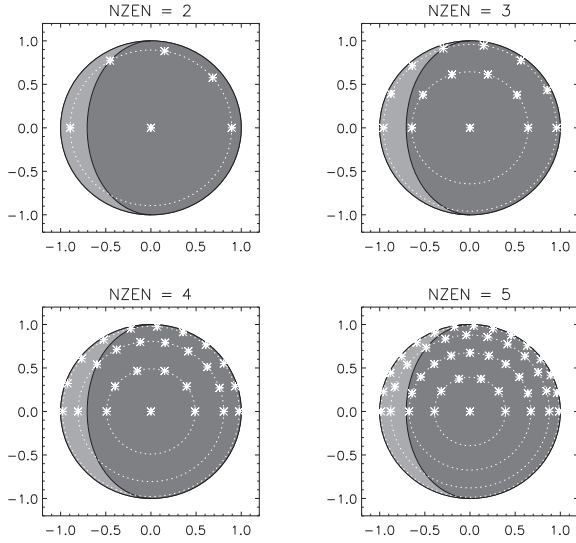


Figure 1. Example of disc-averaging quadrature points chosen for $N_\mu = 2, 3, 4,$ and 5 and phase angle $\zeta = 45^\circ$ (where 0° is nightside and 180° is dayside). The daylight side of the planet is coloured light grey, while the nightside is dark grey. The zenith angles of the chosen Gauss–Lobatto quadrature scheme are shown by the dotted white lines. The azimuth points along each zenith angle line (indicated by the asterisk symbols) are chosen to coincide with the terminator for zenith angles that intersect the terminator and be roughly equally spaced with a fractional disc distance of $1.0/N_\mu$. As we assume north–south symmetry, only points on or above the equator are needed for disc-integration.

chosen to be roughly equally spaced on the day and nightside, as noted above. The reason for doing this is that if we were to observe a planet where the reflected starlight is strong, we would expect that contribution to fall rapidly towards the terminator. By making sure we place a quadrature point at the terminator, we can thus ensure a better estimate of the azimuthal average of the radiance at a given zenith angle. Although in this study we have not included reflected starlight in our models, the disc-integration scheme we have defined means that we will be able to seamlessly and accurately incorporate this in future.

2.1 Computation of contribution of reflected starlight

Assuming a Lambertian reflecting surface, the scattered radiance at any point on the disc is equal to $\alpha F_{\text{star}}/\pi$, where α is the reflectivity and F_{star} is the local stellar irradiance (i.e. W m^{-2}), calculated as $F_{\text{star}} = P \cos \theta / (4\pi d^2)$, where P is the total power of the star, d is the distance of the planet from the star, and θ is the local stellar zenith angle at the point on the disc. At a phase angle of $\zeta = 180^\circ$ the stellar and viewing zenith angles are the same and hence, after equation (4), the planet flux is

$$F_{\text{plan}} = \frac{\alpha P a^2}{4\pi d^2 D^2} \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi/2} \sin \theta \cos^2 \theta \, d\theta \, d\phi \quad (9)$$

or

$$F_{\text{plan}} = \frac{\alpha P a^2}{2d^2 D^2} \int_0^1 \mu^2 \, d\mu = \frac{\alpha P a^2}{6d^2 D^2}. \quad (10)$$

Hence, since $F_{\text{star}} = P/(4\pi D^2)$ we have

$$\frac{F_{\text{plan}}}{F_{\text{star}}} = \alpha \frac{2}{3} \frac{a^2}{d^2} \quad (11)$$

which, at a general phase angle ζ can be approximated to

$$\frac{F_{\text{plan}}}{F_{\text{star}}} = \alpha \frac{2}{3} \frac{a^2}{d^2} \frac{1 - \cos \zeta}{2}. \quad (12)$$

We will return to the likely contribution of reflected starlight to the observed phase curve observations of WASP-43b later.

3 COMPARISON OF 1D AND 2.5D RETRIEVAL SCHEMES

The simplest way of retrieving atmospheric information from phase curve observations, such as those reported by Stevenson et al. (2017), is for each phase angle, ζ , in the set of N_{phase} phase angles to assume that the hemisphere facing Earth has the same temperature and gaseous abundance profiles at all visible parts of the disc. A disc-averaged spectrum can then be computed very simply (in the most simple cases by just computing the spectrum at a viewing zenith angle of 45°) and the temperature/abundance profiles retrieved using schemes such as optimal estimation (e.g. Rodgers 2000; Irwin et al. 2008) or more generally with Bayesian methods, such as nested sampling (e.g. Skilling 2006). Such techniques, which we will here refer to as one-dimensional (i.e. 1D) retrievals, are easy to understand and implement, but suffer from a crucial systematic disadvantage: in order to model the disc-averaged spectra we assume homogeneous conditions across the visible disc, but we then use these retrievals to infer that the atmospheric conditions change markedly with the sub-stellar longitude. Hence, for observations at phase angles sampling mostly the nightside, for example, we assume we cannot see the dayside at all, which is demonstrably incorrect, geometrically, as can be seen in Fig. 1. Here, we aimed to develop a technique where the spatial variation of atmospheric conditions with position on the planet was incorporated directly within the radiative transfer scheme used to retrieve them.

In our new retrieval scheme, for a set of N_{phase} phase curve spectral observations, we fitted all N_{phase} spectra simultaneously with a model that contained the temperature profile and mean gaseous abundances at N_l equally spaced longitudes covering 180°W to 180°E . We then set the value of either temperature or abundance, $x(\Lambda, \Phi)$, at a longitude and latitude (Λ, Φ) to be

$$x(\Lambda, \Phi) = \bar{x} + (x(\Lambda, 0) - \bar{x}) \cos^n \Phi, \quad (13)$$

where $x(\Lambda, 0)$ is the modelled value of the abundance or temperature (at a particular pressure level) at the equator and longitude, Λ , and \bar{x} is mean of the morning and evening terminator values of $x(\Lambda, 0)$, i.e. $\bar{x} = (x(-90^\circ, 0) + x(90^\circ, 0))/2$. For the exponent of $\cos \Phi$ in this equation, n , we used values of $n = 0.25$ (assuming Stefan–Boltzmann-like equilibrium with the stellar irradiation), but also values such as $n = 1$ and $n = 2$ to test for likely latitudinal variation. With this parametrization we can then calculate the disc-averaged radiance over an inhomogenous disc using our numerical disc-integration scheme described above, but more importantly we can retrieve directly the atmospheric conditions on a planet as a function of longitude and height [by retrieving the set of values $x(\Lambda, 0)$] and also have some sensitivity to latitudinal distribution through the assumed $\cos^n \Phi$ dependence. Hence, we call our new technique a 2.5dimensional (i.e. 2.5D) retrieval scheme.

Since we aim to retrieve both temperature and abundances from disc-averaged spectra where we have to perform radiative transfer calculations at a large number of positions across the visible disc for each of the N_{phase} phase angles, our radiative-transfer, or ‘forward model’, is computationally intensive. Hence, for this paper we have used the retrieval method of optimal estimation (e.g. Rodgers

2000). Although less suitable than Bayesian approaches in cases of observations with low SNR, where we need to assess the suitability of a wide range of solutions, the optimal estimation method is very much faster and was thus useful in setting up and testing this approach. For future work, however, we recommend implementing this approach within a Bayesian framework.

The method of optimal estimation is not used as often for exoplanetary work as Bayesian methods (although it has been used several times before, e.g. Lee et al. 2014; Barstow et al. 2017) and some aspects of it may appear counter-intuitive to readers more familiar with techniques such as Monte Carlo Markov chain (MCMC) analysis and nested sampling. The fundamental idea of the approach is that we start with an *a priori* estimate of the atmospheric state and associated *a priori* errors (contained in an error covariance matrix that holds, for example, the expected variances of the temperature profile as a function of height and longitude, and the covariances between these temperatures) and then find a solution that most closely matches the observations without deviating too greatly from the *a priori* estimate by minimizing the cost function ϕ :

$$\phi = (\mathbf{y}_m - \mathbf{y}_n)^T \mathbf{S}_\epsilon^{-1} (\mathbf{y}_m - \mathbf{y}_n) + (\mathbf{x}_n - \mathbf{x}_0)^T \mathbf{S}^{-1} (\mathbf{x}_n - \mathbf{x}_0), \quad (14)$$

where \mathbf{y}_m is the measurement vector, composed of the measured observations, \mathbf{y}_n is the vector of modelled observations calculated from the state vector of model parameters \mathbf{x}_n , \mathbf{S}_ϵ is the measurement covariance matrix, \mathbf{S} is the *a priori* covariance matrix, and \mathbf{x}_0 is the *a priori* state vector. The approach (e.g. Rodgers 2000) was originally developed for analysing Earth observations, where we have very good *a priori* estimates of the expected atmospheric state from climatology and *in situ* observations. However, in its formulation the *a priori* error covariance matrix effectively provides a means of ‘braking’ or ‘smoothing’ the solutions, with small *a priori* error values leading to solutions that differ little from the *a priori* estimates and large values leading to solutions that fit the observed data closely, but are not constrained by the *a priori* estimates at all. The problem with the latter type of solution, known as the ‘exact’ solution, is that it is ‘ill-conditioned’: random noise in the observations can potentially lead to the appearance of large amplitude oscillations in the retrieved vertical profiles. Although such profiles may generate synthetic spectra that match the observations very well, we do not, from atmospheric physics modelling experience, expect the retrieved temperature profiles to vary very rapidly with height. For planetary work, where we do not have a well-known *a priori* based on climatology, models such as our NEMESIS code (Irwin et al. 2008) instead use the *a priori* errors as adjustable ‘tuneable’ parameters to vary the balance of the retrieval between not differing too far from some *a priori* estimate and fitting the observations so closely that the retrieval becomes ill-conditioned. In essence, reducing the *a priori* temperature error in such a scheme reduces the ability of the model to ‘bend’ the fitted temperature profile too greatly from the *a priori* temperature profile to generate spectra that better fit those observed. The diagonal elements of the *a priori* covariance matrix \mathbf{S} are set to the square of the *a priori* error for each state vector element, while to achieve the vertical and horizontal smoothing required by our 2.5D model the off-diagonal elements are set (e.g. Irwin et al. 2008) as

$$S_{ij} = (S_{ii} S_{jj})^{1/2} \exp\left(-\frac{|i-j|}{l}\right), \quad (15)$$

where l is a ‘correlation length’ that adjusts the degree of covariance and thus smoothing. We will return to the specific settings used in our model later.

4 RADIATIVE TRANSFER MODEL

In this paper our aim was to model the phase curve observations of WASP-43b, reported by Stevenson et al. (2017), using our NEMESIS retrieval model (Irwin et al. 2008) in its traditional optimal estimation mode using either the 1D approach or our new 2.5D approach. The measured phase curves of Stevenson et al. (2017) include 15 curves extracted from *HST*/WFC3 observations, binned in equally spaced bins of width 0.035 μm covering the wavelength range 1.1425–1.6325 μm , and two phase curves from *Spitzer*/IRAC from broad channels (see Fig. 6) centred at 3.6 and 4.5 μm . Although Mendonça et al. (2018) and Morello et al. (2019) have independently reprocessed the *Spitzer*/IRAC observations and found significantly different fluxes at 3.6 and 4.5 μm , we have used the measured phase curves of Stevenson et al. (2017) here for ease of comparison of our results. We modelled the spectra using the method of correlated- k (Lacis & Oinas 1991), using k -distribution look-up tables generated from the most recently available line data (e.g. Garland & Irwin 2019), in particular H_2O (Barber et al. 2006), CH_4 (Yurchenko & Tennyson 2014), CO_2 (Tashkun & Perevalov 2011), and CO (Rothman et al. 2010). For high-resolution calculations we used these k -tables at their native resolution of 0.02 μm . However, for our retrievals, where we needed to compute the spectra for multiple iterations, we calculated averaged k -tables for each of the 17 individual ‘channels’. Test calculations proved this to be an excellent approximation, even for broad *Spitzer*/IRAC channels, resulting in residuals far less than the quoted measurement noise of Stevenson et al. (2017). The k -tables were compiled separately for each gas and combined during the radiative-transfer calculation using the random-overlapping-line approximation (Lacis & Oinas 1991). Again, we found this approximation to give residuals insignificant compared with measurement error, as has been found by other authors (e.g. Mollière et al. 2015). Collision-induced absorption of H_2 – H_2 and H_2 – He was modelled using the coefficients of Borysow, Frommhold & Moraldi (1989) and Borysow & Frommhold (1989). Since in this study we assumed thermal emission only, we were able to use NEMESIS in its implicit differentiation mode, where the functional derivatives (i.e. rate of change of computed radiance with atmospheric parameter) are calculated implicitly and thus very rapidly. The planet was modelled to have a mass of 2.052 M_J and radius 1.036 R_J (Gillon et al. 2012), orbiting its star at a distance of 0.015 au. The planet orbits a main-sequence K-type star with a mass of 0.717 M_\odot , radius of 0.667 R_\odot , effective temperature of 4520 ± 120 K and metallicity $[\text{Fe}/\text{H}] = -0.01$ dex (Gillon et al. 2012). These parameters were used to generate a stellar spectrum from the Kurucz ATLAS model atmospheres (Castelli & Kurucz 2004). This was compared with a spectrum extracted for the same star from PHOENIX (Allard, Hauschildt & Schweitzer 2000). While there were some small differences, the effect on the retrieved atmospheric parameters was found to be small.

5 VALIDATION OF RETRIEVAL SCHEMES WITH SYNTHETIC OBSERVATIONS

Before applying our model to the observed WASP-43b phase curve spectra of Stevenson et al. (2017) themselves, we needed to test both the 1D and 2.5D schemes for accuracy and reliability. To do this we used a global thermal structure calculated for WASP-43b from a GCM SPARC/MITgcm (Showman et al. 2009). The specific model for WASP-43b was run based on the set-up of Parmentier et al. (2016) and produces results extremely similar to the published model of Kataria et al. (2015). The model we used was generated

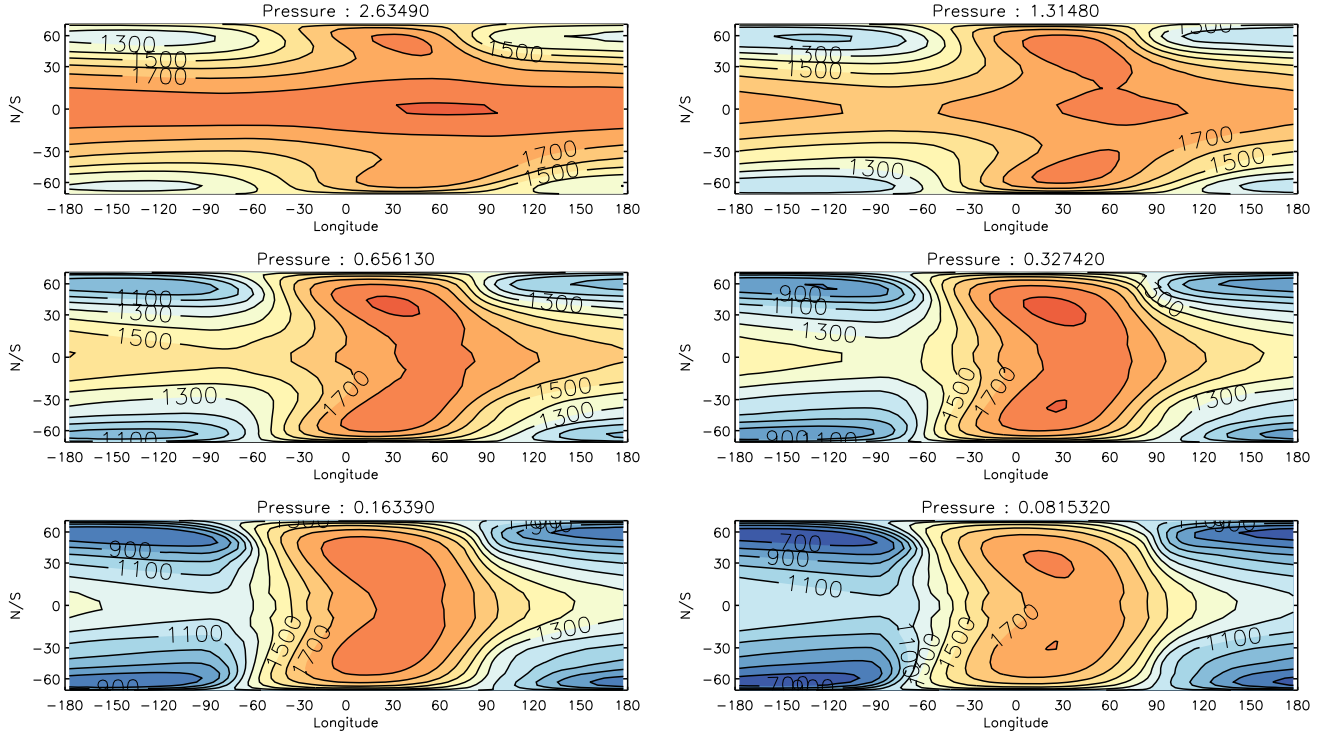


Figure 2. Temperature latitude/longitude contour maps generated by our GCM-based model at six different pressure levels.

assuming no clouds and the gas abundances varied with altitude and location according to equilibrium chemistry leading to rather significant variation with longitude for CH_4 . However, because of horizontal ‘quenching’ (e.g. Cooper & Showman 2006; Agúndez et al. 2014) we do not expect there to be significant longitudinal variation in the abundance of any gas. Hence, to enable a simpler test of our retrieval approaches, the gas abundances for H_2O , CO , CO_2 , and CH_4 were reset at all altitudes and locations to be the latitudinally averaged abundances in the 0.1–1-bar pressure region (using $\cos \Phi$ as the weight, where Φ is the latitude) at the sub-stellar meridian. This gave abundances for H_2O , CO_2 , CO , and CH_4 of 4.8×10^{-4} , 7.4×10^{-8} , 4.6×10^{-4} , and 1.3×10^{-7} , respectively. The modified abundance profiles at all latitudes and longitudes were then used to generate a set of synthetic $HST/Spitzer F_{\text{plan}}/F_{\text{star}}$ spectra that we could then use as our ‘ground-truth’ in our retrieval tests. Although the approach is not fully self-consistent (varying the abundances and thus atmospheric gas opacities will affect the modelled thermal structure also in a self-consistent scheme, see e.g. Drummond et al. 2018; Steinrueck et al. 2019), this is not a concern for the purpose of testing our retrieval schemes since here we retrieve the temperature and abundance profiles independently, with no assumption on the radiative equilibrium state of the atmosphere. Fig. 2 shows the horizontal temperature distribution at six different pressure levels in this model, where we can see that the sub-stellar hotspot is shifted eastwards at all levels, but less so at mid-latitudes and lower pressures. We have plotted these contour maps as a function of longitude and latitude, with the plotted latitudinal distance Φ weighted by $\cos \Phi$ to mimic the fact that we observe WASP-43b from above the equator and thus polar latitudes are foreshortened. Fig. 3 compares the longitudinal variation of the equatorial temperature profile with a latitudinally averaged temperature profile, where the weights are set as $\cos \Phi$. Such an average can be seen to smooth over much




Figure 3. Averaged temperatures and abundances in GCM-based model. Top panel shows the temperature at the equator as a function of longitude and pressure. Bottom panel shows the average temperature profile for each longitude, averaged over all latitudes using $\cos(\text{latitude})$ as a weight.

of the structure seen in the equatorial temperature profiles, more tightly constrains the hotter regions of the planet to the dayside (i.e. between longitudes of -90° and 90°E), and also moves the centre of the hotspot slightly westwards. Finally, Fig. 4 compares the latitudinally averaged temperature profiles for all longitudes together with test retrieved profiles that are described later.

The phase curve observations of WASP-43b of Stevenson et al. (2017) give the $F_{\text{plan}}/F_{\text{star}}$ ratios at fifteen different phases between

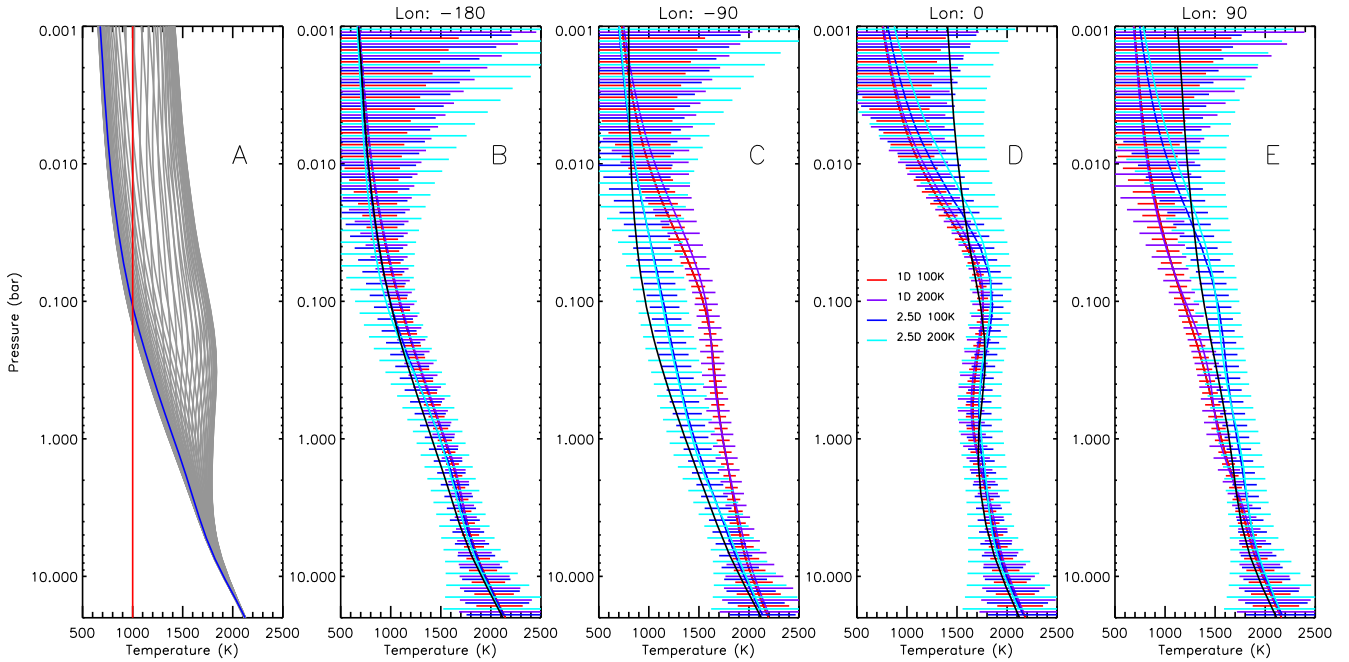


Figure 4. Panel A: Latitudinally averaged temperature profiles for all longitudes from our GCM-based model [using $\cos(\text{latitude})$ as a weight], compared with the two *a priori* temperature profiles used in our study, one set to 1000 K at all levels (red), and the other using the latitudinally-averaged nightside temperature profile at a solar longitude of 180° . The *a priori* errors were set to either ± 50 , ± 100 , or ± 200 K at all levels. Panels B–E: Comparison of temperature profiles retrieved from GCM-based synthetic phase curve spectra at four different central meridian longitudes using our 1D and 2.5D retrieval models: -180° , -90° , 0° , and 90° . The horizontal bars indicate the error due to measurement noise, while the black profile indicates the ‘true’ latitudinally averaged GCM-based temperature profiles at these longitudes. Since the retrieved error, σ_r , tend back to *a priori* error, σ_a , where there is little information, we have here plotted the equivalent raw measurement error $\sigma_m = \sqrt{(1/\sigma_r^2 - 1/\sigma_a^2)}$, which better shows the unconstrained retrieval error.

0.0625 and 0.9375. Hence, we computed the spectra we would expect to see with our model WASP-43b atmosphere at these same phases, assuming thermal emission only. We computed synthetic phase curve spectra using our disc-integration scheme for $N_\mu = 2, 3, 4,$ and 5 zenith angles and found that the computed phase curves converged rapidly as N_μ increased, with χ^2/n (where n is the number of spectral points) reducing by 0.4 going from $N_\mu = 2$ to 3, further falling by 0.02 going from $N_\mu = 3$ to 4 and decreasing by only 0.003 going from $N_\mu = 4$ to 5 (using the measurements errors of Stevenson et al. (2017) for computing χ^2/n). These differences are negligible compared with measurement error for $N_\mu > 2$, but to be sure of numerical accuracy we used $N_\mu = 5$ in all further calculations. As can be seen later in Figs 12 and 13 our model atmosphere provides a reasonable first approximation to the observed phase curves and thus we were satisfied that the GCM-based model provided a reasonably representative data set with which to test our retrieval method. In addition to the disc-averaged spectra, we also generated from the model atmosphere images of the disc of WASP-43b as it would appear at the WFC3/IRAC wavelengths for each phase, i.e. how the planet would appear to an observer who was able to spatially resolve the planet. These images, for the $4.5\text{-}\mu\text{m}$ *Spitzer*/IRAC channel, are shown in Fig. 5. In addition to this primary set of synthetic phase curves, we also generated a secondary set where randomly generated Gaussian noise was added to the spectra with the standard deviation set to the quoted error estimates of Stevenson et al. (2017).

We tested our retrieval model with two *a priori* temperature profiles: (1) one where the *a priori* temperature profile at all longitudes was set to the latitudinally averaged nightside temperature profile of our model at 180°E (using $\cos \Phi$ as a weight); and (2) one where the *a priori* temperature profile was set to a constant temperature

of 1000 K at all altitudes and longitudes. In both cases we found in ‘tuning’ tests that using *a priori* errors of either ± 50 , ± 100 , or ± 200 K at all altitudes, with the vertical correlation length in the *a priori* covariance matrix set to $\Delta H = 1.5$ pressure scale heights, allowed the retrieval model to fit to the observed spectra reasonably well, but were not so large that the retrieval model became ‘ill-conditioned’ (as described in Section 3). These profiles are shown for reference in Fig. 4. It is important to recognize that these assumed *a priori* errors are not hard limits as they might be in models such as MCMC or nested sampling, but instead set the relevant weights in the cost function (equation 14) between fitting the observations well without differing too greatly from the *a priori*. If the data are sufficiently constraining, the fitted temperatures can differ from *a priori* by much more than the *a priori* error. Using either *a priori* temperature profile we found that our retrieval model tended smoothly to consistent temperature solutions at pressures where we have most sensitivity ($\sim 0.1\text{--}1$ bar), and tended smoothly back to *a priori* values at pressures outside this range. However, although retrievals with both *a priori* profiles were consistent, we found the former profile resulted in better behaved retrievals when applied to the spectra measured at all phases since the temperature profile remained close to a physically plausible solution on the nightside, where the SNR is low, and led to more realistic-looking retrievals on the dayside with temperatures tending back to the *a priori* of a smooth increase with pressure at pressures > 1 bar, consistent with expectations. Hence, this was our preferred *a priori* profile in the results we describe.

Our atmosphere had an assumed He/H_2 mole fraction ratio of 0.1765 and contained four additional gases: H_2O , CO_2 , CO , and CH_4 , whose abundances were fitted or assumed. The gas abundances




Figure 5. Simulated appearance of disc of WASP-43b at $4.5 \mu\text{m}$ at different phases for nine different cases: (A) the GCM-based model atmosphere; (B) retrieved appearance from GCM-based synthetic phase curves with 1D model; (C) retrieved appearance from GCM-based synthetic phase curves with 2.5D model, assuming $(\cos \Phi)^{0.25}$ latitude dependence; (D) retrieved appearance from GCM-based synthetic phase curves with 2.5D model, assuming $\cos \Phi$ latitude dependence; (E) retrieved appearance from real observed phase curves with 1D model; (F) retrieved appearance from real observed phase curves with 2.5D model, assuming $(\cos \Phi)^{0.25}$ latitude dependence; (G) retrieved appearance from real observed phase curves with 2.5D model, assuming $\cos \Phi$ latitude dependence; (H) retrieved appearance from real observed phase curves with 2.5D model, assuming $(\cos \Phi)^{0.25}$ latitude dependence and forcing nightside temperatures to 500 K at all altitudes; and (I) retrieved appearance from real observed phase curves with 2.5D model, assuming $\cos \Phi$ latitude dependence and forcing nightside temperatures to 500 K at all altitudes. The phase angle increases from left to right in all cases and is indicated at the top of the figure.

were assumed to have a constant mole fraction (or volume mixing ratio, VMR) with height and when allowing abundance to vary we retrieved logarithmic scaling factors, i.e. $\text{VMR} = \text{VMR}_{\text{apriori}} \exp(x)$, where x is the relevant state vector element, to ensure the mole fractions remained positive during the retrievals. At each iteration the H_2 and He abundances were adjusted (keeping He/ H_2 fixed) to ensure the sum of mole fractions added to unity. The mean molecular weight of the atmosphere was then recomputed in order to calculate the scale height accurately and thus determine the altitudes of each level (above the 20 bar level) assuming hydrostatic equilibrium.

For our first test, we took our synthetic disc-averaged spectrum (with no additional noise) for a phase of 0.5 (i.e. phase angle $\zeta = 180^\circ$, or secondary eclipse) and performed a 1D retrieval assuming thermal emission only. Since we assume the same atmospheric conditions at all locations on the disc in our 1D scheme, the azimuthal part of the integration in equation (8) was not required and so we just integrated over the $N_\mu = 5$ zenith angles. We set our *a priori* scaling factors to $x = 0 \pm 1$, with *a priori* mole fractions set to the averaged dayside abundances of the GCM simulation, namely 4.8×10^{-4} , 7.4×10^{-8} , 4.6×10^{-4} , and 1.3×10^{-7} , for H_2O , CO_2 , CO, and CH_4 , respectively. We then split the atmosphere into 50 layers from 20 to 10^{-5} bar [split equally with respect to $\log(\text{pressure})$] and retrieved for temperature profile and gas abundances. Our *a priori* temperature profile here was set to the latitudinally averaged nightside temperature profile of our

GCM at 180°E , with an error of ± 100 K. Fig. 6 shows the fitted spectral flux ratio spectrum in the seventeen *HST/Spitzer* ‘channels’ compared with a spectrum recalculated from the fitted atmospheric state at reasonably high resolution ($\Delta\lambda = 0.02 \mu\text{m}$) over the whole spectral range, showing the channel response functions of the *HST/WFC3* and *Spitzer/IRAC* observations considered. Panel A of Fig. 6 also shows the highest likely reflected solar component, which was calculated from equation (11) assuming a reflectivity $\alpha = 1$. Comparing this with the observed fluxes we find that at the longer wavelengths the reflected component is likely a relatively small part of the total observed flux. However, it seems that reflected sunlight could potentially be a significant component at the shorter wavelengths. Stevenson et al. (2014) deduce a Bond albedo of $0.18^{+0.07}_{-0.12}$ from the *HST/WFC3* observations and more recently, Keating & Cowan (2017) have reanalysed these data to deduce a near-infrared geometric albedo of 0.24 ± 0.01 and note that reflected light may be significant at the WFC3 wavelengths. However, given the fact that the likely estimated component of reflected starlight at the *HST/WFC3* wavelengths is smaller than the estimated error bars on these observations, we have chosen to ignore any reflected component in this analysis, as was also done by Stevenson et al. (2017).

Fig. 6 also shows the transmission weighting functions computed by our retrieval model (i.e. dT/dz , where T is the transmission to space and z is the altitude), showing that our sensitivity to temperature variations peaks between 10 and 0.05 bar. We can

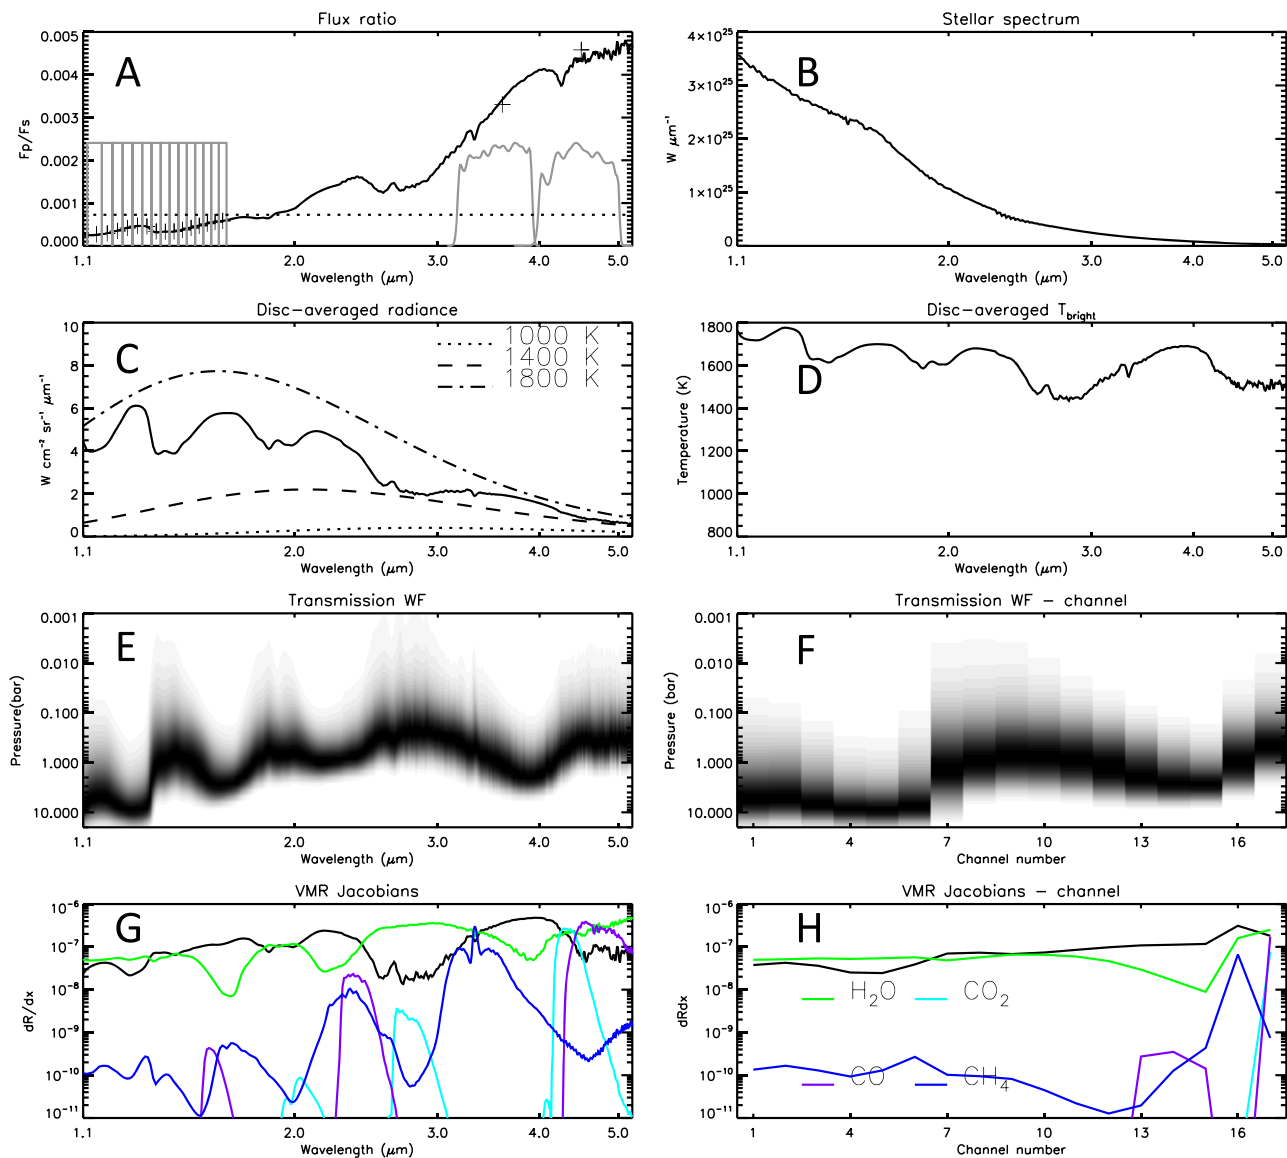


Figure 6. 1D model fit to synthetic data generated from our GCM-based model for a phase of 0.5 (i.e. phase angle $\zeta = 180^\circ$). (A) Modelled $F_{\text{plan}}/F_{\text{star}}$ spectrum at a resolution of $0.02 \mu\text{m}$ using our $N_\mu = 5$ quadrature scheme. The grey lines indicate the WFC3/*Spitzer* filter functions, while the horizontal dotted line shows the maximum possible estimated reflected stellar contribution (assuming unit reflectivity, although since the albedo of WASP-43b is estimated to be ~ 0.24 the likely reflected component is considerably less). (B) Assumed stellar power spectrum. (C) Modelled disc-averaged radiance spectrum, compared with Planck functions for different temperatures. (D) Modelled disc-averaged spectrum plotted as equivalent brightness temperature. (E) Modelled transmission weighting functions (normalized to unity at each wavelength) at high resolution (i.e. $0.02 \mu\text{m}$) plotted against wavelength. (F) Normalized modelled transmission weighting functions integrated over filter function and plotted against channel number. (G) Jacobians (plotted against wavelength) for gas abundances (coloured lines – multiplied by 1000) compared with Jacobian for temperature at a $p \sim 1$ bar (black line). (H) As Panel G, but channel-averaged values plotted against channel number; a key for the gas colour lines is also included.

thus use these observations to retrieve the temperature between these levels, but outside this pressure region the temperatures will tend back to *a priori* in our optimal estimation scheme. Fig. 6 also compares the Jacobians for gas abundance (i.e. the rate of change of measured signal with respect to change in abundance) with that of temperature near 1 bar. It can be seen that CO_2 absorption is only detectable in the $4.5\text{-}\mu\text{m}$ IRAC channel, while CO absorption is significant in the $4.5\text{-}\mu\text{m}$ IRAC channel and the three longer wavelength WFC3 channels. However, since the observation in the $4.5\text{-}\mu\text{m}$ IRAC channel has much lower measurement error than the WFC3 channels, we can see that there will likely be a

degeneracy in the retrieval between the temperature at ~ 0.1 bar and the CO/ CO_2 abundance. The Jacobian for CH_4 absorption is more widely distributed with wavelength, but peaks in the $3.6\text{-}\mu\text{m}$ IRAC channel, which is again measured to a higher precision than the WFC3 channels. Hence, there will likely be a degeneracy in the retrieval between the temperature at just over 0.1 bar and the CH_4 abundance. It can be seen that the Jacobian for water vapour abundance is almost perfectly anticorrelated with the Jacobian for 1-bar-temperature. This is easily understood since water is the main absorber and has features across the $1\text{--}5 \mu\text{m}$ spectral range. By increasing water abundance we shift the weighting functions

to lower pressures where the temperature is lower, producing a change in modelled spectrum almost indistinguishable from simply lowering the temperature. Within our optimal estimation framework we see the same high degree of anticorrelation in the retrieved covariance matrix and at first glance it might be concluded that it is impossible to distinguish between water abundance and temperature. Certainly, where optimal estimation has usually been used before it is common practice to assume the abundance of one gas to be well known such that it can be used as a ‘thermometer gas’ and its absorption features analysed to retrieve atmospheric temperature. Since we do not, *a priori*, know the water vapour abundance, how can meaningful information on temperature and abundance be retrieved? Fortunately, while the anticorrelation is strong at this pressure level for the assumed GCM-based model abundances the anticorrelation becomes less strong at higher and lower pressures due to effects such as the varying widths of the H₂O lines caused by pressure-broadening and the increasing importance of H₂–H₂ and H₂–He collision-induced absorption at depth. Hence, instead of retrieving a single value it is possible to infer a range of abundances for which we achieve similarly good fits to the observations. Analysing the secondary eclipse and primary transit observations of WASP-43b, Kreidberg et al. (2014) report the water vapour mole fraction to be in the range $(3.1–44.0) \times 10^{-4}$ from secondary eclipse observations and $(0.33–14) \times 10^{-4}$ from primary transit analysis, arriving at a joint constraint of $(2.4–21.0) \times 10^{-4}$. This estimate was further refined by Stevenson et al. (2017) who estimate the water vapour abundance to be $(0.25–1.1) \times 10^{-4}$. The approach of optimal estimation is not best designed for exploring a wide parameter space such as this, but with care it can be used to explore the range of solutions consistent with the observations, in this particular case the synthetic observations generated from a GCM-based model atmosphere and for which we know the ‘ground truth’. We can see from Fig. 6 that the signatures of CO₂ and CO are difficult to distinguish between with the WFC3/IRAC data, as was also noted by Stevenson et al. (2017). Hence, we first fixed the mole fraction of CO₂ to 10^{-7} and tried to estimate the abundance of CO only. We then set up a grid of trial CO and H₂O mole fractions varying between 10^{-5} and 5×10^{-3} and then retrieved for each grid combination the vertical profile of temperature and mean methane abundance. The *a priori* methane mole fraction was set to 10^{-5} with the log scaling factor (natural logs) set to 0 ± 10 to allow the model to fit the CH₄ abundance as freely as possible.¹

Fig. 7 shows the atmospheric parameters fitted to the synthetic phase curve data (without additional random Gaussian-generated noise) at all phases with our 1D model, where the *a priori* temperature profile was set to the nightside GCM profile with an error of ± 200 K. Here, we took each of the $N_{\text{phase}} = 15$ synthetic spectra in turn and for each assumed CO and H₂O mole fraction combination fitted the vertical temperature profile and mean CH₄ abundance that gave the closest disc-averaged $F_{\text{plan}}/F_{\text{star}}$ ratio (using $N_{\mu} = 5$ zenith angles and assuming azimuthal symmetry). We then took all N_{phase} individually retrieved temperature profiles and plotted them as a function of the central meridian longitude of WASP-43b (assuming 0°E is the sub-stellar point). The top two rows of Fig. 7 shows how χ^2 varies with fixed CO and H₂O abundance for our models (with the second row also showing the fitted CH₄ abundances at 0°E), while the remaining rows plot the longitudinal

variation of temperature for a single choice of the H₂O and CO abundance, indicated by the green asterisk in the top panel of each column. Looking at the χ^2 contour plot in Fig. 7 it can be seen that we achieve an excellent fit to the synthetic observations ($\chi^2/m < 0.1$) and that the water vapour abundance is reasonably well constrained and consistent with the ‘true’ GCM-based value. However, the CO abundance is much less well constrained. The temperature profiles shown in the first column of Fig. 7 are for CO and H₂O values close to ‘true’, but the remaining two columns explore how the retrieved temperatures vary with the assumed CO and H₂O abundances, by first increasing the assumed H₂O abundance by factor of 10, and then increasing the assumed CO abundance by the same factor. We find that the retrieved temperatures vary with the assumed H₂O abundance as expected, with increased H₂O abundance leading to the higher temperature regions being retrieved at lower pressure levels. However, we also find a significant variation in retrieved temperature with CO abundance, with the sub-stellar hotspot moving to higher altitude and forming a slight inversion at ~ 0.2 bar for large CO values; this is consistent with our observation earlier that the CO abundance and temperature at ~ 0.1 bar are degenerate. All the temperature profiles for the entire grid of CO and H₂O abundances vary moderately smoothly with longitude and altitude, indicating reasonable constraint and while they look different from the ‘true’ temperature profiles, it would be difficult to recognize them to be in error if we did not know what the ‘true’ values were. Hence, it can be seen that although the WFC3/IRAC data, measured with the spectral coverage and error values reported, allow some degree of constraint on the water abundance, they do not allow very strong constraint on the CO abundance due to the degeneracy with temperature. For CH₄, we find that the retrieved abundances are approximately 10 times larger than ‘true’, with retrieved error bars consistent with the spread of retrieved values. To check that the methane retrievals were not being biased by our chosen *a priori* methane abundance of 10^{-5} we repeated our retrievals reducing this *a priori* to 10^{-7} and found essentially identical results. The high CH₄ abundances retrieved with the 1D model are consistent with one of the conclusions of Feng et al. (2016), that the 1D approach ‘overestimates both its (i.e. CH₄) abundance and the precision’. As for temperature, we can see that the retrieved temperatures are hotter near the sub-stellar point, but offset eastwards as is the case in the GCM simulations used to generate the synthetic spectra. However, it can be noticed that the hotter temperatures extend some way past the day/night terminator at 90°E and 90°W, unlike the ‘true’ GCM temperatures. This can be seen most clearly in the bottom left panel of Fig. 7, where the retrieved–true temperatures at 0.1–1 bar show significantly negative differences of ~ -200 K near the sub-stellar longitude, but significantly positive differences of $\sim +200$ K at 90°W and 120°E. This can also be seen in Fig. 4, where we compare the ‘true’ and retrieved temperatures in detail at 180°W, 90°W, 0°E, and 90°E. Here, we can see that the 1D model overestimates the temperature at 90°W by almost 500 K. This behaviour is expected since for 1D retrievals we use the same temperature profile at all points on the disc and to model the slow variation of observed $F_{\text{plan}}/F_{\text{star}}$ with phase angle ζ , we need the temperature to vary more slowly with longitude. This can also be seen in modelled appearance of the planet in the 4.5- μm channel in Fig. 5.

Although we computed these spectra assuming our $N_{\mu} = 5$ Gaussian integration scheme, we checked the necessity of using such a scheme with 1D retrievals for phase curves measured with currently achievable SNR. We did this by modelling the spectra with a single zenith angle of 45° and found that we achieved a similarly close fit to the observations (average differences in

¹Remember that in optimal estimation, the *a priori* errors are weights, not hard limits. Hence, it is possible to retrieve variations larger than the *a priori* errors if the data are sufficiently constraining.

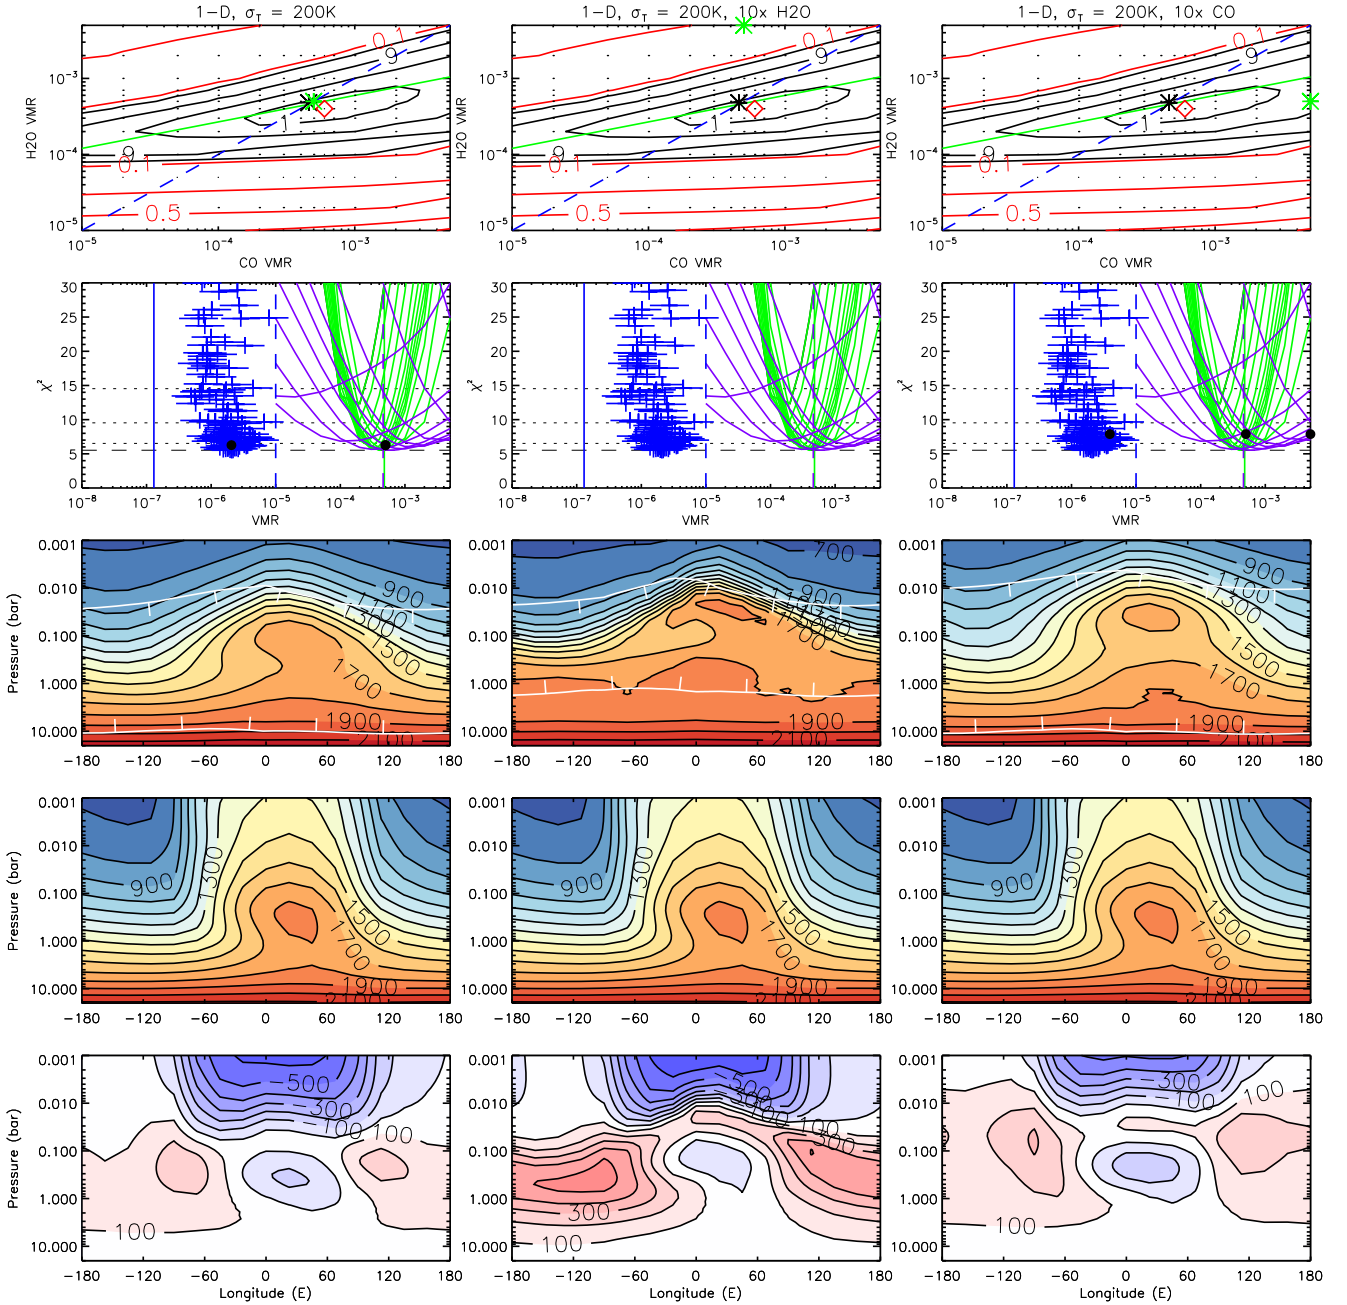


Figure 7. Retrieval of GCM-based synthetic phase curve data (excluding randomly generated synthetic noise) using our 1D model. Figure shows three sets of retrievals: Column (1) 1D retrieval with the CO and H₂O abundance for the temperature plots set to the GCM-based ‘true’ values of $\sim 5 \times 10^{-4}$; Column (2) 1D retrieval with H₂O abundance for the temperature plots increased to 5×10^{-3} ; and Column (3) 1D retrieval with CO abundance for the temperature plots increased to 5×10^{-3} . The top row shows contour plots of the total χ^2 , with contours set at $\min(\chi^2)$, $\min(\chi^2) + 1$, $\min(\chi^2) + 4$, $\min(\chi^2) + 9$, $\min(\chi^2) + 16$. Red contours indicate χ^2/n . For 3σ detection we should be within $\min(\chi^2) + 9$. The blue dashed lines in this row indicates where the abundance of CO is the same as that of H₂O, while the green line indicates the approximate principal axis of the best-fitting solutions for CO and H₂O. In each plot the black asterisk marks the ‘true’ mole fractions, the red diamond indicates the best-fitting mole fractions and the green asterisk indicates the abundances chosen for the temperature plots in the remaining rows. Row 2 shows plots of χ^2 versus H₂O (green) and CO (purple) mole fraction. Also indicated are the retrieved abundances of CH₄ (blue), with the retrieved errors indicated by the horizontal bars. The ‘true’ GCM-based abundances are indicated by the solid vertical lines using the same colour scheme (N.B., the *a priori* CO and H₂O values are effectively indistinguishable for this overall range of abundances at $\sim 5 \times 10^{-4}$, so the H₂O abundance is plotted as solid green and the CO abundance overplotted as dashed purple). The dashed vertical blue line is the *a priori* CH₄ abundance. The horizontal lines in this row show the minimum χ^2 value (dashed lines) and $\min(\chi^2) + 1$, $\min(\chi^2) + 4$, $\min(\chi^2) + 9$ (dotted lines). Finally, the black filled circles in this row indicate the abundances chosen for the temperature plots in the remaining rows. Row 3: Retrieved temperatures as a function of longitude and pressure. The white contour indicates the region where the retrieval is well constrained by the observations and has higher values of the ‘Improvement factor’, I , defined as $I = 1 - \sigma_r/\sigma_a$, where σ_a is the *a priori* error and σ_r is the retrieval error. Row 4: $\cos \Phi$ - weighted latitudinally averaged GCM-based temperature profiles. Row 5: Retrieved temperature minus latitudinally averaged GCM-based temperature (red-blue scale, white = zero difference).

χ^2/n , averaged over all 15 phases, of <0.05) and very similar retrieved properties. This demonstrates that for cases where we assume the same atmospheric conditions at all observable parts of the disc, setting $R(\lambda) = R(\lambda, \theta = 45^\circ)$ is a good approximation. In fact, we found that setting $R(\lambda) = R(\lambda, \theta = 0^\circ)$ we were also able to achieve similarly close fits to the synthetic observations, and similar retrieved atmospheric properties as compared with the full numerical integration scheme. However, in this case, because the transmission weighting functions peak at deeper pressures the retrieved vertical temperature distribution is also shifted to deeper pressures.

While our 1D approach gives an excellent fit to the synthetic observations and is broadly consistent with GCM-based model temperatures and abundances it is oversimplistic to assume that the atmospheric conditions are the same at all points on the disc as can clearly be seen in Fig. 5. We hence applied our 2.5D retrieval scheme (with different assumed values of the latitudinal coefficient n), where we fitted all $N_{\text{phase}} = 15$ spectra simultaneously with a model where the state vector contained the temperature at 20 levels equally spaced in log pressure between 20 and 0.001 bar at 16 different longitudes covering 0°E to 337.5°E in steps of 22.5° . For the temperature profile at each longitude we again set a correlation length of 1.5 scale heights. For horizontal smoothing we applied a correlation length of 22.5° between the temperatures at different longitudes. For the CH_4 abundance, since we assumed horizontal ‘quenching’ we fitted a single mole fraction for all longitudes. Fig. 8 compares our 2.5D model fits to the synthetic phase curve data with our 1D model, but this time for the synthetic phase curve data where additional random Gaussian-generated noise has been added. For our 2.5D model we assumed a latitude temperature dependence of $(\cos \Phi)^{0.25}$ (i.e. $n = 0.25$) and for both 1D and 2.5D retrievals we set the *a priori* temperature profile to the nightside GCM profile with errors of either ± 100 or ± 200 K. It can be seen from Fig. 8 that we achieve similarly close fits to the synthetic phase curve data with either the 1D or 2.5D models, but with the 2.5D scheme our fitted temperatures correspond much more closely with the weighted-mean longitudinal temperatures of the GCM-based model, with the hot region more tightly constrained to the dayside than the 1D cases and so smaller differences seen the retrieved–true contour plots. To explore this more quantitatively, our fitted temperature profiles for longitudes 180°W , 90°W , 0°E , and 90°E are compared with the respective latitudinally averaged GCM-based model profiles in Fig. 4. Here we see that the 2.5D-retrieved temperatures are closer to the GCM-based model profiles than the 1D retrieved profiles and that the retrieved temperature profiles differ from *a priori* by much more than the *a priori* errors, underlining the point made earlier that these *a priori* errors are weights, rather than hard limits. It can also be seen that the differences in the temperature profiles derived by the 1D and 2.5D retrieval models are at some pressure levels larger than the estimated retrieval error, showing these differences to be statistically significant.

With regards to composition, for the 2.5D cases it can be seen in Fig. 8 that the water vapour abundance is again reasonably well constrained and consistent with the ‘true’ GCM-based value. However, the CO abundance is again more poorly constrained. In contrast with the 1D cases, the 2.5D cases recover slightly lower values of CH_4 with larger error bars, more consistent with the assumed GCM-based model values,² and again consistent with

Feng et al. (2016), although the difference compared with the 1D approach is only of 1σ significance at best with the assumed *HST/Spitzer* errors. While we have only shown 2.5D retrievals with $n = 0.25$ [i.e. $(\cos \Phi)^{0.25}$ latitude dependence] we found little difference between the temperatures and abundances retrieved by the 2.5D scheme for latitude dependences $n = 1$ and $n = 2$. However, we did find that the minimum χ^2 values increased monotonically as n increased and that the lowest values were found for the $(\cos \Phi)^{0.25}$ case, indicating that this is our preferred solution. Fig. 5 shows the fitted appearance of the disc of WASP-43b at $4.5 \mu\text{m}$ from our 1D and 2.5D retrievals for the case $\sigma_T = \pm 100$ K (for $n = 0.25$ and $n = 1$ for the 2.5D fits) where it can be seen that the 2.5D approach generates much more realistic-looking disc images, which correspond better with the ‘true’ simulated appearance of this planet. It can also be seen that the $n = 0.25$ case better approximates the north/south distribution of the ‘true’ appearance. We then explored how the retrieved temperatures varied with assumed CO and H_2O abundance for the 2.5D approach. As for the 1D approach we find that increased H_2O abundances raises the vertical location of the higher temperature regions to lower pressure levels. We also again find a significant variation in retrieved temperature with CO abundance, with the sub-stellar hotspot moving to higher altitude and forming an inversion at ~ 0.2 bar for large CO values. However, what is particularly clear from Fig. 8 is that the best fit CO and H_2O values lie along a rough diagonal indicated by the green line in Figs 7 and 8, and that the position along that line of the best fit depends on the assumed *a priori* temperature error and thus the degree of vertical smoothing imposed by the retrieval on the vertical temperature profile. Increasing the *a priori* temperature error leads not only to reduced vertical smoothing and the retrieval of temperature profiles that vary more rapidly with altitude, but also leads to the retrieval of higher H_2O and CO abundances. While this does not lead to very large changes in the best-fitting H_2O abundance, it leads to large changes in the best-fitting CO abundance. This suggests that trying to estimate the metallicity of WASP-43b from the CO absorption for the WFC3/IRAC data is likely to be prone to significant systematic error since the inferred value depends greatly on the assumed/fitted temperature profile. In simple terms this correlation arises from the fact that, as we can see from Fig. 6, the most sensitive wavelength for CO is the $4.5\text{-}\mu\text{m}$ *Spitzer* channel, while the temperature solution is driven by fitting to all the channels simultaneously. Hence, the optimal estimation approach finds that that when fitting to the observations with a temperature profile that is constrained to not to vary too greatly with altitude, a lower CO abundance is required in order to allow the $4.5\text{-}\mu\text{m}$ channel to see to the deeper, warmer levels, where the radiances are higher. If the temperature profile is able to vary more rapidly, higher temperatures are retrieved at the lower pressures and hence higher CO values are needed to push the weighting function for the $4.5\text{-}\mu\text{m}$ channel up to pressure levels where the radiances are consistent with the observations. A similar effect is seen for the CH_4 retrievals, for which the most sensitive wavelength is the $3.6\text{-}\mu\text{m}$ *Spitzer/IRAC* channel. It should be noted that this correlation between CO and H_2O abundance has been seen in previous retrieval analyses using Bayesian approaches (e.g. Kreidberg et al. 2014; Line et al. 2016).

For both 1D and 2.5D retrievals, we see that the difference between retrieved and ‘true’ is worst for pressures less than 0.05 bar on the dayside. This is because the retrievals tend towards the *a priori* estimates at pressures where there is little information. Figs 7 and 8 also plot contours of the temperature ‘Improvement factor’, I , for the temperature retrievals. This is defined as $I = 1 -$

²In an additional test, we found that when fitting to the synthetic phase curve data without additional noise, the retrieved CH_4 abundances were slightly more consistent with ‘true’.

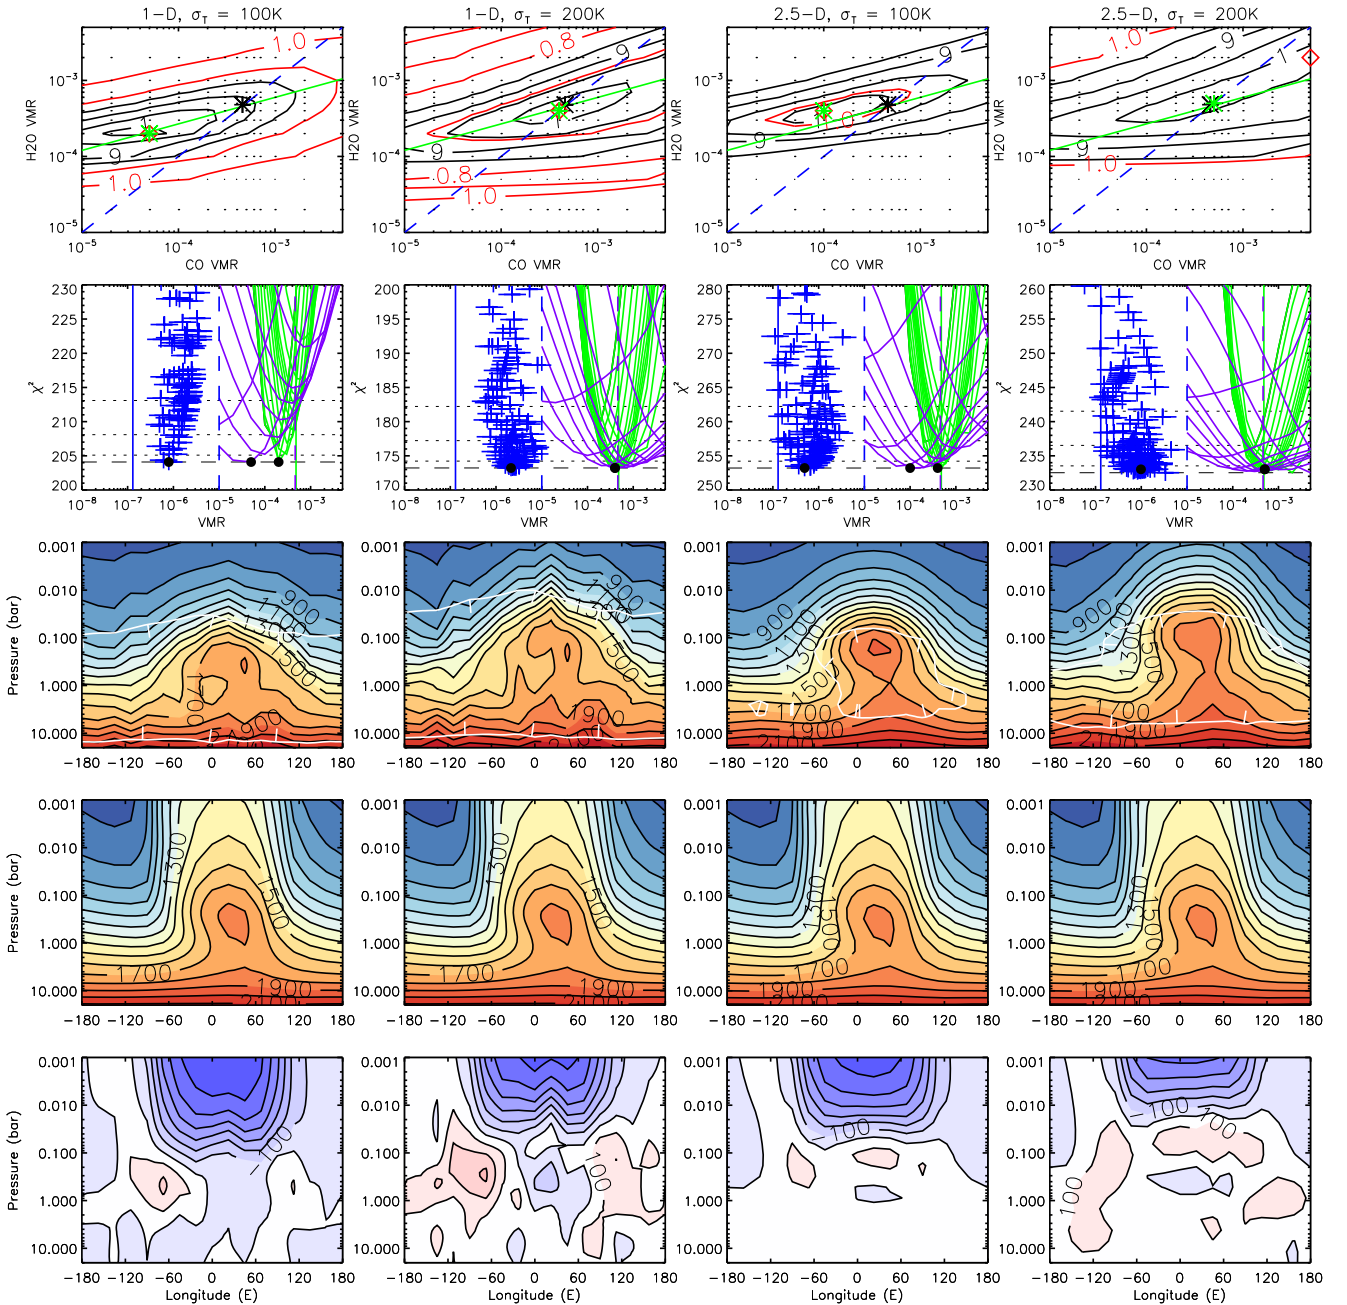


Figure 8. Retrieval of GCM-based synthetic phase curve data (including randomly generated synthetic noise) using both our 1D model and 2.5D models. The rows are the same as described in Fig. 7. Four sets of retrievals are shown: Column (1) 1D retrieval with *a priori* temperature error ± 100 K; Column (2) 1D retrieval with *a priori* temperature error ± 200 K; Column (3) 2.5D retrieval ($n = 0.25$) with *a priori* temperature error ± 100 K; and Column (4) 2.5D retrieval ($n = 0.25$) with *a priori* temperature error ± 200 K. The rows are as described for Fig. 7. Here the H_2O and CO abundances chosen for the temperature plots are different from the ground truth and are indicated by the green asterisks in the plots on the top row. N.B., the reduced area of significant improvement for the 2.5D retrieval with temperature error ± 100 K compared with ± 200 K is easily explained: by reducing the *a priori* temperature error, the retrieval cannot vary as far from *a priori* on the nightside since the data are less precise. Hence, the improvement factor contours surround only the dayside regions, where the data have a higher SNR.

σ_r/σ_a , where σ_a is the *a priori* error and σ_r is the retrieval error. Higher values of I indicate where the retrieval is more constrained by the measurements and thus where the retrieved temperatures are most reliable. We can see here that this region corresponds to the ~ 0.05 –10 bar region for the 1D approach and ~ 0.2 –10 bar region for the 2.5D approach (with diminished improvement factors

on the nightside for the 2.5D retrievals since the temperatures are lower here leading to reduced SNR and solutions that can deviate less from the *a priori* estimate). Outside this pressure region, the temperature profiles revert smoothly to *a priori* and thus differences between ‘true’ and retrieved should not cause concern. We repeated our retrievals using our other *a priori* temperature

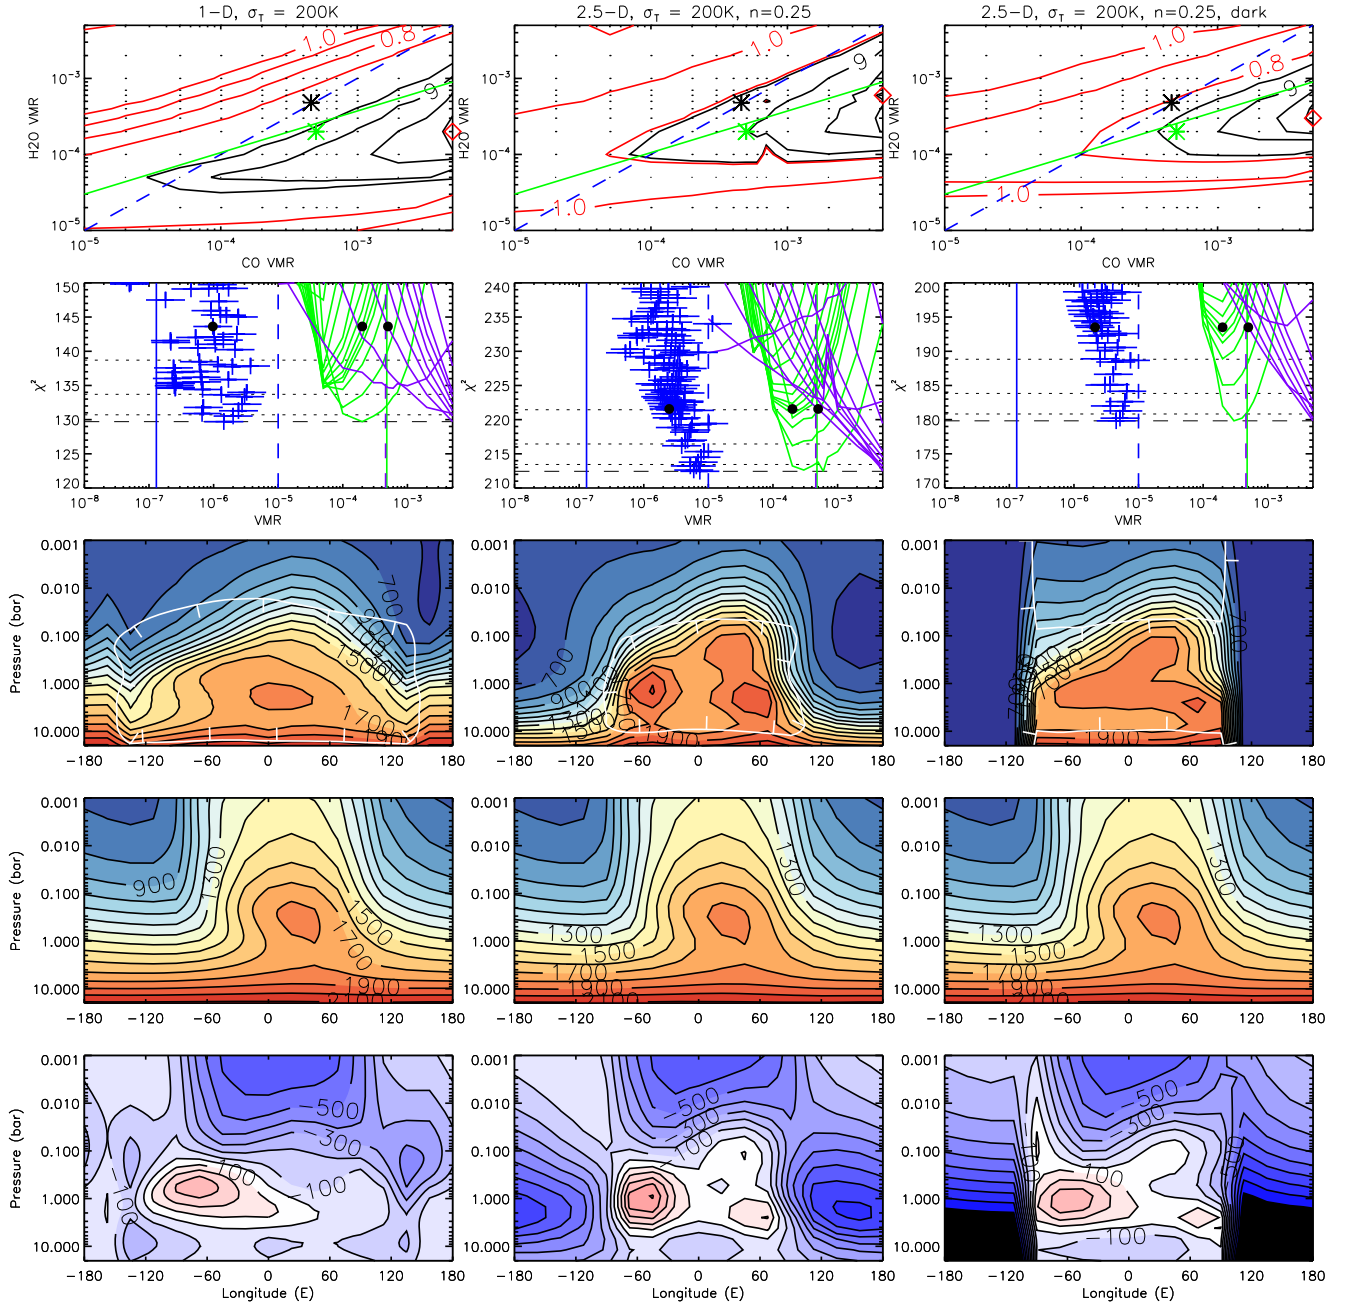


Figure 9. Application of our 1D and 2.5D models ($n = 0.25$) to the measured phase curves of Stevenson et al. (2017), assuming an *a priori* temperature error of ± 200 K. The rows are as described for Fig. 7. For the temperature maps, we have chosen H₂O and CO abundances (indicated by green asterisks in row 1) that balance the requirement to have the lowest χ^2 values with the expectation to have roughly equal abundances (2×10^{-4} and 5×10^{-4} , respectively for all cases).

profile ($T = 1000$ K at all altitudes) and found consistent results within the pressure regions where the ‘Improvement factor’ was significant.

6 APPLICATION OF RETRIEVAL SCHEMES TO REAL WFC3+IRAC OBSERVATIONS

Having demonstrated the efficacy and reliability of our modelling approach on a set of simulated observed data, where we knew what the ‘ground truth’ was, we then applied both our 1D and 2.5D models (using identical set-ups to the validation tests) to the

observed phase curves of Stevenson et al. (2017). Our results are shown in Figs 9, 10, and 11 for *a priori* temperature errors of ± 200 , ± 100 , and ± 50 K, respectively.

Comparing the 1D and 2.5D retrievals it can be seen that the nightside of WASP-43b appears much colder than the GCM-based simulation and that the 2.5D approach finds that the hot region of the planet is much more tightly constrained to the dayside of the planet than the 1D approach. For the dayside we retrieve the highest temperatures at ~ 200 mb, significantly shifted to the east, but at lower altitudes the hotter regions appear to be centred more to the west, mirroring the GCM simulations, where the hotter regions are

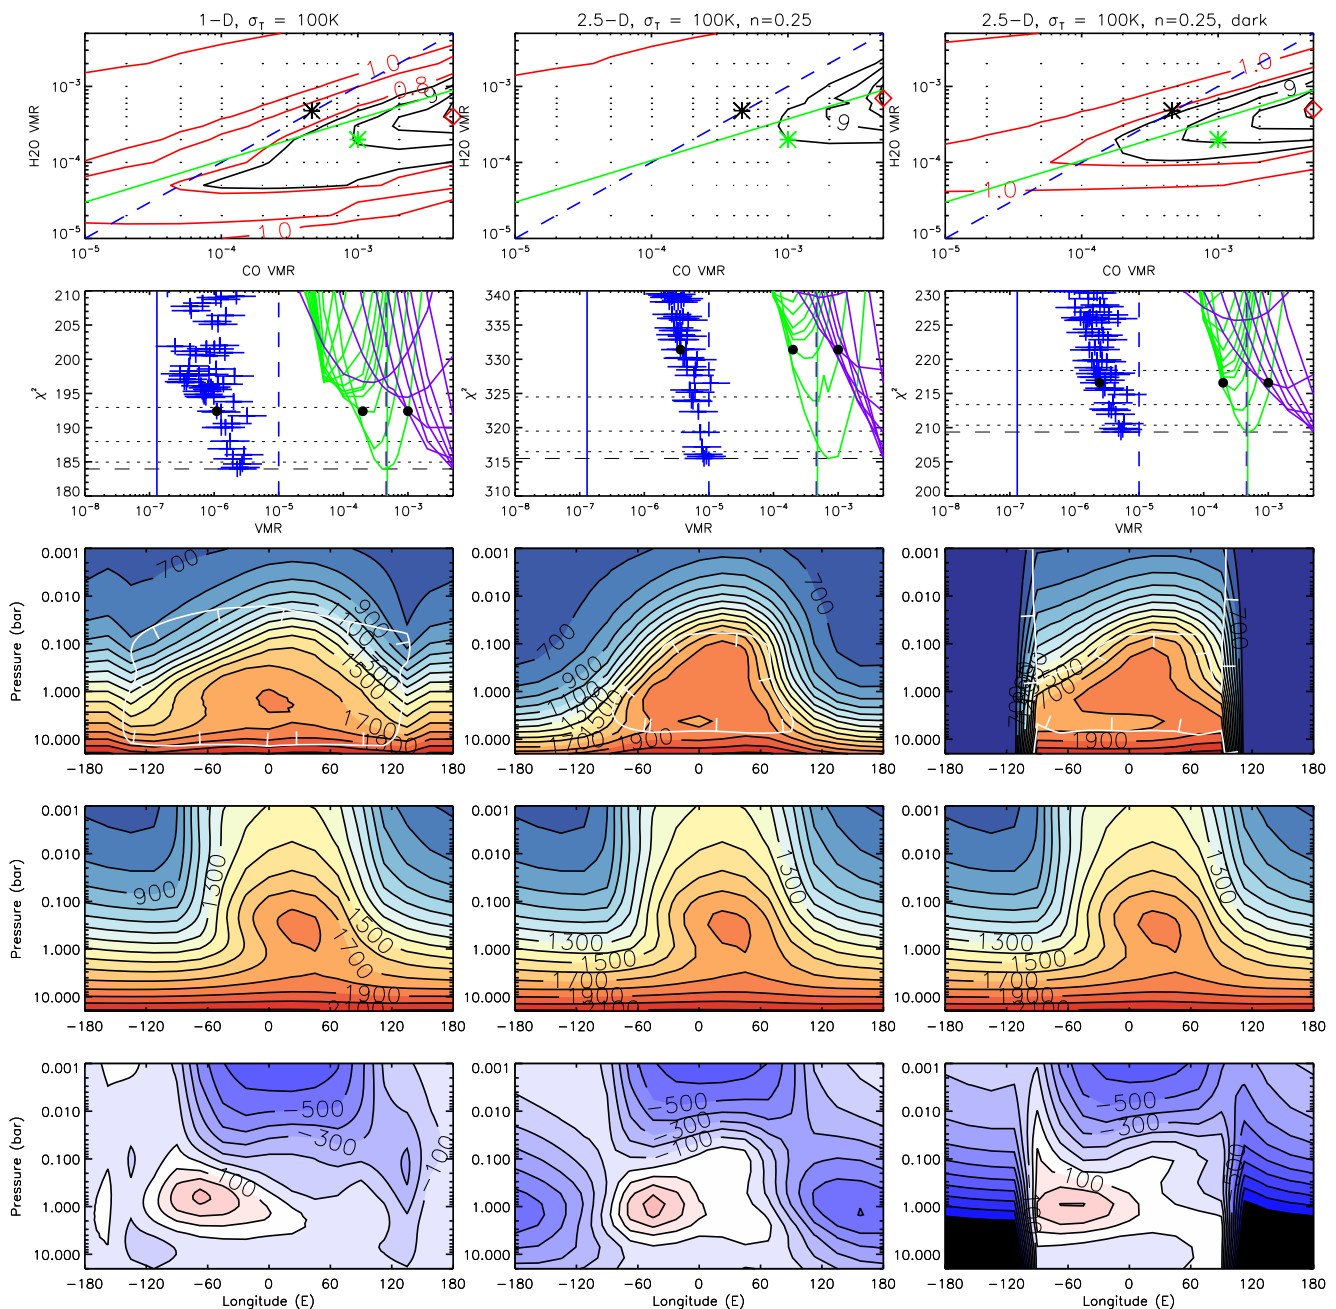


Figure 10. Application of our 1D and 2.5D models ($n = 0.25$) to the measured phase curves of Stevenson et al. (2017), assuming an *a priori* temperature error of ± 100 K. The rows are as described for Fig. 7. For the temperature maps, we have chosen H₂O and CO abundances (indicated by green asterisks in row 1) that balance the requirement to have the lowest χ^2 values with the expectation to have roughly equal abundances (2×10^{-4} and 1×10^{-3} , respectively for all cases).

shifted more to the east near the equator, but further westward at mid to polar latitudes (Fig. 2). It seems in our retrievals that this pattern may be even more enhanced in the actual atmosphere of WASP-43b. This day–night temperature difference becomes much more marked for the 2.5D case where the modelling finds the temperature at all longitudes simultaneously. We repeated these retrievals for different values of the latitude dependence coefficient, n , and as for the GCM simulations found the best fits for $n = 0.25$. Fig. 5 shows the fitted appearance of the disc of WASP-43b at $4.5 \mu\text{m}$ from our 1D and 2.5D retrievals for the case $\sigma_T = \pm 100$ K (for $n = 0.25$ and $n = 1$)

where it can be seen that the 2.5D approach generates much more realistic-looking disc images, which correspond much better with the expected appearance of this planet.

The much colder temperatures retrieved on the nightside of WASP-43b (especially for the 2.5D case), which is difficult to reproduce in GCM simulations, means that emission from the nightside contributes much less to the disc-averaged radiance for phases that can see both sides of the planet. Although the Improvement Factors are found to be lower on the nightside than the dayside for the 2.5D retrieval, we should not think that the retrieval

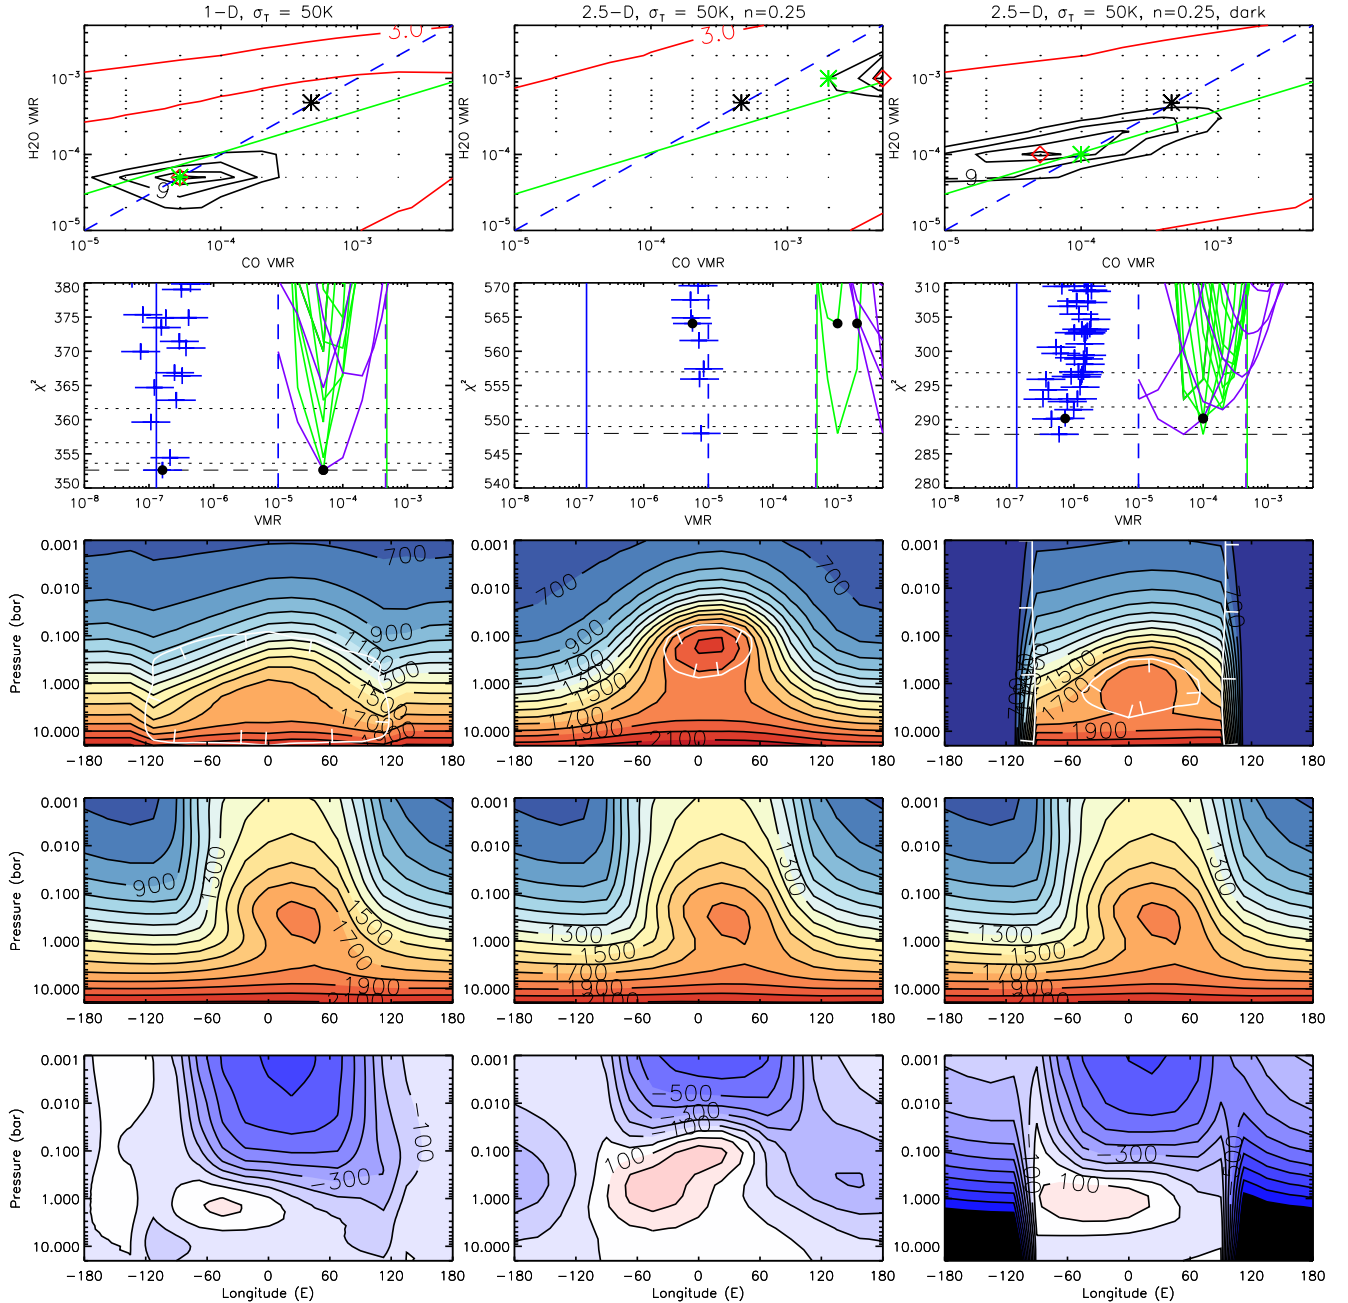


Figure 11. Application of our 1D and 2.5D models ($n = 0.25$) to the measured phase curves of Stevenson et al. (2017), assuming an *a priori* temperature error of ± 50 K. The rows are as described for Fig. 7. For the temperature maps we have chosen H_2O and CO abundances (indicated by green asterisks in row 1) that balance the requirement to have the lowest χ^2 values with the expectation to have roughly equal abundances. In this case, rather different solutions are chosen for the three cases shown.

does not return information on the nightside. Rather its best solution comes from driving the nightside temperatures so low that it cannot specify exactly how cold they should be, only that they should be low enough not to contaminate too much the radiance from the dayside. This phenomenon has been noted before (Kataria et al. 2015; Parmentier et al. 2016) and it has been used to suggest that if the temperatures fall low enough on the nightside that a thick cloud forms at a level which effectively obscures any emission. From the transmission weighting functions shown in Fig. 6 we can see that such a cloud must have a cloud top pressure less than ~ 200 mb

in order to suppress emission at all the WFC3/IRAC wavelengths. To explore whether such a scenario could be supported within our 2.5D retrieval framework, we repeated our retrievals with the same temperature *a priori* on the dayside, but fixing the temperatures to be arbitrarily low (500 K) on the nightside, so that we only retrieved the dayside temperature variation and dayside abundances of H_2O , CO and CH_4 . Our results are shown in column 3 of Figs 9–11 for $\sigma_T = \pm 200$, ± 100 , and ± 50 K, respectively. Here, it can be seen from the fitted χ^2/n values that we achieve a much closer fit to the data when we force the nightside emission to be very low than

when we try to fit the temperature on both dayside and nightside. This suggests that in the former retrievals the temperatures on the nightside were limited from falling as low as the retrieval model wanted by the *a priori* constraints. By forcing the nightside to be very cold (and thus simulating the presence of a cloud with a top at <200 mb and cloud-top temperature of ~ 500 K) we obtain an excellent fit to the observations and can also place tighter constraints on the gaseous abundances with the H_2O mole fraction found to be in the range $(2-10) \times 10^{-4}$. Hence, we conclude that the WASP-43b phase curve observations are consistent with the nightside of this planet being obscured by cloud. The simulated appearance of WASP-43b at $4.5 \mu\text{m}$ with this ‘cloudy’ 2.5D model is also shown in Fig. 5. Figs 12 and 13 show that our retrievals (for $\sigma_T = 100$ K) achieve very close fits to the observed phase curves of WASP-43b.

Although the retrieved temperatures appear plausible and the observations are consistent with the presence of cloud on the nightside, the gas abundances inferred from these observations is more puzzling. It can be seen that for all retrievals from the real observations of WASP-43b (Stevenson et al. 2017) a CO abundance of 5×10^{-3} is favoured, unless the *a priori* temperature error is reduced to <100 K, in which case the stronger *a priori* constraint leads to noticeably poorer fits and higher χ^2/n values. We also find that the methane abundance lies in the range $(1-10) \times 10^{-6}$ for the 2.5D cases. At high temperatures (such as those found in the atmosphere of WASP-43b) we expect carbon to be predominantly in the form CO and oxygen to be split between CO and H_2O . We thus expect the CH_4 abundance to be very low at such temperatures. For a solar composition atmosphere the C/O ratio is ~ 0.53 and so we expect roughly equal abundances of H_2O and CO and very low CH_4 abundance, which is what is found in the GCM-based model. However, for the WASP-43b observations we find the best fit CO and CH_4 abundances to be 10 times greater than expected, from which we infer a C/O ratio of ~ 0.91 , since $\text{CO}/\text{H}_2\text{O} \sim \text{C}/(\text{O}-\text{C})$. The question of the C/O ratio in the atmospheres of giant exoplanets has been addressed by a number of authors (e.g. Öberg, Murray-Clay & Bergin 2011; Madhusudhan, Amin & Kennedy 2014; Espinoza et al. 2017; Cridland et al. 2019). When planetesimals form they trap more oxygen than carbon in solid form, meaning that the gas is enriched in C/O. Authors such as Öberg et al. (2011) argue that this means that giant planets should be carbon-rich since they accrete more gas than planetesimals, but authors such as Espinoza et al. (2017) note that if such planets accrete both gas and solids at the same time, then they should generally be oxygen-rich. Hence, there is considerable interest in the C/O ratio of planetary atmospheres and so our retrieved C/O ratio is potentially significant. However, the derivation of this ratio is prone to several sources of systematic error, which need to be fully addressed. One possible source of systematic error we considered was whether we might be using an incorrect stellar spectrum (needed to convert the measured $F_{\text{plan}}/F_{\text{star}}$ observations to radiance). However, comparing our stellar spectrum from the Kurucz ATLAS model atmospheres (Castelli & Kurucz 2004) with the stellar spectrum for WASP-43b assumed by other authors from PHOENIX (Allard et al. 2000) we found insignificant differences. Instead, we believe our main source of error in determining C/O is that the abundance of CO comes almost entirely from the $4.5\text{-}\mu\text{m}$ *Spitzer*/IRAC observations and the CH_4 abundance from the $3.6\text{-}\mu\text{m}$ *Spitzer*/IRAC observations, while the temperature comes from the combined *HST*/*Spitzer* observations. We already noted earlier that there is considerable degeneracy between atmospheric temperature near the tropopause and the abundances of these two constituents. However, in addition to this, we must remember that if there is any systematic error in the $F_{\text{plan}}/F_{\text{star}}$ ratios between the

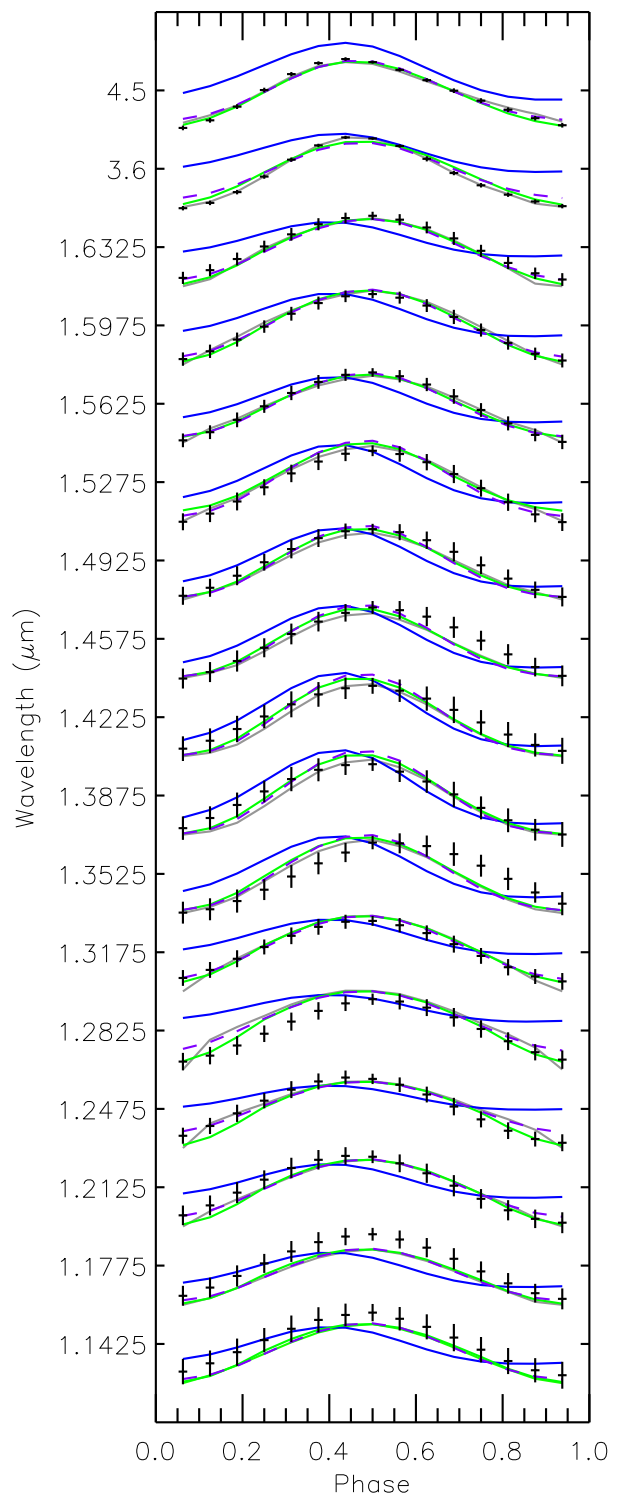


Figure 12. Fit to the phase curves ($\sigma_T = 100$ K) for each wavelength bin reported by Stevenson et al. (2017), normalized for ease of reference. The measured, normalized phase curves and errors are indicated by the points and error bars. The line colours are: grey – fits to these curves using our 1D model; purple-dashed – fits to these curves using our 2.5D model, assuming $(\cos \Phi)^{0.25}$ latitude dependence; and green – fits to these curves using our 2.5D model, assuming $(\cos \Phi)^{0.25}$ latitude dependence and forcing nightside temperatures to 500 K. For reference, the normalized synthetic phase curves calculated from our GCM-based model atmosphere are shown in blue.

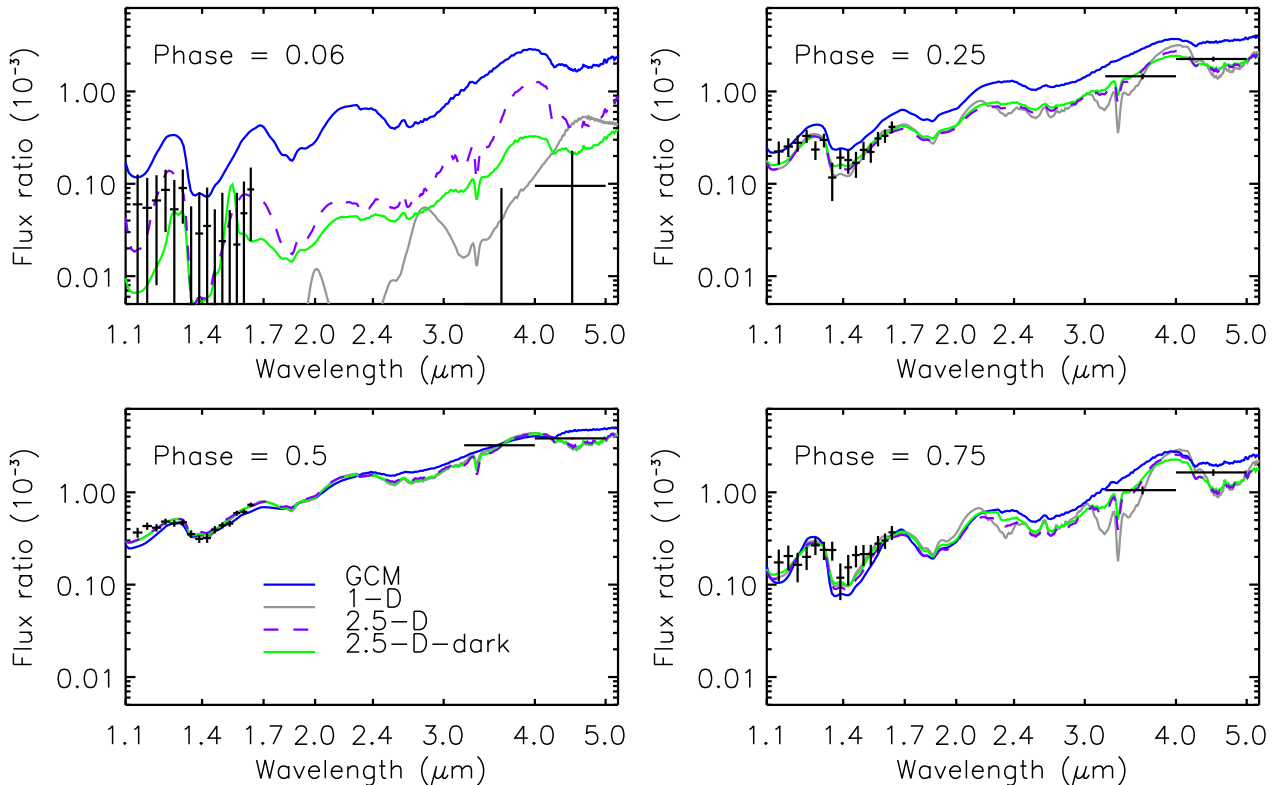


Figure 13. Phase curves for each wavelength bin reported by Stevenson et al. (2017), compared with moderate resolution spectra ($\Delta\lambda = 0.02 \mu\text{m}$) calculated from the fitted atmospheric temperatures and abundances ($\sigma_T = 100 \text{ K}$) for the same geometry. Curves are compared for phases 0.06, 0.25, 0.5, and 0.75. The observations are compared to the *ab initio* GCM-based simulations (blue), our 1D model fit (grey), our 2.5D model, assuming $(\cos \Phi)^{0.25}$ latitude dependence (purple-dashed) and our 2.5D model, assuming $(\cos \Phi)^{0.25}$ latitude dependence and forcing nightside temperatures to 500 K (green).

HST/WFC3 and *Spitzer*/IRAC observations, we could very easily retrieve systematically offset CO abundances. For example, if the *Spitzer* radiances flux ratios are either lower than they should be compared with the *HST* observations or, alternatively, if the *HST* flux ratios are higher than they should be, then the optimal fit to the combined data would be obtained by increasing the CO and CH_4 abundance to lower the *Spitzer* radiances relative to the *HST* radiances, which may be what is happening here. We note that in two independent reprocessings of the *Spitzer* flux ratios Mendonça et al. (2018) and Morello et al. (2019) found significantly different values from Stevenson et al. (2017) for the *Spitzer*/IRAC channels, which suggests that there may indeed be systematic differences between the *HST*/WFC3 and *Spitzer*/IRAC channels in the reported phase curve spectra. The fact that we find both CO and CH_4 to be larger than we might expect for a hot solar composition atmosphere leads us to suspect that these differences do exist. Finally, we note that the larger CO abundances we infer here compared with the lower abundances retrieved from the same data by Stevenson et al. (2017) are most likely caused by the differences in temperature profile parametrization. In this work we have allowed the atmospheric temperature to vary at all altitudes and we have seen that there is a degeneracy between how rapidly the temperature is allowed to vary with altitude and the retrieved CO abundance. The temperature profile used in retrieval model employed by Stevenson et al. (2017), CHIMERA (Line et al. 2014), is parametrized (as is necessary for Bayesian models that explore a large parameter space) and so has an implicit limitation in vertical variability; any such constraint on vertical variability could lead to systematic differences in the inferred CO abundance.

Although the best fits are for rather large CO values, it can be seen from Figs 9–11 that a wide range of CO and H_2O values gives fits to the data with a $\chi^2/n < 1.0$. If we assume that the metallicity of WASP-43b is similar to that of the Sun then we would expect the abundances of H_2O and CO to be very similar and thus the solution to lie along the dashed blue lines in the χ^2 contour plots of Figs 9–11. Looking at the right-hand column (i.e. where the nightside emission has been forced to be very low) for Fig. 10, for $\sigma_T = 100 \text{ K}$, it can be seen that solutions where the mole fraction of both H_2O and CO are in the range $(1 - 7) \times 10^{-4}$ fit the observed spectra to $\chi^2/n < 0.9$. For these values, the mole fraction of CH_4 is $(1 - 3) \times 10^{-6}$. Since we fix the abundance of CO_2 to 10^{-7} our estimate for $\text{CO} + \text{CO}_2$ is effectively the same as that for CO.

For comparison, with their 1D Bayesian approach and parametrized temperature profiles, Stevenson et al. (2017) report mole fractions of $10^{(-4.28 \pm 0.32)} = (0.25 - 1.1) \times 10^{-4}$ for H_2O , $10^{(-3.5 \pm 0.4)} = (1.3 - 7.9) \times 10^{-4}$ for $\text{CO} + \text{CO}_2$, and $> 10^{-5.3} = 5 \times 10^{-6}$ for CH_4 . These results overlap with our determinations, but we find higher abundances of H_2O are consistent with the observations.

7 DISCUSSION

We have developed a novel retrieval technique for modelling the phase curves of exoplanets that we have shown is robust and which provides an excellent fit to the observations of WASP-43b reported by Stevenson et al. (2017). Our approach uniquely models the radiance contributed to the disc-averaged radiance from both dayside and nightside simultaneously with a full treatment of

observation angles and disc integration. We find that the nightside temperatures are required to be much cooler than those predicted from the traditional 1D approach to be consistent with observations and our fitted temperatures and gaseous compositions agree much better with the ‘true’ values of our GCM-based model atmosphere when tested on GCM-based synthetic observations. Indeed, we find nightside temperatures of WASP-43b to be so low that we can model the observations effectively from the dayside radiances only, suggesting that the nightside of WASP-43b may be obscured by a thick cloud with a cloud top at <0.2 bar.

We also find considerable degeneracy in our retrievals between gaseous abundance and temperature. While the abundance of water vapour is reasonably well constrained, the retrieval of the abundance of CO and CH₄ comes effectively from the two *Spitzer*/IRAC observations at 3.6 and 4.5 μm , co-retrieved with temperatures derived to be most consistent between the *HST* and *Spitzer* channels. It is found that allowing different degrees of vertical smoothing in the retrieved temperature profiles gives substantial differences in the retrieved CO abundance. This behaviour helps to explain why our retrievals appear to favour considerably more CO than the study of Stevenson et al. (2017), which used a temperature profile that was more vertically constrained. We do not believe the difference is caused by systematic radiative transfer differences in the NEMESIS model used here and the CHIMERA model (Line et al. 2014) used by Stevenson et al. (2017), since a recent study (Barstow et al. 2018) has shown that these models generate consistent spectra given identical inputs. However, we note that possible systematic differences between the radiometric calibration between the *HST*/WFC3 and *Spitzer*/IRAC datasets could easily lead to large differences in the retrieved CO/H₂O ratio.

Although our approach is a significantly more geometrically realistic than previous methods, there are potentially a number or areas in which it could be improved and/or developed, which we will discuss now.

(1) We used an optimal estimation retrieval model here because such an approach is fast, especially in implicit differentiation mode (where the Jacobians are calculated analytically rather than by numerical differentiation), which NEMESIS can employ for thermal emission calculations. In addition, we used channel-averaged k -tables, which gives an additional improvement in computational speed. Even then, however, the retrievals are time consuming as we have to model $N_{\text{phase}} \times N_{\text{wave}} = 15 \times 17 = 255$ flux ratio measurements using a disc-integration scheme that for $N_{\mu} = 5$ comprises 45–47 points on the disc for each phase angle. Hence, $255 \times 45 = 11,475$ radiative transfer calculations are done for each iteration of the model. Since our optimal estimation model typically takes 10–15 iterations to converge, this gives $\sim 138,800$ calculations per retrieval. While the optimal estimation method is fast and efficient, it is less easy to explore the full solution space than Bayesian techniques such as MCMC or nested sampling. However, in this study we have mapped a wide range of acceptable solutions by performing retrievals for a grid of assumed H₂O and CO abundances and find that such an approach allows us to give conservative estimates on the likely gaseous abundances and leads to more realistic-looking temperature distributions than the 1D approach. Nevertheless, although Bayesian techniques need $\sim 10,000$ iterations before they converge, which is considerably more than 10–15 iterations needed for optimal estimation, we plan in the near future to incorporate the new 2.5D approach into our nested sampling version of NEMESIS to more fully explore solutions that are consistent with the observed phase curves.

(2) We have modelled these observations assuming no clouds. It is possible that the reason we see such a large day/night contrast in the WASP-43b observations is not that the temperatures vary greatly with longitude, but that the nightside, which is cooler, forms clouds that obscure the thermal emission from lower altitudes (e.g. Kataria et al. 2015; Parmentier et al. 2016). We have assessed this possibility with our model and find that it gives improved fits to the observations. As we note, looking at the weighting functions in Fig. 6 it can be seen that such a cloud would have a cloud top at pressure level of <0.2 bar in order to suppress the radiation at all 17 wavelengths considered here, with a cloud top temperature of ~ 500 K. In this paper we have modelled the effect of clouds by forcing the nightside temperatures to arbitrarily low values. However, in future work we could instead add a cloud into our radiative transfer scheme. To do this, we would need to make some estimate of its complex refractive index spectrum in order to compute, with Mie theory, its extinction cross-section spectrum. Since we need a cloud that is opaque over the whole range considered here of 1–5 μm , it is likely that the particles must have radii on the order of a few μm .

(3) If we are in the future to include clouds in our model, we should then also consider the scattered stellar contribution they might make. We noted earlier and showed in Fig. 6 that reflection from cloud could potentially make a significant contribution to the observed flux ratios at wavelengths less than 2 μm . We have designed our disc-averaging scheme to be able to incorporate scattering, but such computations are hugely computationally challenging. Not only are the radiative transfer calculations themselves computationally expensive, but our implicit differentiation scheme is inefficient for tracking the partial derivatives of our Matrix-Operator scattering model (Plass, Kattawar & Catchings 1973) and hence the functional derivatives are calculated numerically. In our current model we have 16 longitudes, 20 pressure levels for temperature, and 4 gas abundances giving 384 state vector elements. Hence, if calculating the functional derivatives numerically, our radiative transfer calculations would be at least 385 times slower, even before factoring in the at least 10 times slower speed of our scattering model. Hence, such modelling will be challenging.

(4) We have here assumed the temperature profile to be continuous and retrieve at all pressures in our model. The use of a parametrized profile would simplify and further increase the speed of our model, but at the disadvantage of restricting the range of solutions we can explore to those covered by the parametrization. Such an approach could be optimal for Bayesian realizations of our 2.5D retrieval technique, but as we have seen there is considerable degeneracy between temperature and the abundance of CO and H₂O. By constraining the temperature profile more strongly, we may bias our estimates of the abundances of the gaseous constituents of the atmosphere of WASP-43b. It can also be seen from Figs 9 to 11 and previous Bayesian retrievals from these data (e.g. Stevenson et al. 2017), that the prior and method chosen to fit to the observations can significantly change the retrieved temperatures also. Hence, the accuracy with which we can estimate how the temperature varies with longitude and altitude and also the molecular abundances is strongly limited by the inherent degeneracy in these *HST*/WFC3 data.

(5) We have assumed here that there is a single n coefficient the $(\cos \Phi)^n$ latitude dependence. However, looking at the brightness of the disc simulated from the GCM atmosphere in Fig. 5 it can be seen that while $(\cos \Phi)^{0.25}$ is generally most appropriate, at some phases $\cos \Phi$ or $(\cos \Phi)^2$ might be a better approximation. It would be interesting to explore whether a parametrization that combines different n -dependencies might give more reliable retrievals.

(6) Finally, in this study we assumed the abundance of CH₄ to be the same at all longitudes for the 2.5D model, because of the expected horizontal ‘quenching’. However, our model set-up is also able to retrieve abundances that do vary with longitude, as noted earlier. In early tests with such a model using synthetic phase curves generated from GCM simulations that included longitudinal abundance variation (rather than the hybrid test model used here, where the abundances were set to longitudinal averages) we found that we were able to recover significant longitudinal variation in CH₄, approximately matching the synthetic model variation. However, adding this additional degree of freedom led to some instability in the temperature retrievals, which is easily understood given the degeneracy between CH₄ abundance and temperature at pressures greater than 0.1 bar, already noted. However, future measurements covering a wider wavelength range at higher SNR might be able to break such degeneracy, and we could then use this approach to determine longitudinal variations in abundance also.

8 CONCLUSION

We have developed a novel retrieval technique, which we call a 2.5D retrieval approach that attempts to model more realistically and reliably phase curves observations of exoplanets. In our 2.5D approach we retrieve the vertical temperature profile and mean gaseous abundance at all longitudes and latitudes simultaneously, assuming that the temperature or composition, x , at a particular longitude and latitude (Λ, Φ) is given by $x(\Lambda, \Phi) = \bar{x} + (x(\Lambda, 0) - \bar{x}) \cos^n \Phi$, where \bar{x} is the mean of the morning and evening terminator values of $x(\Lambda, 0)$, i.e. $\bar{x} = (x(-90^\circ, 0) + x(90^\circ, 0))/2$, and n is an assumed coefficient. We find that this model fits the synthetic flux ratio phase curve observations of our WASP-43b GCM-based simulations more realistically than the traditional 1D approach, and provides a more geometrically realistic retrieval of the observed WASP-43b phase curve observations of Stevenson et al. (2017). Our retrieved gaseous abundances are broadly consistent with previous determinations and we find the mole fraction of H₂O lies in the range $(1-10) \times 10^{-4}$, which is consistent with the atmosphere of this planet having an approximately solar composition. In tests where we retrieve from synthetic phase curves generated from a GCM-based model, we find that the current *HST/Spitzer* WASP-43b data provide much poorer constraint on the abundance of CO, with the choice of *a priori* temperature profile influencing the best-fitting parameters. Finally, when fitting to the measured phase curves reported by Stevenson et al. (2017) we find that the best-fitting CO and CH₄ abundances both seem biased to larger values than we might expect, with CH₄ ~ 10 times larger than the expected thermodynamically-expected value and CO/H₂O ~ 10 , suggesting a C/O ratio of ~ 0.91 . While the CO/H₂O results alone might lead us to conclude than C/O ~ 1 , the fact that the CH₄ abundance also seems overestimated leads us to suspect that there might be a systematic difference in the radiometric calibrations between the *HST/WFC3* and *Spitzer/IRAC* observations. We conclude that the uncertainty on the *Spitzer* points is probably underestimated by the error bars assumed in these retrievals as three different data reductions of the same phase curve (Stevenson et al. 2017; Mendonça et al. 2018; Morello et al. 2019) lead to flux ratio differences larger than 1σ . Since the inferred C/O ratio is highly dependent on the relative *HST* versus *Spitzer* flux ratios (and the errors on these ratios) a robust determination will not be possible until a consensus is arrived at in the community on the radiometric calibrations of these data sets, which includes all possible contaminating effects (e.g. stellar variability, Knutson et al. 2012).

Having demonstrated the efficacy of this 2.5D approach for the *HST/Spitzer* phase curve observations of WASP-43b, the next step will be to apply them also to recent observations of WASP-103b (Kreidberg et al. 2018) and WASP-18b (e.g. Arcangeli et al. 2019) and possible new observations of other exoplanets. Further into the future, the *James Webb Space Telescope (JWST)* plans to make phase curve observations of several exoplanets over a wide wavelength range with much higher SNR than the *HST/Spitzer* observations analysed here. It is hoped that these improved observations will break many of the degeneracies inherent in existing observations and allow us to retrieve temperature profiles that are much less dependent on the prior and method chosen to fit to the observations. In particular, the acquisition of high-SNR observations covering the absorption features of H₂O, CO, and CH₄ should allow us to more effectively discriminate between the abundances of these gases and temperature in the atmosphere of WASP-43b and other exoplanets and thus enable us to more reliably determine the C/O ratios and enrichment of heavy elements in these exotic worlds. The new 2.5D approach described here should be particularly effective in analysing these improved data and we intend to start testing this approach on synthetic *JWST* observations in the near future.

ACKNOWLEDGEMENTS

This work was supported by United Kingdom’s Science and Technology Facilities Council. The authors would like to sincerely thank the anonymous reviewer for their careful and detailed review, and very helpful suggestions on how to improve the presentation of this paper.

REFERENCES

- Agúndez M., Parmentier V., Venot O., Hersant F., Selsis F., 2014, *A&A*, 564, A73
- Allard F., Hauschildt P. H., Schweitzer A., 2000, *ApJ*, 539, 366
- Arcangeli J. et al., 2019, *A&A*, 625, A136
- Barber R. J., Tennyson J., Harris G. J., Tolchenov R. N., 2006, *MNRAS*, 368, 1087
- Barstow J. K., Aigrain S., Irwin P. G. J., Fletcher L. N., Lee J. M., 2013, *MNRAS*, 434, 2616
- Barstow J. K., Aigrain S., Irwin P. G. J., Hackler T., Fletcher L. N., Lee J. M., Gibson N. P., 2014, *ApJ*, 786, 154
- Barstow J. K., Aigrain S., Irwin P. G. J., Sing D. K., 2017, *ApJ*, 834, 50
- Barstow J., Garland R., Line M., Rocchetto M., Waldmann I., 2018, European Planetary Science Congress, Vol. 12, EPSC2018
- Blecic J., Dobbs-Dixon I., Greene T., 2017, *ApJ*, 848, 127
- Borysov A., Frommhold L., 1989, *ApJ*, 341, 549
- Borysov A., Frommhold L., Moraldi M., 1989, *ApJ*, 336, 495
- Burrows A., Budaj J., Hubeny I., 2008, *ApJ*, 678, 1436
- Castelli F., Kurucz R. L., 2004, *A&A*, 419, 725
- Cooper C. S., Showman A. P., 2006, *ApJ*, 649, 1048
- Cowan N. B., Agol E., 2008, *ApJ*, 678, L129
- Cridland A. J., van Dishoeck E. F., Alessi M., Pudritz R. E., 2019, *A&A*, 632, A63
- Drummond B., Mayne N. J., Manners J., Baraffe I., Goyal J., Tremblin P., Sing D. K., Kohary K., 2018, *ApJ*, 869, 28
- Espinoza N., Fortney J. J., Miguel Y., Thorngren D., Murray-Clay R., 2017, *ApJ*, 838, L9
- Evans T. M. et al., 2018, *AJ*, 156, 283
- Feng Y. K., Line M. R., Fortney J. J., Stevenson K. B., Bean J., Kreidberg L., Parmentier V., 2016, *ApJ*, 829, 52
- Gandhi S., Madhusudhan N., 2018, *MNRAS*, 474, 271
- Garland R., Irwin P., 2019, preprint ([astro-ph/0405087](https://arxiv.org/abs/1904.05087))
- Gillon M. et al., 2012, *A&A*, 542, A4

- Gillon M. et al., 2014, *A&A*, 562, L3
 Hellier C. et al., 2009, *Nature*, 460, 1098
 Hellier C. et al., 2011, *A&A*, 535, L7
 Irwin P. G. J. et al., 2008, *J. Quant. Spectrosc. Radiat. Transfer*, 109, 1136
 Kataria T., Showman A. P., Fortney J. J., Stevenson K. B., Line M. R., Kreidberg L., Bean J. L., Désert J.-M., 2015, *ApJ*, 801, 86
 Keating D., Cowan N. B., 2017, *ApJ*, 849, L5
 Knutson H. A. et al., 2012, *ApJ*, 754, 22
 Kreidberg L. et al., 2014, *ApJ*, 793, L27
 Kreidberg L. et al., 2018, *AJ*, 156, 17
 Krissansen-Totton J., Garland R., Irwin P., Catling D. C., 2018, *AJ*, 156, 114
 Laci A. A., Oinas V., 1991, *J. Geophys. Res.*, 96, 9027
 Lee J.-M., Irwin P. G. J., Fletcher L. N., Heng K., Barstow J. K., 2014, *ApJ*, 789, 14
 Line M. R. et al., 2016, *AJ*, 152, 203
 Line M. R., Knutson H., Wolf A. S., Yung Y. L., 2014, *ApJ*, 783, 70
 Madhusudhan N., Amin M. A., Kennedy G. M., 2014, *ApJ*, 794, L12
 Maxted P. F. L. et al., 2013, *MNRAS*, 428, 2645
 Mendonça J. M., Malik M., Demory B.-O., Heng K., 2018, *AJ*, 155, 150
 Mollière P., van Boekel R., Dullemond C., Henning T., Mordasini C., 2015, *ApJ*, 813, 47
 Morello G., Danielski C., Dickens D., Tremblin P., Lagage P. O., 2019, *AJ*, 157, 205
 Parmentier V., Fortney J. J., Showman A. P., Morley C., Marley M. S., 2016, *ApJ*, 828, 22
 Pinhas A., Madhusudhan N., Gandhi S., MacDonald R., 2019, *MNRAS*, 482, 1485
 Plass G. N., Kattawar G. W., Catchings F. E., 1973, *Appl. Opt.*, 12, 314
 Rodgers C. D., 2000, *Inverse Methods for Atmospheric Sounding: Theory and Practice*, World Scientific Publ., Singapore, New Jersey, London, Hong Kong
 Rothman L. S. et al., 2010, *J. Quant. Spectrosc. Radiat. Transfer*, 111, 2139
 Showman A. P., Fortney J. J., Lian Y., Marley M. S., Freedman R. S., Knutson H. A., Charbonneau D., 2009, *ApJ*, 699, 564
 Sing D. K. et al., 2016, *Nature*, 529, 59
 Skilling J., 2006, *Bayesian Anal.*, 1, 833
 Steinrueck M. E., Parmentier V., Showman A. P., Lothringer J. D., Lupu R. E., 2019, *ApJ*, 880, 14
 Stevenson K. B. et al., 2014, *Science*, 346, 838
 Stevenson K. B. et al., 2017, *AJ*, 153, 68
 Tashkun S. A., Perevalov V. I., 2011, *J. Quant. Spectrosc. Radiat. Transfer*, 112, 1403
 Yurchenko S. N., Tennyson J., 2014, *MNRAS*, 440, 1649
 Öberg K. I., Murray-Clay R., Bergin E. A., 2011, *ApJ*, 743, L16

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.