MNRAS **491,** 5126–5139 (2020) Advance Access publication 2019 November 15

Discarding the disc in a changing-state AGN: the UV/X-ray relation in NGC 4151

Ra'ad D. Mahmoud[®] and Chris Done

Centre for Extragalactic Astronomy, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK

Accepted 2019 November 12. Received 2019 November 12; in original form 2019 August 14

ABSTRACT

Recent monitoring campaigns designed to map the accretion regime in active galactic nuclei (AGN) show major discrepancies with models where the optical/ultraviolet (UV) is produced by X-ray-illuminated, optically thick disc material within a few hundred gravitational radii. However, these campaigns only monitored X-rays below 10 keV, whereas the bolometric luminosity for most of these AGN peaks above 50 keV. We use data from the recent multiwavelength campaign by Edelson et al. on NGC 4151 – the only AGN bright enough to be monitored at higher energies with Swift Burst Alert Telescope (BAT). We develop a spectral-timing model with a hot corona, warm Comptonization, and outer standard disc. This fits the time-averaged spectrum well, but completely fails to match the UV variability predicted from the X-ray light curve. However, it reveals that NGC 4151 had a bolometric luminosity around 1.4 per cent of the Eddington luminosity during this campaign, close to the luminosity at which AGN show a 'changing-state' transition, where the broad optical lines disappear. Stellar mass black holes show a state transition at a similarly low Eddington fraction, which is broadly interpreted as the inner disc being replaced by an optically thin flow. We find that the UV light curve can instead be matched by reprocessing of the X-ray flux on size scales of the broad-line region (BLR; 1.5–20 light-days) and rule out there being optically thick material inwards of this, as expected if the thin disc is replaced by the flow below the inner radius of the BLR. These results emphasize the need for even longer time-scale, multiwavelength monitoring campaigns on variable AGN.

Key words: accretion, accretion discs – black hole physics – galaxies: active – galaxies: individual: NGC 4151.

1 INTRODUCTION

Active galactic nuclei (AGN) are powered by accretion on to a supermassive black hole, and the energy released makes them the most luminous persistent sources in the sky in the optical, UV, and X-ray bands. In the standard picture of accretion, the gravitational energy is dissipated in an optically thick, geometrically thin disc that extends down to the innermost stable circular orbit (ISCO). In this disc, the temperature increases radially inwards, so that the total spectrum is the sum over all radii of the local blackbody spectra produced by each annulus (Shakura & Sunyaev 1973). The peak disc temperature is linked to the mass of the black hole and its luminosity by $T_{\text{peak}} \sim 2 \times 10^5 (\dot{m}/m_8)^{1/4}$ K, where $m_8 = M_{\text{BH}}/10^8 \text{ M}_{\odot}$ and \dot{m} is the mass accretion rate as a fraction of the Eddington rate. These typically far-UV temperatures mean that the bulk of the accretion power in bright AGN ($\dot{m} > 0.01$) is hidden by interstellar absorption. In the absence of this absorption, the optical/near-UV

* E-mail: ra'ad.d.mahmoud@durham.ac.uk

should be dominated by a sum of blackbodies from the disc, giving a differential flux $F_{\nu} \propto \nu^{1/3}$. However, most optical/UV spectra of AGN are redder than this simple prediction, with mean $F_{\nu} \propto \nu^{-1/2}$ (Richards et al. 2003). Some of this may be due to circum-AGN dust (Davis, Woo & Blaes 2007; Baron et al. 2016), but there is also strong evidence of a marked downturn in the intrinsic far-UV spectrum (see e.g. Zheng et al. 1997; Telfer et al. 2002; Shull, Stevans & Danforth 2012; Laor & Davis 2014).

The X-ray spectra reveal another discrepancy with simple disc models. First, there is an X-ray power-law tail that extends out to between a few tens and a few hundred keV. There is a general consensus that the presence of the X-ray tail indicates that some fraction of the gravitational power is released in an optically thin region, distinct from the optically thick disc. Electrons in this 'hot corona' can be heated to much higher temperatures than in the dense disc, producing the power-law tail by Compton upscattering of seed photons from the disc (as is also seen in the black hole binaries; see e.g. the review by Done, Gierliński & Kubota 2007). However, there is additional soft X-ray emission below 1 keV: the soft X-ray excess (e.g. Porquet et al. 2004). The origin of this is not well understood, 2015; Jin et al. 2012; Petrucci et al. 2018). One way to understand this structure is to use variability timescales to map the different regions. The hard X-ray emission varies rapidly with time-scales of less than 10–100 R_g/c (where $R_g = GM_{\rm BH}/c^2$; Galeev, Rosner & Vaiana 1979; Frank, King & Raine 2002), indicating that the size scale of the hard coronal region is small. Time-scales for variability in the standard disc models are very much longer, so any rapid optical/UV variability should be produced by reprocessing of the variable X-ray illumination. Thus tracking the relation of the reprocessed optical/UV emission to the driving X-ray flux variability gives a diagnostic of the structure of the outer accretion disc.

This reverberation technique is well established for deriving the size scale of the broad-line region (BLR; see e.g. the review by Peterson & Horne 2004), but the recent intensive optical/UV/X-ray monitoring campaigns with the *Neil Gehrels Swift Observatory* (hereafter *Swift*) have revolutionized the search for *continuum* reverberation from the outer disc. However, the first campaign on NGC 5548 showed the level of disconnect between reprocessing models and observations. First, it showed there was very little correlation between the X-ray variability and the optical/UV variability. In addition, while the UV and optical variability were well correlated with each other, the lag between them was larger than predicted by the disc models by a factor of ~ 2 (Edelson et al. 2015).

Subsequent intensive monitoring campaigns on other AGN have only reinforced these two problems, highlighting our lack of understanding of the accretion disc structure rather than revealing it (Arévalo et al. 2009; Mehdipour et al. 2011; Lawrence 2018). The X-ray variability is generally poorly correlated with the optical and UV variability, in conflict with the idea that the optical and UV are produced by reprocessed X-ray illumination (Edelson et al. 2017, hereafter E17; Buisson et al. 2018), and while the optical is well correlated with the UV, the lag between them is larger than predicted by the standard disc models (E17).

However there is much more information in the monitoring campaign data than just the lags between the variability seen in different wavebands. Incorporating the spectral information is important, as it constrains the power in the X-rays relative to the optical/UV. Given an assumed disc geometry, this also predicts how much X-ray luminosity intercepts the disc, at which point one can model the disc reverberation. Gardner & Done (2017, hereafter GD17) developed the first full spectral-timing model and used it to predict the reprocessed variability from the observed X-ray variability in the Edelson et al. (2015) *Swift* NGC 5548 data set. These predictions completely fail to reproduce the observed optical and UV light curves. Even allowing the disc size scale to vary does not help.

GD17 proposed two possible solutions for the first problem that the observed X-ray variability did not look like the observed UV variability. One was that the model is fundamentally incorrect, and the hard X-rays do not illuminate the optical/UV disc, possibly due to vertical structure on the disc that shields the outer disc from direct irradiation. The second possibility was that the model is essentially correct in assuming that the optical and UV variability is driven by reprocessing of the variable hard X-ray emission, but that the observed *Swift* X-ray Telescope (XRT) X-ray light curve did not track the true bolometric X-ray flux variations. This is clearly possible in NGC 5548 as the bolometric flux peaks at 100 keV (Mehdipour et al. 2011), whereas the *Swift* XRT band extends only up to ~10 keV. However, the longer than expected lags between the



Figure 1. Mean normalized light curves from the NGC 4151 *Swift* campaign in *Swift* BAT (panel a) and *Swift* UVOT *UVW*1 (panel b; E17).

optical and UV then also required that the material being illuminated has a different size and geometry than expected from an outer disc.

Here, we apply the procedure of GD17 to the intensive monitoring campaign data from one of the brightest hard X-ray AGN, NGC 4151 (E17). This is the only AGN where the higher energy *Swift* Burst Alert Telescope (BAT) can get reasonable constraints at 15–50 keV, simultaneously with its optical/UV from the *Swift* UV/Optical Telescope (UVOT) and the 0.5–10 keV from the XRT. This allows us to track most of the bolometric luminosity directly, so we can for the first time test the fundamental assumption of reprocessing. These data span 69 d with sampling of 6 h in the UVOT and ~2 d in the BAT; light curves for both the BAT and the representative *UVW*1 band are shown in Fig. 1.

In Section 2, we model the mean broad-band spectrum. This is dominated by the hard X-ray luminosity rather than a UV disc component and the inferred bolometric luminosity is below $0.02L_{\rm Edd}$. Stellar mass black hole binaries are observed to undergo a dramatic transition at this luminosity, from disc-dominated to Compton-dominated spectra. This transition is generally interpreted as the inner thin disc evaporating into an optically thin hot flow (e.g. Done et al. 2007). While the bright AGN spectra do not match as well to disc models as the stellar mass black hole binaries, they do appear to make a similar transition to Compton-dominated spectra at a similar luminosity, with the soft X-ray excess collapsing, correlated with the loss of the characteristic broad optical lines (changing state AGN; Noda & Done 2018; Ruan et al. 2019). It seems most likely that NGC 4151 was close to a changing state event at the time of the monitoring campaign, as has been seen in the past from this source (Penston & Perez 1984; Shapovalova et al. 2008).

None the less, there is still optical/UV emission from the AGN, which we model with an outer disc/warm Comptonization region. A strong UV contribution should be produced at radii of $100-400R_g$ in this structure, i.e. at a light travel time of less than 1 d from the hot inner flow. We predict the UV light curve resulting from illumination of this structure by the observed variable hard X-ray light curve from the *Swift* BAT, and compare this to the observed UVOT light curve both directly and via cross-correlation (Section 3). As in NGC 5548, the match is extremely poor. The reprocessing model predicts much higher levels of correlation and much shorter lags than those seen in the data, However, we find that (now *unlike* NGC 5548) there does exist a linear transfer function that allows the UV to be produced by the variable X-ray flux (Section 4). The impulse response function (IRF) required is broadly distributed

between lags of 1.5 and 20 d. If this is due to a light travel time, the responding material is cospatial with the inner BLR. We rule out there being optically thick material in a disc geometry on size scales less than 1.5 light-days, as this would produce too much UV fast variability. This variability cannot be hidden by self-shielding as proposed for NGC 5548, as even the inner edge of the warm Comptonization region contributes substantially to the UV flux, so its reprocessing signal would be seen directly, but is not.

We conclude that there is very little evidence for a standard disc structure in NGC 4151 at the time of the intensive monitoring campaign, from either its UV spectrum or light curve. Instead, the X-ray spectra are consistent with the inner disc being replaced by a hot flow, and the UV spectra and variability are consistent with being produced by reprocessing in dense gas in the BLR as first suggested by Korista & Goad (2001). This is now strongly supported by more recent work (Lawther et al. 2018; Chelouche, Nuñez & Kaspi 2019). It remains to be seen whether similar campaigns on objects at higher L/L_{Edd} will give results that connect better to expectations of standard disc models or whether continuum reverberation from inner BLR scales always dominates the UV variability.

2 ENERGETICS OF DISC REPROCESSING AND SOFT EXCESS COMPTONIZATION

The spectra of AGN are complex, typically showing three components that can be modelled by an outer disc emitting the optical and UV continuum, a hot corona producing the hard X-ray power law, with additional emission over and above this below 1 keV (the soft X-ray excess). There is now a growing consensus that much of this soft excess is produced by Comptonization in warm (kT_e) \sim 0.2 keV), optically thick ($\tau \sim$ 10–20) material that is separate from the hot, optically thin X-ray corona (e.g. Magdziarz et al. 1998; Petrucci et al. 2013; Matt et al. 2014; Mehdipour et al. 2015; Boissay, Ricci & Paltani 2016; Porquet et al. 2018). The best current model proposes that this is produced by the accretion energy being released in the upper layers of the disc rather than preferentially in the mid-plane (although there are difficulties with this picture; see García et al. 2019). This produces a sandwich geometry with a heated layer on top of a thin disc of cool, 'passive', material (Petrucci et al. 2018). Gravitational energy is radiated in the upper layers of the disc. Half of these photons are intercepted and thermalized by the underlying passive disc, to then be reradiated into the warm Compton layer as seed photons. This ties the seed photon luminosity to the warm Comptonization power, so the shape of the spectrum is fixed by energy balance (Haardt & Maraschi 1993; Petrucci et al. 2018).

Kubota & Done (2018, hereafter KD18) took this model for the warm Comptonization and extended it into a fuller picture of the accretion flow: AGNSED. This assumes that the energy released at each radius is set by the Novikov–Thorne disc emissivity (Novikov & Thorne 1973), but this is only thermalized to a blackbody in the outer disc, where $r > r_{warm}$ (all radii hereafter being dimensionless, where $r = R/R_g$). Furthermore, at $r_{warm} > r > r_{hot}$, the energy is instead emitted from the warm Compton medium above the passive disc. This warm Comptonization shape is characterized by its electron temperatures, kT_{warm} , the spectral index, Γ_{warm} , and the seed photon temperatures across this zone, which are set by the local gravitational dissipation and hard X-ray illumination. This is unlike the original model of Petrucci et al. (2018), where the seed photon temperature was a free parameter. Energy balance in a slab geometry hard wires $\Gamma_{warm} = 2.7$, so the only free parameter in the



Figure 2. Schematic of the geometry assumed here using AGNSED. The blue region within r_{hot} denotes the hard X-ray corona. The region between r_{warm} and r_{hot} consists of warm, optically thick Comptonizing material (green) sandwiching the thin, passive disc (grey). Hard X-rays irradiate the underlying thin disc, in turn changing the seed photon temperature and normalization of the soft Compton emission from the warm material. Between r_{warm} and r_{out} , only cool thermal disc emission is produced (red), found to be negligible for the *UVW1* band in Fig. 3.

warm Compton shape is kT_{warm} . Then there is a final transition to the hot coronal zone for $r_{\text{in}} > r > r_{\text{hot}}$, which has a vertical scale height h_{cor} , and which likely has no underlying optically thick disc. This component's shape is characterized by its electron temperature kT_{hot} and spectral index Γ_{hot} . The overall normalization is set by the black hole mass (fixed by either the scaling relations or direct BLR reverberation mapping), the mass accretion rate as a fraction of the Eddington rate, $\dot{m} = L/L_{\text{Edd}}$, and the black hole spin.

AGNSED also includes irradiation of the outer blackbody disc and disc-like warm Compton region by the X-ray hot inner region, so that the total flux at each radius $r > r_{hot}$ is $F_{rep}(r) + F_{grav}(r)$, where

$$F_{\rm rep}(r) = \frac{\cos(n)L_{\rm cor}}{4\pi(lR_{\rm g})^2},\tag{1}$$

where $l^2 = h_{cor}^2 + r^2$ and $cos(n) = h_{cor}/l$ (Zyċki, Done & Smith 1999). The hard coronal flux typically peaks at 100 keV, so Compton recoil is significant, meaning that reflection cannot be efficient. The maximum albedo is of order ~0.5 so at least half of the illuminating flux will heat the upper few Thompson optical depths of the outer disc and warm Compton region. In the outer disc, this gives rise to blackbody emission with temperature

$$T_{\rm eff}(r) = f_{\rm col} T_{\rm grav}(r) \left(\frac{F_{\rm rep}(r) + F_{\rm grav}(r)}{F_{\rm grav}(r)} \right)^{1/4}, \tag{2}$$

where the colour temperature correction factor $f_{\rm col}$ is fixed to unity since the temperatures here are low enough that no colour correction is required (Done et al. 2012). In the warm Compton region, the additional flux input from irradiation adds to the illumination of the underlying passive disc, so modifies the seed photon luminosity and temperature such that $T_{\rm seed}(r) = T_{\rm eff}(r)$ and the spectral index remains fixed at $\Gamma_{\rm warm} = 2.7$.

2.1 Spectral model fitting

We use AGNSED in XSPEC version 12.10.0 (Arnaud, Borkowski & Harrington 1996) to model the accretion flow, with this accretion structure sketched in Fig. 2. The standard disc (red) extends only from r_{out} to r_{warm} . From r_{warm} to r_{hot} the energy is instead released in the upper layers (green), leaving the mid-plane disc as a passive reprocessor (grey). This truncates at r_{hot} , leaving the energy to be dissipated in the hot flow (blue).

The black hole mass is fixed to $M_{\rm BH} = 4 \times 10^7 \,\rm M_{\odot}$ (as used by E17). We also fix the distance of NGC 4151 to D = 19 Mpc and assume an inclination angle of $i = 53^{\circ}$ (i.e. $\cos(i) = 0.6$). The electron temperature of the hot Compton component, kT_{hot} , is fixed at 100 keV, while its spectral index, Γ_{hot} , is free. The warm Comptonization spectral index, Γ_{warm} , is fixed to 2.7 by the reprocessing in the assumed slab geometry (Petrucci et al. 2018), while we also fix the temperature, kT_{warm} , to 0.2 keV, as observed (Gierliński & Done 2004).

For the component size scales, we fix the outer disc size at $r_{out} = 10^5$, as there is likely to be optically thick material that extends out to torus scales. This assumption matches fairly well with the observed slope of the optical spectrum in the longest wavelength archival data from the *Hubble Space Telescope (HST*; see Section 2.2), but is much larger than the self-gravity radius at $r_{sg} = 390$. None the less, we expect that material still connects from torus scales to the disc even if it is clumpy rather than smooth. The *relative* normalization of the thermal disc and two Compton components are then allowed to vary by allowing r_{hot} and r_{warm} to be free parameters, while the overall normalization of AGNSED is then specified by the (free) mass accretion rate.

However, the SED of NGC 4151 also shows complex absorption and reflection from material surrounding the AGN. Beuchert et al. (2017) perform a careful analysis of the broad-band X-ray spectrum and characterize the absorption by two neutral components, both of which are variable: one with $N_{\rm H} \sim 13-25 \times 10^{22}$ cm⁻² that partially covers ~50 per cent of the source, while the other fully covers with $N_{\rm H} \sim 6-10 \times 10^{22}$ cm⁻². A small fraction of the nuclear flux is scattered around these low ionization absorbers by gas on large scales, which also emits a mix of photoionized and collisionally ionized lines. There is also a clear reflection signature from distant material (a narrow Fe K α line core; Shu, Yaqoob & Wang 2010), although the existence of inner disc reflection is debated (e.g. Keck et al. 2015).

Our much more limited *Swift* XRT data do not allow us to constrain all of these components. We instead take our AGNSED continuum and add a single cold reflection component (PEXMON; Nandra et al. 2007) with spectral index tied to that of Γ_{hot} , and blurred by low velocities (RDBLUR; Fabian et al. 1989). We also apply two absorption components (PCFABS + PCFABS) with column densities $N_{\rm H} = 9 \times 10^{22}$ and 25×10^{22} cm⁻², and covering fractions $f_{\rm cov} = 1$ and 0.86, respectively, in order to match the complex curvature. We include a small fraction of the hot coronal emission from AGNSED to model the scattered emission at low energies, together with hot plasma emission to fit the stellar/hot gas lines (MEKAL + MEKAL; Mewe, Gronenschild & van den Oord 1985; Arnaud & Raymond 1992). All of these components are absorbed by the Galactic column (PHABS) fixed at $N_{\rm H} \sim 0.02 \times 10^{22}$ cm⁻².

Our data also extend down into the UV and optical so we also include dust extinction (ZREDDEN), with E(B - V) fixed at 0.03 from our Galaxy. We also allow for a contribution from host galaxy starlight in the UVOT apertures with an Sb template (HOSTPOL; Ezhikode et al. 2017). It is clear that there is substantial host galaxy contamination in the optical photometric bands, but there is much less in the *UVW1* (green band in Fig. 3).

The final model is hostpol + phabs*zredden* [pcfabs*pcfabs* (rdblur*pexmon + agnsed_1)

+ $agnsed_2$ + mekal + mekal]. The $agnsed_2$ term is the scattered continuum, and so contains only the hard Compton component. We fix the normalization of this scattered emission to 3 per cent of that of the main continuum, for consistency with Beuchert et al. (2017; see their table 3).

The resulting fit is shown in Fig. 3, with the outer disc (red solid line) from $r = 10^5$ –390, warm Comptonization (green solid line) from r = 390–90, and hot Comptonization (cyan solid line) from



Figure 3. Fit to data with the model hostpol + phabs*zredden*
[pcfabs*pcfabs* (rdblur*pexmon + agnsed)

+ $agnsed_2$ + mekal + mekal]. The soft and hard Compton components are denoted by the solid green and cyan lines, respectively, while the outer disc component is shown by the solid red line. The pink solid line denotes the distant reflection component. Constant components (stars at optical/UV from an Sb template, and hot gas in soft X-rays that both emits lines and scatters continuum into our line of sight) are shown in grey. The absorbed data are shown as crosses with colours denoting the bandpasses of Edelson et al. (2017; v, b, u, UVW1, UVM2, UVW2, X1-4, BAT). The green and blue strips denote the UVW1 and Swift BAT bandpasses, respectively.

r = 90-6 shown separately. This gives a mass accretion rate of $log(\dot{m}) = -1.86$, i.e. 1.4 per cent of the Eddington rate. We hereafter refer to this model and parameter set as our 'canonical' model.

The hard Compton corona within r_{hot} is fixed in this modelling to be at a standard scale height of $h_{cor} = 10$, as expected from averaging over an optically thin diffuse source with volume emissivity given by gravitational energy release (GD17). We show the effect of this assumption in Fig. 4, where we compare the canonical fit (red, solid line) with one where all parameters are the same except using $h_{cor} = r_{hot} = 90$ (black, solid line). Clearly, increasing the scale height of the corona increases the fraction of reprocessing in the optical/UV due to the larger solid angle now subtended by the disc/warm Compton region. More reprocessing will result in a stronger reverberation signal, yet there is a fairly low correlation by eye between the *UVW*1 and BAT curves in Fig. 1. We therefore keep $h_{cor} = 10$ for our canonical model, to minimize the variable component in *UVW*1 from the hard X-ray illumination.

For the same reason, we also fix the black hole spin to zero in this fit, thus pushing the ISCO and all characteristic radii as far out as possible. The effect of this is to maximize the disc reprocessing time lag and minimize the disc reprocessing amplitude. The only effect of increasing the spin would be to pull the ISCO and all transition radii inward, resulting in shorter time delays and stronger reprocessing on small scales, which is clearly disfavoured by the shapes of the UV/BAT light curves.



Figure 4. Comparison of different AGNSED parameter sets with reprocessing on (solid lines) or off (dashed lines). The red colour denotes the canonical model we use in this paper, based on the 2016 *Swift* campaign on NGC 4151, with $\log(\dot{m}) = -1.86$, $r_{out} = 10^5$, $r_{hot} = 90$, and $h_{cor} = 10$. The black colour denotes the canonical model as in the red, only now with the scale height of the corona, h_{cor} , set to the coronal outer radius, $r_{hot} = 90$, instead of the fiducial $h_{cor} = 10$.

2.2 Comparison to archival data

Archival *HST* spectra give a much better view of the optical/UV continuum. In Fig. 5, we show Space Telescope Imaging Spectrograph (STIS) spectra as a function of wavelength collected in both 1998 February (blue dots; ObsIDs: O42302070, O42302080, and O423020A0) and 2000 May (red dots; ObsIDs: O59701010, O59701020, and O59701040). All of these are taken with 0.1 arcsec slit width, so galaxy contamination is negligible, as confirmed by the variability across the entire wavelength range.

Shapovalova et al. (2008) estimate the host galaxy contribution using ground-based data at different apertures. They estimate a galactic contribution at 5100 Å of 1.1×10^{-14} erg cm⁻² s⁻¹ Å⁻¹ for their 4.2 × 19.8 arcsec² aperture, closest in area to the 5 arcsec radius circle used by E17. We scale our Sb galaxy template to this value at 5100 Å and subtract the resulting count rates from each UVOT data point resulting in the grey points and error bars in Fig. 5.

Bentz et al. (2013) instead used the *HST* images to estimate the host galaxy contribution in their $5 \times 12 \operatorname{arcsec}^2$ aperture as 1.7×10^{-14} erg cm⁻² s⁻¹ Å⁻¹ at 5100 Å. Again scaling our Sb galaxy template to this value at 5100 Å and subtracting the resulting count rates from each UVOT data point yields the green UVOT spectrum in Fig. 5.

We overlay our canonical fit (black line), which is a fairly good match to the observed slope of the STIS data in both high and low states at wavelengths longer than 5000 Å. This also matches the (grey) UVOT data for the 2016 campaign when the host galaxy continuum derived from Shapovalova et al. (2008) has been subtracted. In contrast, the magenta line shows the effect of changing the outer disc radius to the self-gravity radius. This is the same as r_{warm} , so there is no standard outer disc at all in that model, and we see that it passes through the (green) UVOT data where host subtraction has instead been applied according to the method of Bentz et al. (2013). Taking the extrema of possible host galaxy corrections for our data to be given by the Shapovalova- and Bentz-subtracted UVOT data sets (grey and green, respectively), we find that models either with or without an outer disc are permitted by

our data, although we note that the slope of the longer wavelength STIS data is better matched by models that include an outer disc. Crucially however, both models (and host-subtracted UVOT data sets) converge in the *UVW*1 band, and so the timing model tested in Section 3 will be completely unaffected by the level of host galaxy contamination and thus the presence or absence of an outer disc.

Shapovalova et al. (2008) also show the long-term 5117 Å light curve of NGC 4151. They identify low states in 2000-2002 and 2004-2006 (where their data end) where the AGN flux drops below 3×10^{-14} erg cm⁻² s⁻¹ Å⁻¹. Our data set is in this category, with a flux of $1-2 \times 10^{-14}$ erg cm⁻² s⁻¹ Å⁻¹ depending on which host galaxy subtraction is used. Fig. 1 in that paper shows a comparison of spectra from 1996 and 2005 that are rather similar to our two HST spectra in terms of their continua. These show that the broad base of the H β line has reduced by much more than the continuum flux change in these low states. While simultaneous HST spectra are not available for the 2016 Swift campaign period, fortunately there are Liverpool Telescope Fibre-fed RObotic Dual-beam Optical Spectrograph (FRODOSpec) data taken coincident with the 2016 Swift monitoring. While these cannot easily be flux calibrated, a continuum-subtracted spectrum around $H\beta$ is shown in the inset on Fig. 5 in black. The broad component is low, similarly to those seen in the low states of Shapovalova et al. (2008). Comparison with the 1998 HST F_{λ} spectrum (also shown in the inset and continuum subtracted) confirms that the 2016 H β line profile was much lower than in 1998, while the O III flux was mildly higher. The size of this change in H β could not be accounted for by a simple flux calibration, indicating that there was likely a major change in the structure of the BLR between the 1998 and 2016 epochs; this suggests that NGC 4151 may have been at the edge of - or even undergoing - a 'changing state' episode at the time of the 2016 campaign. The same can also likely be said for the 2000 epoch, which shows a very similar line profile to that in 2016. This is all most likely real variability rather than simply changes in line of sight obscuration, as Lubiński et al. (2002) show that the optical continuum correlates with the hard X-rays over these long timescales, while Kishimoto et al. (2013) also confirm that the infrared follows the optical continuum change.

3 LIGHT CURVES FROM WARM COMPTONIZATION OF DISC-REPROCESSED PHOTONS

We now take the soft Compton and hard X-ray components from the canonical model of the previous section, and follow a similar procedure to GD17 in order to extract predicted light curves in the UV. We choose here to use only the UVW1 band as there is very little host galaxy contamination in this band. Since there is also little contribution from the outer standard disc in this bandpass (see Fig. 5), we model reverberation only from the warm Comptonization zone, for clarity. Tests performed a posteriori indeed confirm that inclusion of the outer disc reverberation has no impact on the predicted light curves, because the temperature at which the emission thermalizes at large radii is too low to contribute to UVW1. In Fig. 6, the soft Compton spectra for 10 annuli spanning the full range of warm zone radii are shown as dashed lines; it is the sum of these spectra over all the passive disc/warm Compton annuli that produces the overall warm Comptonization spectrum. We now quantify how these warm Compton annuli respond to variability in the hard X-ray flux.

The IRF accounts for the light travel times to different radii and azimuthal angles within a given annulus, and so for a given



Downloaded from https://academic.oup.com/mnras/article/491/4/5126/5626354 by University of Durham user on 13 October 2022

Figure 5. Comparison of *HST* STIS data taken in 1998 February and 2000 May, with the 2016 *Swift* UVOT data that have here been corrected for host galaxy emission. We show *HST* STIS data taken on 1998 February 10 (blue dots; ObsIDs: O42302070, O42302080, and O423020A0) and on 2000 May 24 (red dots; ObsIDs: O59701010, O59701020, and O59701040). The grey and green data points and errors denote the *Swift* UVOT data with host galaxy emission subtracted according to the predictions of Shapovalova et al. (2008; grey data) and Bentz et al. (2006; green data). The black solid line denotes the canonical model fit to the 2016 UVOT data, with the host galaxy component omitted. The magenta solid line denotes the canonical model fit, but with the outer blackbody disc (as well as the host galaxy) omitted, which is more consisted with the Bentz-corrected data. Inset: Left-hand axis (black line) shows continuum-subtracted Liverpool Telescope Fibre-fed RObotic Dual-beam Optical Spectrograph (FRODOSpec) count rates (ObsID: b_e_20160422_3_2_1_2, dated 2016 April 22), coincident with the *Swift* UVOT campaign. Right-hand axis shows the H β –O III line profiles from the *HST* STIS spectra (1998 again in blue, 2000 in red), also with constant continuum subtractions for comparison of the line profiles. We see a dramatic drop in the broad H β line profile between 1998 and both 2000 and 2016, at mildly increased O III luminosity, indicating a major change in the structure of the broad-line region (BLR) between these epochs.

radius, this function describes what fraction of the reprocessed flux from that radius has time delay τ with respect to the illuminating continuum. We calculate this as in GD17 (based on Welsh & Horne 1991), giving the reprocessed flux from an annulus at radius *r* and time *t* as

$$F_{\rm rep}(r,t) = \frac{\cos(n)}{4\pi (lR_{\rm g})^2} \int_{\tau_{\rm min}}^{\tau_{\rm max}} {\rm IRF}(r,\,\tau) L_{\rm cor}(t-\tau) \,\mathrm{d}\tau. \tag{3}$$

This causes the effective temperature of the seed photons supplying the warm Comptonization in that annulus to vary as

$$T_{\text{seed}}(r, t) = f_{\text{col}} T_{\text{grav}}(r) \left(\frac{F_{\text{rep}}(r, t) + F_{\text{grav}}(r)}{F_{\text{grav}}(r)} \right)^{1/4}.$$
 (4)

We stress again that these equations assume a 'passive' underlying disc, i.e. there is no intrinsic variability in the underlying disc. The only source of variability in the warm Comptonization here is from external heating by the illuminating hard X-rays. We assume that this change in heating is followed quickly by the change in cooling via reprocessing in the passive mid-plane material, so that the spectral index and temperature of the spectrum from a given annulus do not change, while its seed photon temperature does change. By computing the seed photon temperature at each time step for each annulus, our code not only accounts for the changing normalization of the overall warm Comptonization due to the driving continuum, but also for the changing *shape* of this component as different annuli are illuminated, changing their seed photon temperatures. We therefore stress that no single radius dictates the optical/UV luminosity.

The spectral model assumes that the seed photons for the hard X-rays are from the warm Comptonization zone, so the hot corona emission extends down into (and hence directly contributes to) the UVW1 band (see Fig. 6). We include this direct variable emission, so the UVW1 light curve is the sum of the contributions in this band from the primary hard X-ray source, the constant, intrinsic



Figure 6. Model spectrum: cyan line denotes the total hard Compton emission from the corona, which illuminates the disc. The solid green line shows the total soft-Compton emission from the warm material on the disc, which reprocesses all thermal emission from the underlying thin disc. The dashed lines show the soft-Compton emission from the 10 constituent annuli in the soft-Compton region. The seed photon temperature (the blackbody temperature of the underlying disc) and normalization of these annuli are modulated by the illuminating hard X-ray emission. The green and blue bands denote the *UVW*1 and *Swift* BAT bandpasses, respectively.



Figure 7. Top panel: rebinned and interpolated *Swift* BAT light curve. Bottom panel: rebinned and interpolated *UVW1* light curve in blue. The predicted *UVW1* light curve resulting from the model assuming disc illumination driving warm Comptonization is denoted in red. The lack of agreement between the data and model is clear.

dissipation in the warm Comptonization region, and the variable reprocessed flux.

Rather than simulating the driving hard X-ray fluctuations, we use the observed *Swift* BAT light curve of NGC 4151 (shown in E17 and in Fig. 1a) as the driving light curve. We interpolate this on a grid of 0.05 d, and then resample it to produce give an even sampling of dt = 0.5 d, shown in Fig. 7(a). This light curve is fed into our disc reprocessing–warm Comptonization model, allowing us to predict the *UVW*1 light curve from our physical model of the accretion geometry.

Fig. 7(b) shows our resulting *UVW*1 light curve (red), compared to that observed (blue). Our simulated *UVW*1 light curve is very clearly not like the real *UVW*1 data, with much more fast variability than is observed. In particular, the peaks at MJDs 57460 and 57485 in the modelled *UVW*1 curve are completely absent in the observed



Figure 8. Normalized cross-correlation functions (CCFs) between the *Swift* BAT light curve and *UVW*1 observed light curve (blue), and the *Swift* BAT light curve and modelled *UVW*1 light curve (red).

*UVW*1 curve, replaced instead by a single broad peak centred at MJD 57470.

The discrepancy is even clearer in Fig. 8, showing the crosscorrelation functions (CCFs) between the hard X-ray curve and the modelled (solid red line) or observed (solid blue line) UVW1 curves. These CCFs are computed according to the procedure of GD17, which accounts for the error bar variance. It is clear that the predicted lag of ~ 1 d for the simple reprocessing/warm Comptonization model where the Comptonization takes place between 90 and $390R_{\odot}$ is far shorter than that observed, and so the model predicts far more correlation between the BAT and UVW1 light curves than observed. Instead, the data show much less correlation, with positive lags between 2 and 15 d. E17 calculates that the Swift BAT band and UVW2 bands are correlated with a significance of only 72 per cent when the red noise character of AGN variability is taken into account. We must therefore conclude that the processes generating the emission in each of these bands are, at best, weakly related, unlike the extremely strong correlation between all the UVOT bands (E17).

We test the robustness of these timing model results to the input spectral model in Appendix A. There we show an extreme model, where including additional reddening allows the intrinsic SED of these data to be fit with *no* warm Compton zone. In this case, $r_{warm} = r_{hot} = 60-70R_g$, so the UV emitting material is even closer, resulting in an even stronger and faster reverberation signal that is even more inconsistent with the observed variability. We conclude that if the UV is powered by optically thick emitting material, then it must be present at radii significantly below 1 light-day (440 R_g), irrespective of whether it emits as blackbody or warm thermal Comptonization. Yet such material in a disc-like geometry is strongly ruled out by the variability. The results presented in this section are therefore robust to the SED decomposition, as all reprocessing taking place in this class of models must be on sub-day time-scales.

3.1 The disconnect between the hard X-rays and the UV continuum

Our spectral model gives a good description of the broad-band continuum in NGC 4151 during this campaign, where the UV continuum comes from a region of warm Comptonizing material. It also predicts the *UVW*1 variability from the assumed disc geometry

and emissivity, which place this material at r = 90-390. However this picture is wholly ruled out by the observed *UVW*1 light curve, which exhibits much less of the fast variability predicted by the model in Fig. 7, and is instead dominated by variability on longer time-scales, up to 15 d (Fig. 8).

Instead, we conclude that our model of the accretion flow is wrong; it overestimates the amount of fast variability, and underestimates the slow variability. This all points to there being an additional spectral component in *UVW*1, which is slowly variable, and which replaces part of the fast-responding outer disc/warm Comptonization in the current spectral decomposition. Fig. 8 shows that there is a broad lag signature from 2 to 15 d, a time-scale that neatly encompasses the size scale of the BLR in this source, with He II measured at 2–3 light-days, H β at 6 light-days, and H α at 11 light-days (Peterson & Horne 2004), all of which are substantially smaller than the IR torus size scale of 40 light-days (Oknyansky et al. 2014).

A potential contribution from diffuse emission from the BLR was first identified by Korista & Goad (2001), and given strong impetus by the intensive monitoring campaign results that pointed to longer response time-scales than expected from the disc (E17, see their fig. 5). Similar results are also seen in the sparser data of Chelouche et al. (2019), where they derive a spectrum of the slowly reverberating component in the optical and identify this with the BLR. Lawther et al. (2018) show predictions for the diffuse BLR flux for clouds of different densities. It is mainly composed of a series of radiative recombination continua - most obviously at the Balmer and Paschen limit - and free-free continua (as was suggested by Barvainis 1987!). This non-blackbody shape is important, as a key problem with blackbody thermalization is that this gives temperatures that are too low to contribute to UVW1 for the large size scales required by the lags (GD17). Conversely, material at high enough temperature to contribute blackbody radiation to UVW1 would have too small an area for the required luminosity (GD17).

4 RECOVERING THE IMPULSE RESPONSE FUNCTION

The conclusion that there is an additional optical/UV component from the BLR means that the accretion continuum contribution in our spectral fit is overestimated. Here we attempt to separate out the contribution of this new component. We assume that the *UVW1* flux can be split into a constant component, carrying a fraction f_c of the average *UVW1* flux, a slowly variable fraction due to reprocessing from large scales, f_s , while the remaining fraction, $(1 - f_c - f_s)$, is fast varying from a direct contribution of the hard X-rays in the *UVW1* band its reprocessing on the optically thick disc/warm Compton (lagged and smoothed by the light travel time of less than 1.5 d). For simplicity we combine both of these fast variability components together. This means that the *UVW1* emission is described by

$$\frac{F_{\rm UVW1}(t)}{\langle F_{\rm UVW1}(t)\rangle} = f_{\rm c} + f_{\rm s} \, \text{IRF}(t) \circledast \frac{F_{\rm BAT}(t)}{\langle F_{\rm BAT}(t)\rangle} + (1 - f_{\rm c} - f_{\rm s}) \frac{F_{\rm BAT}(t)}{\langle F_{\rm BAT}(t)\rangle}.$$
(5)

We first assume that the slow response function is a Gaussian, so $IRF(\tau) = e^{\frac{(\tau-\mu)^2}{2\sigma^2}}$, where μ is the centroid and σ^2 is the variance. The observed slow response has a contribution up to ~ 20 d (see Fig. 8) so we calculate the IRF out to 30 d. This means that we

have to discard the initial 30 d of the *UVW*1 light curve, as it depends on the X-ray light curve prior to the monitoring campaign. We calculate the resulting truncated *UVW*1 light curve, minimizing χ^2 between our predicted light curve and the data to determine the best-fitting f_c , f_s , and IRF parameters. However for a single Gaussian IRF this gives an unacceptable fit (χ^2 /d.o.f. = 316/75), so we instead use the observed double-peaked CCF (Fig. 8) to guide our choice of a double Gaussian for the IRF. This requires three additional parameters (a relative normalization, K_2 , and the additional Gaussian centroid and variance). This now gives a good description of the light curve, with the model shown as the red line in Fig. 9(c.ii), with associated IRF in Fig. 9(a), and χ^2 /d.o.f. = 52/72.

We also test a more sophisticated IRF with more free parameters. Motivated by the shape of the observed CCF in Fig. 8, we replace the Gaussian functions with exponentially modified Gaussians, characterized by rates λ_1 and λ_2 , so that

$$IRF(\tau) = I \left[e^{\frac{\lambda_1}{2} (2\mu_1 + \lambda_1 \sigma_1^2 - 2\tau)} \operatorname{erfc} \left(\frac{\mu_1 + \lambda_1 \sigma_1^2 - \tau}{\sqrt{2}\sigma_1} \right) + K_2 e^{\frac{\lambda_2}{2} (2\mu_2 + \lambda_2 \sigma_2^2 - 2\tau)} \operatorname{erfc} \left(\frac{\mu_2 + \lambda_1 \sigma_2^2 - \tau}{\sqrt{2}\sigma_2} \right) \right].$$
(6)

Here erfc is the complementary error function, and *I* normalizes the IRF such that its integral is unity as required. This gives a better fit to the *UVW*1 light curve (IRF in Fig. 9b and green line in Fig. 9c.ii), but only by $\Delta \chi^2 = 10$ for the two additional free parameters. Overfitting the data is clearly an issue, but we show this in order to show the size of systematic uncertainty on the IRF, f_c and f_s .

To test the validity of the derived fractions and IRF, we now include 40 d of UVW1 data prior to the beginning of the 2016 campaign by extending the Swift BAT light curve back in time using the standard all sky monitoring data. Since the Swift BAT instrument is a large sky monitor (covering a quarter of the sky at any given time), we can do this without sacrificing signal-to-noise ratio at earlier times. We then interpolate within this light curve to produce an X-ray light curve on an equivalent time grid to the 2016 campaign data, facilitating the IRF fit. This extension of the BAT data to earlier times allows us to compare our predictions to the most important feature in the UVW1 light curve, namely the large rise around MJD 57460. We rerun both IRF models to fit the entire UVW1 light curve. The exponentially modified double Gaussian IRF yields $\Delta \chi^2 = 30$ better than the double Gaussian case, so we show this alone in Fig. 10. The fit now gives a good description of the main features of the entire UVW1 light curve (χ^2 /d.o.f. = 145/126).

4.1 Physical origin of the derived IRF

The IRFs we have derived in Section 4 show that the slow variability can be matched by reverberation, with 90 per cent of the UVW1 flux predominantly coming from BLR size scales of 1.5-20 light-days. Only $\sim 5-10$ per cent of the UVW1 band emission can be correlated with the hard X-rays on short time-scales. This is a factor of 3 smaller than that predicted by the model explored in Section 3, but is consistent with the 9 per cent of UVW1 flux predicted to be contributed directly by the low energy extension of the hard X-ray component. The hard Compton component shape – well constrained by the peak in Swift BAT – requires that its seed photons come from UV energies or below, and so this hard component contribution to UVW1 must be present. This leaves very little room for any UVW1



Figure 9. Modelled *UVW*1 light curve using phenomenological impulse responses to the BAT curve. Panel (a): red, solid line denotes the impulse response function (IRF) used to fit the *UVW*1 light curve, composed of a sum of two Gaussians. The dashed lines denote the two constituent Gaussians that compose this function. Panel (b): more sophisticated IRF used to fit the *UVW*1 light curve, with functional form in equation (6). Panel (c.i): rebinned and interpolated *Swift* BAT light curve. Panel (c.ii): rebinned and interpolated observed *UVW*1 light curve in blue. Modelled *UVW*1 light curves denoted in red and green, where the UV emitting component is directly modulated by the *Swift* BAT light curve convolved with either the IRF of panel (a) (red) or panel (b) (green).

contribution from warm Comptonization or thermal disc emission from within 1.5 light-days ($\leq 650R_g$). The accretion structure within 1.5 light-days is therefore far more consistent with being a pure radiatively inefficient accretion flow (RIAF; Yuan, Quataert & Narayan 2003), where seed photons are generated internally by cyclo-synchrotron emission (Veledina 2016; Inoue & Doi 2018).

We reiterate that in this picture, the lack of short (<1 light-day) response is due to there being little-to-no optically thick material on these scales. We illustrate this geometry in Fig. 11, where the inner disc has evaporated into a hot, optically thin flow that extends out to $\sim 600R_g$. This result is indeed corroborated by the very recent work of Zoghbi, Miller & Cackett (2019), who find no evidence for a relativistically broadened iron line component after accounting for absorption effects, and thus no evidence for an inner disc.

Lawther et al. (2018) performed detailed calculations of diffuse emission from the BLR. They show that while 'standard' BLR cloud densities of $n_{\rm H} \sim 10^{10}$ cm⁻³ efficiently produce the observed lines, higher density clouds emit continuum (via free–free and bound–free interactions) more efficiently than lines. They show that a population of denser clouds within the BLR produces a substantial diffuse continuum. Additional evidence for these clouds comes from BLR reverberation studies of NGC 5548. The emission line 'holiday', where the high-ionization BLR lines change their response to the continuum variability, can be explained by changes in dense clouds on the inner edge of the BLR filtering the continuum seen by the BLR (Dehghanian et al. 2019a,b). In both NGC 4151 and in NGC 5548, this material is probably the same as that which produces the time variable X-ray absorption. This is seen as a blueshifted



Figure 10. Modelled *UVW1* light curve using exponentially skewed Gaussian response to the BAT curve input, for the entirety of the pointed campaign. Panel (a): IRF used to fit the *UVW1* light curve, with functional form in equation (6). Panel (b.i): rebinned and interpolated *Swift* BAT light curve from the 2016 pointed campaign denoted in blue. BAT light curve prior to the pointed campaign is shown in cyan. The inclusion of this data, covering 40 d prior to the beginning of the 2016 campaign, allows us to make a prediction for the entire *UVW1* curve from the beginning of the pointed campaign at MJD 57435 onward. Panel (b.ii): rebinned and interpolated betweed *UVW1* light curve in blue. Modelled *UVW1* light curve denoted in green, where the UV emitting component is directly modulated by the extended *Swift* BAT light curve of panel (b.i), convolved with the IRF of panel (a).

absorption line against the broad H β emission, indicating that it is part of a dense outflow; as the H β line is not a resonance transition, absorption requires substantial collisional excitation to the n = 2level. Dehghanian et al. (2019b) show again that the ratio of continuum to lines increases as the density increases, with densities higher than $n_{\rm H} \sim 10^{12}$ cm⁻³ producing substantial optical/UV continuum (both free–free and Balmer/Paschen recombination continua; see their fig. 3). If the density becomes too high however, the emission thermalizes to a temperature of $T < L_{\rm cor}/(4\pi R^2 \sigma_{\rm SB}) \sim 3500$ K – too low to contribute to the UV band.

In this picture, the double response IRF shape we see in Fig. 9(b) does not describe two separate mechanisms, but the distributed response of partially ionized clouds, from the near and far side of the AGN. This is very similar to the geometry invoked for the clumpy molecular torus material, but on smaller scales, and we note the similarity of our IRF to those of the infrared reverberation simulations of Almeyda et al. (2017, see e.g. their fig. 7).

To check whether BLR reprocessing of hot flow emission is energetically feasible, we can consider the energy absorbed by the obscurers in the canonical SED fit. By computing the hard coronal luminosity before and after obscuration, we find that 35 per cent of the hot coronal power should be absorbed by the obscurers in the canonical SED decomposition. According to *HST* extended narrowline region studies (see e.g. Kaiser et al. 2000), the narrow-line region cone in NGC 4151 has an opening angle of approximately 60°. The BLR should fill in the remainder of the sky, so that it subtends a solid angle of around 11 sr on the hard corona, or 86 per cent of the sky from the perspective of the central engine. This would imply that $0.86 \times 35 = 30$ per cent of the hard coronal power is absorbed by the BLR, or the wind at its inner edge. Of course since the density of the obscurer is likely to be higher closer to the disc plane, this intercepted (and reprocessed) power is actually likely to be higher, but it is already sufficient to reproduce the power in the warm Compton SED component, which is $L_{warm}/L_{hard} = 27.5$ per cent of the hard Compton luminosity. There can also be additional power from reflection of the higher energy flux, so reprocessing can work energetically.

The fraction of flux that is constant is $f_c < 0.05$, according to the error ranges shown in Fig. 10. We calculate the contribution from the host galaxy in the SED model (Section 2.1). This is ~3 per cent of the *UVW1* band in the canonical SED fit, which is fully consistent with the contribution preferred by the exponentially modified Gaussian IRF fit to the final light curve ($f_c = 0^{+0.05}_{-0.00}$). A larger constant UV flux is only allowed in the models that do not include the first 40 d of data (Fig. 9), but we note that any constant UV disc contribution invoked to explain these higher levels of f_c must be accompanied by a comparable level of fast reprocessing (as established by Sections 2 and 3 and Appendix A), which is emphatically disfavoured by all IRF fits.

However, one issue with the UV continuum arising mainly from diffuse BLR emission is that standard BLR reverberation mapping works! That is, the BLR lines lag the UV continuum, in apparent contradiction to the idea that the optical/UV continuum is produced within the BLR. For instance, Bentz et al. (2006) and De Rosa et al. (2018) find a lag of ~ 6 d between the H β line and the underlying continuum in NGC 4151, implying a physical separation between the optical continuum and $H\beta$ emission zones. That said, the lines and the diffuse continuum may have very different radiation patterns and the observed lags can be produced if the line emission is produced more efficiently from the illuminated front face of the BLR clouds (the magenta photons in Fig. 11), while the diffuse continuum is predominantly produced on the shaded side. This weights the line response towards material on the far side of the black hole, while the continuum is weighted to the near side, with a smaller light travel time delay (red and green photons in Fig. 11; see e.g. Almeyda et al. 2017, comparing their figs 6 and 7). Indeed the example of a H β lag of 6 d is well within the <10 d lag margin attributable to anisotropic clump illumination according to the simulations of Lawther et al. (2018; see their fig. 7). On the other hand, if some lines produce a similar radiation pattern to the diffuse continuum from our perspective, we would predict a negligible lag for these lines relative to the optical; this is also supported by some observations. Taking the cases analysed by Peterson & Cota (1988) for instance, the combined HeII lag measured by Antonucci & Cohen (1983, -1 ± 5 d), and the Mg II and C IV lags measured by Ulrich et al. (1984, 12 ± 20 d, 5 ± 8 d) are all consistent with zero, as one may expect if these lines are produced cospatially and aligned with the diffuse continuum.

4.2 Revisiting NGC 5548

GD17 instead proposed that the warm Comptonization region is vertically extended (see their fig. 7b) so that it shields the outer disc from the hard X-ray illumination. This geometry appeared feasible as the models were based on OPTXAGNF that uses only a single warm Compton component rather than considering its radial structure as



Figure 11. The geometry we propose to explain both the observed SED and derived impulse response functions (IRFs). The blue zones denote the radiatively inefficient accretion flow that emits strongly in the hard X-rays. The red/white circles denote the broad-line region (BLR) clumps that reprocess the hard X-ray emission and on which line emission is excited. The grey zones denote outer disc material from which the reprocessor clouds may be launched via e.g. a dust-driven wind. The existence of this disc material down to $\sim 500R_g$ has been suggested by Miller et al. (2018) from the velocity width of the Fe K α line. Wiggly lines denote photons in the optical (red), UV (green), and hard X-rays (blue) bands, and line emission (purple) that is preferentially produced on the illuminated (white) side of the BLR clumps.

we do here. In this case our Fig. 6 shows that – while the warm Compton annuli are more luminous at smaller radii – UVW1 has equal contributions from the all radii. The inner radii could shield the outer radii from the hard X-rays, but we would still see the fast reprocessed variability from the inner warm Compton radii. Self-shielding by the warm Comptonization zone in NGC 4151 cannot hide the reprocessed signal in our model. Similar considerations show that (with hindsight) this could not work in NGC 5548 either.

None the less, in NGC 5548 the requirement for X-ray shielding was not so compelling as the X-ray bandpass only extended up to 10 keV, and was clearly affected by time variable absorption (Mehdipour et al. 2016). Thus it was always possible in this object that the observed X-ray variability did not track the variability of the bolometric flux, which peaks at 100 keV, and is unaffected by obscuration. The lack of a linear IRF between these signals supports this idea (see Appendix B). Thus in NGC 5548 the variable *UVW*1 and optical emission could both be from reprocessing on a disc/warm Compton region of the (unseen) 100 keV X-ray flux. However, the lag between the optical/UV behind the *HST* far-UV was also too long by a factor of 2.

GD17 explicitly tested models for the BLR contribution to the lags, and showed that this did not work to explain the UV–optical lag. However, they used a much longer lag time of 15 light-days – as expected for the H β BLR size scale given the observed luminosity – but not consistent with the actual BLR lag of a few days measured during the NGC 5548 campaign (Pei et al. 2017, their fig. 13; De Rosa et al. 2018). They also assumed the spectrum of the BLR was dominated by the observed spectral features, rather than being a diffuse continuum from higher density clouds, and thus they underestimated the contribution of the reprocessed BLR flux. Their reverse engineered transfer functions showed that the lag of the UV and optical required material at a few light-days, consistent with the observed BLR size scale in these data. However, for blackbody

(rather than diffuse) emission, the thermalization temperature is too low, which caused them to discount this possibility.

In retrospect it seems more likely that the NGC 5548 monitoring campaign was sampling an AGN with a similar geometry to that shown in Fig. 11; it was at a similarly low Eddington luminosity of $\sim 0.03L_{Edd}$, and its broad H β line is almost as weak as that in NGC 4151 (De Rosa et al. 2018).

5 CONCLUSIONS

We have built a full spectral-timing model to predict the UV light curve from the observed hard X-ray light curve in NGC 4151. The model is based on a (truncated) disc and warm Comptonization region illuminated by the hard X-ray source. Energetically, the X-ray luminosity requires all of the gravitational energy released within the inner $90R_g$, whereas the UV can be produced by the same accretion rate through the truncated disc/warm Compton region at larger radii, between 90 and $390R_g$. We use this geometric/energetic model to predict the light curve in the UVW1 band arising from illumination of the warm Compton/thin disc structure by the coronal hard X-rays. The changing hard X-ray luminosity drives variations in heating of the warm Comptonizing material above the disc, as well as variations in its seed photons due to reprocessing. We show that such a geometry entirely fails to reproduce observed UVW1 variability from the observed hard X-ray (15-50 keV) variability. In particular, a disc geometry strongly overpredicts the coherence between these signals and underpredicts the lag. Optically thick material within a few hundred R_g should respond on time-scales of less than 1-2 d and be smoothed on this time-scale by the light travel time across the region. This predicts that the UV and X-ray light curves should be strongly correlated on time-scales shorter than 1.5 d, yet the data show only a weak correlation. The Swift BAT band traces the bolometric luminosity of the source (unlike the 'hard' X-ray band of the Swift XRT that can be affected by the

Instead, we find that the UV can be produced by reprocessed X-ray flux in material on size scales of the BLR. This material cannot be emitting blackbody radiation as the resulting temperatures would be too low at this size scale. However, dense clumps (n > n) 10^{12-15} cm⁻³) that are optically thin rather than thick could produce a diffuse recombination continuum rather than lines or blackbody emission, as suggested by Korista & Goad (2001) and Lawther et al. (2018). Such material is seen in the line of sight as the complex, mainly neutral, X-ray absorber and is also imprinted on the optical emission lines as the H β absorption. Following Dehghanian et al. (2019a,b), we identify this with a wind on the inner edge of the BLR. We suggest that this is also the inner edge of the optically thick accretion disc, so that inwards of this there is only the hot, optically thin gas from which the hard X-rays originate. This interpretation is supported by *Chandra* grating data that suggest that the 'narrow' Fe K α line is produced by thermal disc material extending no further in than $500R_g$ at this luminosity (Miller et al. 2018), as well as by the recent work of Zoghbi et al. (2019), who find no evidence of relativistic Fe K α line broadening associated with an inner disc. Higher spectral resolution observations with the X-ray Imaging and Spectroscopy Mission (XRISM; JAXA/NASA, probable launch 2022) and the Advanced Telescope for High-Energy Astrophysics (ATHENA; ESA, probable launch 2030) will undoubtedly reveal this structure in more detail.

ACKNOWLEDGEMENTS

The authors thank the anonymous referee for their very helpful comments and suggestions that served to improve the paper. The authors thank Rick Edelson for initiating the awesome monitoring campaigns, and more specifically for the NGC 4151 light curves and spectral data. The authors also thank Bożena Czerny, Swayamtrupta Panda, Gary Ferland, and many attendees of the 2019 Guilin Conference on Mapping the Central Regions of Active Galactic Nuclei for illuminating conversations about the BLR and reverberation mapping. RDM acknowledges the support of a Science and Technology Facilities Council (STFC) studentship through grant ST/N50404X/1. CD acknowledges the STFC through grant ST/P000541/1 for support.

REFERENCES

- Almeyda T., Robinson A., Richmond M., Vazquez B., Nikutta R., 2017, ApJ, 841, 1
- Antonucci R. R. J., Cohen R. D., 1983, ApJ, 271, 564
- Arévalo P., Uttley P., Lira P., Breedt E., McHardy I. M., Churazov E., 2009, MNRAS, 397, 2004
- Arnaud M., Raymond J., 1992, ApJ, 398, 394
- Arnaud K., Borkowski K. J., Harrington J. P., 1996, ApJ, 462, L75
- Baron D., Stern J., Poznanski D., Netzer H., 2016, ApJ, 832, 8
- Barvainis R., 1987, ApJ, 320, 537
- Bentz M. C. et al., 2006, ApJ, 651, 775
- Bentz M. C. et al., 2013, ApJ, 767, 149
- Beuchert T. et al., 2017, A&A, 603, A50
- Boissay R., Ricci C., Paltani S., 2016, A&A, 588, A70
- Buisson D. J. K. et al., 2018, MNRAS, 475, 2306
- Chelouche D., Nuñez F. P., Kaspi S., 2019, Nat. Astron., 3, 251
- Davis S. W., Woo J.-H., Blaes O. M., 2007, ApJ, 668, 682

- Dehghanian M. et al., 2019a, ApJ, 877, 119
- Dehghanian M. et al., 2019b, ApJ, 882, L30
- De Rosa G. et al., 2018, ApJ, 866, 133
- Done C., Gierliński M., Kubota A., 2007, A&AR, 15, 1
- Done C., Davis S. W., Jin C., Blaes O., Ward M., 2012, MNRAS, 420, 1848
- Edelson R. et al., 2015, ApJ, 806, 129
- Edelson R. et al., 2017, ApJ, 840, 41
- Ezhikode S. H., Gandhi P., Done C., Ward M., Dewangan G. C., Misra R., Philip N. S., 2017, MNRAS, 472, 3492
- Fabian A. C., Rees M. J., Stella L., White N. E., 1989, MNRAS, 238, 729
- Frank J., King A., Raine D. J., 2002, Accretion Power in Astrophysics, 3rd edn. Cambridge Univ. Press, Cambridge
- Galeev A., Rosner R., Vaiana G., 1979, ApJ, 229, 318
- García J. A. et al., 2019, ApJ, 871, 88
- Gardner E., Done C., 2017, MNRAS, 470, 3591
- Gierliński M., Done C., 2004, MNRAS, 349, L7
- Haardt F., Maraschi L., 1993, ApJ, 413, 507
- Inoue Y., Doi A., 2018, ApJ, 869, 114
- Jin C., Ward M., Done C., Gelbord J., 2012, MNRAS, 420, 1825
- Kaiser M. E. et al., 2000, ApJ, 528, 260
- Keck M. L. et al., 2015, ApJ, 806, 149
- Kishimoto M. et al., 2013, ApJ, 775, L36
- Korista K. T., Goad M. R., 2001, ApJ, 553, 695
- Kubota A., Done C., 2018, MNRAS, 480, 1247
- Laor A., Davis S. W., 2014, MNRAS, 438, 3042
- Lawrence A., 2018, Nat. Astron., 2, 102
- Lawther D., Goad M. R., Korista K. T., Ulrich O., Vestergaard M., 2018, MNRAS, 481, 533
- Lubiński P., Zdziarski A. A., Walter R., Paltani S., Beckmann V., Soldi S., Ferrigno C., Courvoisier T. J.-L., 2010, MNRAS, 408, 1851
- Magdziarz P., Blaes O. M., Zdziarski A. A., Johnson W. N., Smith D. A., 1998, MNRAS, 301, 179
- Matt G. et al., 2014, MNRAS, 439, 3016
- Mehdipour M. et al., 2011, A&A, 534, A39
- Mehdipour M. et al., 2015, A&A, 575, A22
- Mehdipour M. et al., 2016, A&A, 588, A139
- Mewe R., Gronenschild E. H. B. M., van den Oord G. H. J., 1985, A&AS, 62, 197
- Miller J. M. et al., 2018, ApJ, 865, 97
- Nandra K., O'Neill P. M., George I. M., Reeves J. N., 2007, MNRAS, 382, 194
- Noda H., Done C., 2018, MNRAS, 480, 3898
- Novikov I. D., Thorne K. S., 1973, in DeWitt C., DeWitt B., eds, Black Holes (Les Astres Occlus). Gordon & Breach, New York, p. 343
- Oknyansky V. L., Metlova N. V., Taranova O. G., Shenavrin V. I., Artamonov B. P., Gaskell C. M., Guo D.-F., 2014, Odessa Astron. Publ., 27, 47
- Pei L. et al., 2017, ApJ, 837, 131
- Penston M. V., Perez E., 1984, MNRAS, 211, 33
- Peterson B. M., Cota S. A., 1988, ApJ, 330, 111
- Peterson B. M., Horne K., 2004, Astron. Nachr., 325, 248
- Petrucci P.-O. et al., 2013, A&A, 549, A73
- Petrucci P. O., Ursini F., De Rosa A., Bianchi S., Cappi M., Matt G., Dadina M., Malzac J., 2018, A&A, 611, A59
- Porquet D., Reeves J. N., O'Brien P., Brinkmann W., 2004, A&A, 422, 85
- Porquet D. et al., 2018, A&A, 609, A42
- Richards G. T. et al., 2003, AJ, 126, 1131
- Ruan J. J., Anderson S. F., Eracleous M., Green P. J., Haggard D., MacLeod C. L., Runnoe J. C., Sobolewska M. A., 2019, ApJ, 883, 76
- Shakura N. I., Sunyaev R. A., 1973, A&A, 24, 337
- Shapovalova A. I. et al., 2008, A&A, 486, 99
- Shu X. W., Yaqoob T., Wang J. X., 2010, ApJS, 187, 581
- Shull M. J., Stevans M., Danforth C. W., 2012, ApJ, 752, 162
- Telfer R. C., Zheng W., Kriss G. A., Davidsen A. F., 2002, ApJ, 565, 772
- Ulrich M. H. et al., 1984, MNRAS, 206, 221
- Veledina A., 2016, ApJ, 832, 181
- Welsh W. F., Horne K., 1991, ApJ, 379, 586

Yuan F., Quataert E., Narayan R., 2003, ApJ, 598, 301

Zheng W., Kriss G. A., Telfer R. C., Grimes J. P., Davidsen A. F., 1997, ApJ, 475, 469

Zoghbi A., Miller J., Cackett E., 2019, ApJ, 884, 26

Zyċki P. T., Done C., Smith D. A., 1999, MNRAS, 305, 231

APPENDIX A: PHYSICAL MODELLING WITH NO SOFT EXCESS REGION

In this appendix we highlight that the modelled *UVW*1 curve we recovered in Section 3 is insensitive to the form of accretion disc assumed at small radii, such that even if the disc within 1 light-day does not have a warm Compton atmosphere, the fact that it must dominate the UV band means that the resultant *UVW*1 curve is still strongly dissimilar to the data.

If our data have undergone more complex reddening than we anticipated in Section 2, it becomes possible that there is no warm Compton component at all, and the intrinsic SED is actually dominated by only a thermal Shakura–Sunyaev disc and a hard Compton component. By including a second, multiplicative ZREDDEN component in addition to the other components listed in Section 2.1, and tying r_{warm} to r_{hot} , we find that we can in fact fit the spectrum with the same quality as the canonical model, now omitting the warm Compton zone. We show the intrinsic SED (not including the complex absorption, host galaxy dilution, and reddening) in Fig. A1, panel (a.i), where we have fixed the coronal height to $h_{cor} = 10$, and the fit inner disc radius is $60R_g$. We see in panel (a.iii) that the predicted light curve from this geometry is nothing like that which is observed. The disparity becomes far more pronounced when the coronal height is increased to $h_{cor} = 100$ in panels (b.i–iii), where the variability amplitude and shape are now entirely wrong.

Combined with the results of Sections 2 and 3, we conclude with a high degree of certainty that, as long as optically thick material at sub-light-day size scales dominates the UV band – whether that material is thermal or weakly Comptonized – these geometries will fail to correctly predict the UV light curve, and so should be ruled out.



Figure A1. Panel (a.i) shows the intrinsic SED for NGC 4151 in the case that there is no warm Compton zone, and the spectrum is only composed of a thermal disc illuminated by a central hard Compton source at a height of $h_{cor} = 10$ above the black hole. Constituent thermal disc annuli are shown as dashed lines. Like the canonical model, this geometry also produces a very poor match to the true (blue) *UVW*1 curve in panel (a.iii), where we see that the modelled (red) curve has a lot more of the hard Compton variability on short time-scales than we find in the data. Panels (b.i–iii) show the same model set-up only now with the central source at $h_{cor} = 100$, showing an even greater disparity with the *UVW*1 data in panel (b.iii).

APPENDIX B: IMPULSE RESPONSE RECOVERY IN NGC 5548

Here we apply the impulse response recovery method of Section 4 to the 2014 *UVW*1 and *Swift* XRT light curves of NGC 5548 (Edelson et al. 2015), originally investigated with physical modelling efforts in GD17. In that work it was found that the sharply peaked variations in the 0.8–10 keV band and lack thereof in the *UVW*1 band resulted in the simple disc reverberation model failing to replicate the *UVW*1 curve from the hard X-ray input. We therefore test whether the phenomenological model used here can recover the *UVW*1 curve (and associated impulse response) when the driving signal is associated with the 0.8–10 keV band in NGC 5548.

We assume integrated hard Compton-to-disc contributions in the *UVW1* band using the SED of fig. 3(a) of GD17, and that the 0.8–10 keV band is contributed to exclusively by the hard Compton component. Equation (5) therefore applies. We also prescribe the

exponentially modified double-Gaussian functional form for the impulse response for maximum fit freedom, so that equation (6) applies.

We find that this technique is unable to reproduce the *UVW1* light curve from the hard X-ray input, as shown by the highly contrasting *UVW1* light curves in Fig. B1(b.ii). No matter the shape of the IRF, the highly peaked hard X-ray light curve cannot be smoothed sufficiently to reproduce the *UVW1* light curve without inducing a strong lag, which the observed light curves do not appear to exhibit. We therefore end up with a model *UVW1* curve that exhibits only a short lag as required, but with insufficient smoothing. Furthermore, other features in the observed *UVW1* curve (e.g. the peak at 115 d) have absolutely no counterpart in the hard X-ray light curve, adding to the weight of evidence that suggests the sources of emission in these bands do not continuously communicate in NGC 5548, seemingly unlike in NGC 4151.



Figure B1. Modelled *UVW*1 light curve using phenomenological impulse responses to the *Swift* XRT X3 curve for the NGC 5548 data of Gardner & Done (2017; from Edelson et al. 2015). Panel (a): impulse response function (IRF) used to fit the *UVW*1 light curve, with functional form in equation (6). Panel (b.i): interpolated *Swift* XRT X3 light curve. Panel (b.ii): interpolated observed *UVW*1 light curve denoted in blue. Modelled *UVW*1 light curves denoted in green, where the UV emitting component(s) is directly modulated by the *Swift* XRT X3 light curve convolved with the IRF of panel (a). Panel (c): cross-correlation functions (CCFs) between *Swift* XRT X3 and *UVW*1 light curves. The observed CCF is denoted in blue. The modelled CCF resulting from the IRF in panel (a) is denoted in green.

This paper has been typeset from a TEX/IATEX file prepared by the author.