

Modelling continuum reverberation in active galactic nuclei: a spectral-timing analysis of the ultraviolet variability through X-ray reverberation in Fairall 9

Scott Hagen   and Chris Done

Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK

Accepted 2023 February 13. Received 2023 January 17; in original form 2022 October 7

ABSTRACT

Continuum reverberation mapping of active galactic nuclei (AGNs) can provide new insight into the nature and geometry of the accretion flow. Some of the X-rays from the central corona irradiating the disc are absorbed, increasing the local disc temperature. This gives an additional reprocessed contribution to the spectral energy distribution (SED) which is lagged and smeared relative to the driving X-ray light curve. We directly calculate this reverberation from the accretion disc, creating fully time-dependent SEDs for a given X-ray light curve. We apply this to recent intensive monitoring data on Fairall 9, and find that it is not possible to produce the observed ultraviolet (UV) variability by X-ray reprocessing of the observed light curve from the disc. Instead, we find that the majority of the variability must be intrinsic to the UV emission process, adding to evidence from changing-look AGN that this region has a structure which is quite unlike a Shakura–Sunyaev disc. We filter out this long time-scale variability and find that reprocessing alone is still insufficient to explain even the fast variability in our assumed geometry of a central source illuminating a flat disc. The amplitude of reprocessing can be increased by any vertical structure such as the broad-line region and/or an inner disc wind, giving a better match. Fundamentally though the model is missing the major contributor to the variability, intrinsic to the UV/EUV emission rather than arising from reprocessing.

Key words: accretion, accretion discs – black hole physics – galaxies: active – galaxies: individual: Fairall 9.

1. INTRODUCTION

Active galactic nuclei (AGNs) are powered by accretion on to a supermassive black hole. This accretion flow is generally interpreted in the context of a Shakura & Sunyaev (1973) disc, where energy dissipated through the flow gives rise to a radial temperature profile, increasing as the disc approaches the black hole. This leads to a spectral energy distribution (SED) composed of a sum of blackbodies, which for AGN peaks in the ultraviolet (UV).

However, observations paint a more complex picture than a simple disc. Firstly, AGN spectra are always accompanied by a high energy X-ray tail (Elvis et al. 1994) originating from the innermost regions of the accretion flow. Additionally there is also a soft X-ray excess, observed as an upturn below 1 keV which is remarkably constant in shape across a broad range of AGN (e.g. Gierliński & Done 2004; Porquet et al. 2004). Finally, the UV is often redder than expected from a standard disc, with a turnover again at a remarkably constant energy (e.g. Laor & Davis 2014). Clearly then, AGN SEDs are more complex than a simple disc.

One way to match the data is to assume that the accretion flow is radially stratified, where the energy only thermalizes to the standard Shakura & Sunyaev (1973) blackbody temperature at large radii, and inside of this the accretion energy instead emerges as warm or

hot Comptonized emission [Done et al. 2012; Kubota & Done 2018 (hereafter KD18)]. This is designed to allow the models enough flexibility to match the broad-band SED, with enough constraints from energetics to fit the data without too much degeneracy.

We can use spectral-timing to test this idea of a radial transition in the flow. Observations of AGN display strong X-ray variability, which can be used to constrain the nature and geometry of the accretion flow (e.g. Uttley et al. 2014). In particular, illumination of the disc (blackbody or warm Compton emitting) by the fast variable hot Compton X-ray flux leads to a fast variable reprocessed component, which correlates with but lags behind the X-rays by a light traveltime. This reverberation mapping was originally proposed by Blandford & McKee (1982) as a means to measure the size scale of the broad-emission line region (BLR) in AGN. The lines respond to changes in the ionizing (UV to soft X-ray) light curve in a way that can be described as a convolution of the driving light curve with a transfer function which contains all the light-traveltime delays from the specific geometry. Observing the driving light curve and its delayed and smoothed emission line response gives information on the size scale and geometry of the BLR. This is often simply condensed to a single number, which is the mean lag revealed in cross-correlation (e.g. Welsh & Horne 1991; Peterson 1993; Horne et al. 2004; Peterson et al. 2004).

However, the technique of reverberation mapping is quite general, and can be used to map continuum components as well as emission lines. Recently, there has been a drive for intensive multiwave-

* E-mail: scott.hagen@durham.ac.uk

length monitoring campaigns of AGN to map the accretion disc geometry from observations of the reprocessed continuum emission produced by the fast variable X-ray source illuminating the disc (e.g. McHardy et al. 2014, 2018; Edelson et al. 2015, 2019). In particular, the use of space telescopes (especially *Swift*) in these campaigns has allowed high-quality and near-continuous monitoring over extended periods of time, with simultaneous data of both the (assumed) driving hard X-ray and the disc reprocessed UV/optical emission.

The results from these continuum reverberation campaigns are very surprising. In general, the UV variability is not well correlated with the X-ray variability which is meant to be its main driver. Instead, all wavebands longer than the far-UV correlate well with the far-UV light curve, but on a time-scale which is longer than expected from a standard disc (e.g. the compilation of Edelson et al. 2019).

Here, instead of working backwards to a geometry from cross-correlation time lags, we work forward from a geometry given by the new SED models. Crucially, as well as predicting the smoothing/lag from light-traveltime delays, this also allows us to predict the *amplitude* of the response, giving a predicted UV light curve which can be compared point by point to the observed UV light curve.

This approach was first used by Gardner & Done (2017) (hereafter GD17) to model the light curves in NGC 5548, but here the strong extrinsic X-ray variability from an unusual obscuration event made comparison difficult (e.g. Mehdipour et al. 2016; Dehghanian et al. 2019a). NGC 4151 gave a cleaner comparison as although this also shows strongly variable absorption, it is bright enough to be monitored by the *Swift* BAT instrument, sensitive to the higher energy 20–50 keV flux which is unaffected by the absorption. This showed clearly that there was a radial transition in the flow, with no reverberating material within 500–1000 R_g (Mahmoud & Done 2020). Such a hole in the optically thick material in this object is also now consistent with the X-ray iron line profile (Miller et al. 2018) and its reverberation (Zoghbi, Miller & Cackett 2019), despite previous claims to the contrary (e.g. Zoghbi et al. 2012; Cackett et al. 2014; Keck et al. 2015). The new SED models had indeed predicted truncation of the optically thick disc material, though on somewhat smaller size scales of 50–100 R_g . The size of the response also showed that there was substantial contribution to the reprocessed flux from above the disc plane. This is almost certainly the same material as is seen in the variable absorption, which is clearly a wind launched on the inner edge of the BLR (Kaastra et al. 2014; Chelouche, Pozo Nuñez & Kaspi 2019; Dehghanian et al. 2019b). An additional diffuse UV contribution from X-ray illumination of this wind/BLR (Korista & Goad 2001; Lawther et al. 2018) means that the SED model fits to the data overestimated the UV from the accretion flow itself, which probably led to the underestimate of the truncation radius in NGC4151.

In this paper, we follow the approach of GD17 and Mahmoud & Done (2020), but for Fairall 9 (hereafter F9). This explores a very different part of parameter space in terms of mass accretion rate. Both NGC 4151 and NGC 5548 are at 0.01–0.03 L_{Edd} , close to the changing state transition so the disc should be quite strongly truncated, with a hot flow replacing the inner disc (Noda & Done 2018; Ruan et al. 2019). By contrast, F9 has $L \sim 0.1L_{\text{Edd}}$, so should have much more inner disc. F9 is also most likely an almost face on AGN as it shows very little line-of-sight absorption from either cold or ionized material (bare AGN: Patrick et al. 2011). This means that the X-ray light curve is more likely representative of the intrinsic variability, rather than being heavily contaminated by extrinsic absorption variability (though we note there are occasional obscuration dips: Lohfink et al. 2016).

We construct a full spectral-timing code, AGNVAR, which predicts variability at any wavelength from the new SED models (AGNSED in XSPEC). We make this publicly available as a PYTHON module,¹ and apply it to the recent intensive monitoring data on F9 (Hernández Santisteban et al. 2020).

We describe the model in Section 2. The underlying SED is necessary to understand the overall energetics: A source where the X-ray luminosity is as large as the UV luminosity can give a much stronger UV response to a factor 2 change in X-ray flux than a source like F9 where the X-ray power is 10x smaller than the UV (KD18). The radial size scale of the transition regions in the disc is likewise set by the SED fits, which determines the response light traveltime. We outline our method for evolving the SED along with the light curve, and explore how our model system responds to changes in X-ray illumination. We apply this to the mean SED of F9, and form the full time-dependent light curves from reprocessing of the observed X-rays in Section 4. Our model fails to describe the data, most clearly as it produces much less variability amplitude than observed for any reasonable scale height of the X-ray source. This is clearly a consequence of simple energy arguments from the SED as the UV luminosity is considerably larger than the X-ray luminosity, so even a factor 2 change in X-ray flux has only limited impact on the UV, especial in the geometry assumed here of a central source illuminating a flat, truncated disc. This clearly shows that most of the variability in the UV is intrinsic (assuming the observed X-rays are isotropic), which is at odds with standard Shakura–Sunyaev disc models, as these should only be able to vary on a viscous time-scale (e.g. Lawrence 2018). Our work highlights the lack of understanding of the structures which emit most of the accretion power in AGN.

2. MODELLING THE RESPONSE OF THE ACCRETION FLOW

Throughout our analysis we fix the black hole mass and distance to $M = 2 \cdot 10^8 M_{\odot}$, $d = 200$ Mpc (Bentz & Katz 2015), and assume an inclination angle of $\cos(i) = 0.9$. We will also adopt the standard notation for radii, where r is the dimensionless gravitational radius, and R is the physical distance from the black hole. These are related by $R = rR_G$, where $R_G = GM/c^2$.

2.1 The intrinsic SED

We start by considering the intrinsic disc emission. We will follow the AGNSED model described in KD18, and give a brief summary here.

The accretion flow is assumed to be radially stratified, but with a standard Novikov–Thorne emissivity profile $\epsilon_{\text{NT}}(R)$ (Novikov & Thorne 1973; Page & Thorne 1974). We divide the flow into annuli of width ΔR , which emits luminosity $L(R)\Delta R = 2 \times 2\pi R\epsilon_{\text{NT}}\Delta R$ (where the factor 2 is for each side of the disc). Each disc annulus emits as a blackbody for $R_{\text{out}} < R < R_w$, with temperature $T_{\text{NT}} = (\epsilon_{\text{NT}}/\sigma_{\text{SB}})^{1/4}$ (Shakura & Sunyaev 1973).

Inwards of this, for $R_w < R < R_h$, thermalization is incomplete, leading to the luminosity instead being emitted as warm, optically thick Comptonization. Following KD18, we assume that this warm Comptonization region overlies a passive (non-emitting) disc (Petrucci et al. 2018), so its seed photons are set by reprocessing the warm Comptonization power on the passive disc material. Unlike the models of Petrucci et al. (2018), this sets the seed photon temperature

¹<https://github.com/scotthgn/AGNvar>

self-consistently from the size scale, which is just the same as the expected disc temperature, $T_{\text{NT}}(R)$. This means that the seed photon temperature changes in the warm Comptonization region from R_w to R_h , again unlike the Petrucci et al. (2018) models, where it is only a single temperature blackbody.

We model the warm Compton spectrum using NTHCOMP (Zdziarski, Johnson & Magdziarz 1996; Zycki, Done & Smith 1999), normalizing the output luminosity of each annulus to the intrinsic disc luminosity of the annulus. The seed photon temperature is set as above, leaving two additional parameters: the electron temperature, $kT_{e,w}$, setting the high-energy roll over of the resulting spectral component, and the photon index Γ_w , which determines the spectral slope. Reprocessing of optically thick, warm Comptonization from a passive disc gives $\Gamma_w = 2.5\text{--}2.7$ (Petrucci et al. 2018).

This still does not explain the full SED, as AGN spectra are always accompanied by a high energy X-ray tail (e.g. Elvis et al. 1994). As in KD18, we consider this emission originates from the innermost accretion flow, where the disc has evaporated into an optically thin, geometrically thick corona (e.g. Narayan & Yi 1995; Liu et al. 1999; Zdziarski, Lubiński & Smith 1999; Rózańska & Czerny 2000) extending down to the innermost stable circular orbit, R_{isco} . We still assume that this is heated by the Novikov–Thorne emissivity, giving L_{diss} , but with seed photons from the inner edge of the warm Comptonization region, so that the total X-ray luminosity is increased by this seed photon contribution giving total power $L_x = L_{\text{diss}} + L_{\text{seed}}$ (see KD18 for details on calculating these). The model assumes that the inner disc is replaced by the hot Comptonizing plasma, so unlike the warm Comptonization region, there is no underlying disc to provide a source for these seed photons. Instead, we assume that the seed photons originate from the truncated disc beyond R_h . We assume the seed photons come mainly from the inner edge of the truncated disc, so if there is a warm Comptonization region the seed temperature is $T_{\text{NT}}(R_h)\exp(y_w)$, where y_w is the Compton y -parameter for the warm Comptonization region. If there is no warm Comptonization region, then the seed temperature is simply the temperature of the inner disc, $T_{\text{NT}}(R_h)$. Again, like the warm region, we model this with NTHCOMP and leave the electron temperature, $kT_{e,h}$, and photon index, Γ_h , as parameters within the code. This sets the spectral shape of the warm Compton region, which we then normalize to the total luminosity L_x .

We now have a model, identical to the one given in KD18, consisting of three regions, with absolute size scale set by the black hole mass via R_g . A standard outer disc, located between r_{out} and r_w , a warm Comptonization region, where the disc emission fails to thermalize, located between r_w and r_h , and a hot Comptonizing corona replacing the optically thick disc from r_h to r_{isco} . This geometry is sketched in Fig. 1a. The total SED is then the total contribution from each region added together, and forms the basis of our spectral timing model.

2.2 Contribution from reprocessing

So far we have only considered the intrinsic emission from the accretion flow. However, a portion of the photons emitted by the corona will be incident on the disc, and a fraction of these will be absorbed and re-emitted. This will give a contribution to the local temperature at a point on the disc $\propto F_{\text{rep}}^{1/4}$, which is dependent on both the geometry of the disc (e.g. Zycki et al. 1999; Hartnoll & Blackman 2000), and the corona. In our current picture, we consider the inner corona to be an extended sphere, with luminosity per unit volume which depends on radius. The flux at a given point on the disc from illumination requires integrating the diffuse emission over the entire

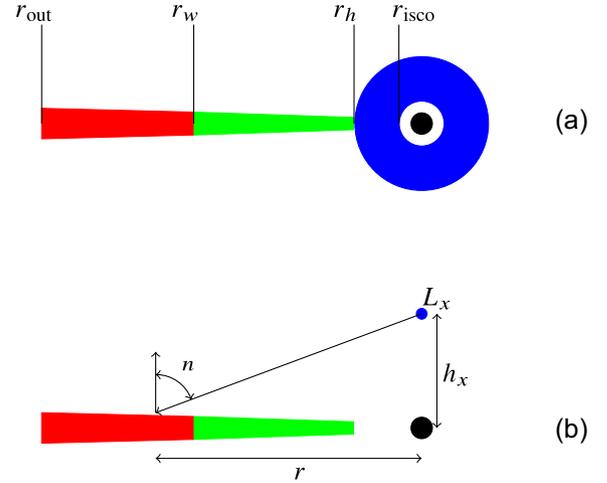


Figure 1. (a): Schematic of the model geometry considered in this paper. Between r_{out} and r_w we have a standard accretion disc (red region), emitting like a multicolour blackbody. From r_w to r_h we have the warm Comptonization region (green), where the disc fails to thermalize leading to Comptonization of the underlying disc photons. Finally, between r_h and r_{isco} the disc has evaporated into the hot Comptonization region, which we consider as a spherical corona (blue). (b): Schematic of the geometry we use to calculate the reprocessed emission. Here, we approximate the spherical corona as a point source located a height h_x above the spin axis; following GD17.

corona. Repeating this for the entire disc is clearly computationally expensive. However, GD17 showed that there is little difference in the illumination pattern between a radially extended corona powered by Novikov–Thorne emissivity and a point source located a height $h_x \sim 10$ above the black hole. This removes the need for the expensive integration, and also makes the calculation of time delays in the next section much simpler, giving the picture illustrated in Fig. 1b. The flux seen by a point on a disc then takes the simple form (Zycki et al. 1999)

$$F_{\text{rep}}(r) = \frac{L_x \cos(n)}{4\pi(r^2 + h_x^2)R_G^2} \quad (1)$$

where $\cos(n) = h_x/(r^2 + h_x^2)^{1/2}$ is the angle between the incident ray on the disc and the surface normal, and L_x is the coronal luminosity. The effective temperature at a given radius will then be

$$\sigma T_{\text{eff}}^4(r) = \sigma T_{\text{NT}}^4(r) + F_{\text{rep}}(r)(1 - A) \quad (2)$$

where σ is the Stefan–Boltzman constant, and A is the disc albedo. As in KD18 we fix the albedo to 0.3. We note that unlike KD18 we use L_x rather than $0.5 L_x$ for the luminosity seen by the disc. This is because we allow the disc to be heated from both sides. Due to symmetry, this is the same as using only one side for the geometry but letting it see the full X-ray luminosity.

2.3 Time-dependent reprocessing

We now extend the reprocessed contribution to the SED into a time-dependent SED model by considering the light-traveltime between the X-ray source, accretion disc, and observer as in GD17 and Mahmoud & Done (2020).

Here, we follow the method in Welsh & Horne (1991). The direct coronal emission will have a shorter path-length to the observer, than the reprocessed emission that first has to travel via the disc. This is illustrated in Fig. 2, where we can see that the indirect emission has a travel path increased by $l_1 + l_2$ with respect to the direct emission,

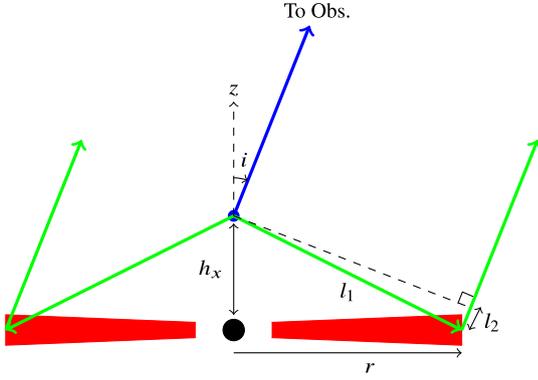


Figure 2. Illustration of the geometry considered when calculating the time delay between the direct and reprocessed emission. For simplicity we use a lamppost geometry. The direct coronal emission is shown by the blue line, while the indirect emission via the disc is given in green. The labels l_1 and l_2 indicate the portions of the indirect travel path that contribute to the time delay.

which leads to the indirect emission lagging behind the direct by a time delay $\tau = (l_1 + l_2)/c$. Clearly l_1 and l_2 depend on the disc position being considered, hence we can rewrite our path difference in terms of the disc coordinates r and ϕ . The result is given in equation (3), which is similar to that given in Welsh & Horne (1991), however with the addition of the coronal height h_x .

$$\tau(r, \phi) = \frac{R_G}{c} \left\{ \sqrt{r^2 + h_x^2} + h_x \cos(i) - r \cos(\phi) \sin(i) \right\} \quad (3)$$

We define a grid across the accretion disc, with default spacing $d \log r = 0.01$ and $d\phi = 0.01$ rad, and use equation (3) to construct time-delay surfaces; or iso-delay surfaces. These delay surfaces are used to map an observed X-ray light curve $F_{x, \text{obs}}(t)$ on to the disc. For each time t within the light curve, a point (r, ϕ) on the disc will see the X-ray luminosity from time $t - \tau$. Assuming the X-ray coronal luminosity varies exactly like the observed light curve, such that $L_x(t)/\langle L_x(t) \rangle = F_{x, \text{obs}}/F_{x, \text{obs}}$, then we have that the disc temperature will vary as

$$\sigma T_{\text{eff}}^4(r, \phi, t) = \sigma T_{\text{NT}}^4(r) + \frac{\cos(n)(1 - A)}{4\pi(r^2 + h_x^2)R_G^2} L_x(t - \tau(r, \phi)). \quad (4)$$

To calculate the time-dependent SEDs, we start by calculating the disc temperature within each grid-point across the disc, for every time-step in the input light curve. In the interest of computational efficiency, and since $\Delta R \ll c\Delta t$ across the extent of our accretion disc, we flatten the two-dimensional (2D) grid into a radial grid by calculating the mean temperature within each annulus; based off the grid points within that annulus. The SED for each time-step is then calculated following Section 2, resulting in a series of model SEDs covering the duration of the light curve.

We then extract light curves in a given energy band by defining a midpoint energy, E_{mid} , and a bin width, dE . The model light curve is simply the integral of the SED flux density within the energy bin centred on E_{mid} at each time-step t . This can be directly compared to the observed fluxes. However, we choose instead to show light curves in terms of the mean normalized flux $F/\langle F \rangle$ so that both data and model are dimensionless.

This all assumes that the effects of general relativity are small, unlike the code of Dovčiak et al. (2022) as used for reverberation by Kammoun, Papadakis & Dovčiak (2019) and Kammoun et al. (2021). A fully general relativistic treatment is required for very small corona

height; however, Kammoun et al. (2021) show that these corrections are negligible for a large coronal height.

2.4 Understanding the disc response

We illustrate the model by considering a short X-ray flash, with amplitude $L_{x, \text{max}}/\langle L_x \rangle = 2$, modelled by a top hat with width $\Delta t = 0.2$ d. We consider how this propagates down through the SED, as the flash travels across the disc. We fix the mass and inclination to that of F9, set the outer edge of the disc at $r_{\text{out}} = 10^3$ and assume zero spin and a mass accretion rate which is 10 per cent of Eddington. We assume for simplicity that there is no warm Compton region, so the standard disc extends from r_{out} to r_h , and take $r_h = 12$ and 33 in order to see how increasing truncation of the inner disc changes both the spectrum and the response.

The left-hand panel of Fig. 3 shows snapshots of the model for a face on observer. The illuminated ring propagates outwards, with times $t = 0.5$ d (blue), 1 d (orange), 2 d (green), 4 d (red), and 8 d (magenta) after the X-ray flash. The middle panel shows the SED (upper), with the change in spectrum at each time (lower). The spatial disc response is plotted on a log scale, so the constant width 0.2 light day traveltime of the step function is progressively smaller on the log scale at larger radii. This then also explains why the amplitude of the fluctuation drops as the step function sweeps across the disc. Both X-ray irradiation and intrinsic flux go as $F \propto r^{-3}$, so the steady-state disc temperature goes as $T \propto r^{-3/4}$ i.e. drops by a factor 2 for every factor 2.5 increase in radius. The disc can be approximated by a series of rings, each with temperature a factor 2 lower for radius increasing by a factor 2.5. The innermost ring is completely illuminated by the fixed width 0.2 light day flash, so it responds to the entire factor 2 increase in X-ray flux. However, at the outer radius, the light flash ring only covers a small fraction of the lowest temperature emitting region, so the amplitude of the response is much smaller. Thus, the largest amplitude reverberation signal is always expected on the inner edge of the truncated disc, and the change in disc SED at all energies is dominated by the inner disc. This explains the shape of the light curves in *UVW2* (magenta) and *UVW1* (green) shown in the right-hand panel. These both peak on a time-scale corresponding to the light traveltime to the inner edge of the disc of $[(r^2 + h_x^2)^{1/2} + h_x]R_g/c = 0.29$ d for $r_h = 12$. The decay is the exponential, with a time-scale roughly given by the time-scale at which the illuminating flash reaches the radius with temperature which peaks in each waveband (~ 6 eV for *UVW2* and ~ 4 eV for *UVW1*). It is this exponential decay rather than the peak response which encodes the expected increasing time-scale behaviour from a Shakura–Sunyaev disc, where $\tau \propto \lambda^{4/3}$ (Collier et al. 1999; Cackett, Horne & Winkler 2007), so that the decay in *UVW2* is ~ 0.7 d, while that in *UVW1* is 1.1 d.

The lower panel of Fig. 3 shows the effect of increasing the disc truncation radius to $r_h = 33$. Here, the light traveltime to the inner edge of the disc is 0.5 d, so the blue ring showing the position of the flash on the disc after 0.5 d does not exist. Other differences are the SED (middle panel) shows a stronger hard X-ray tail, as expected as the higher r_h means a larger fraction of the accretion power is dissipated in the hot Compton component. This power is taken from the inner edge of the disc, so the disc SED peaks at lower energies as well as being less luminous. This shifts the predicted SED peak from being in the unobservable FUV, highlighted by the pink shaded region, for $r_h = 12$, to emerging into the observable *UV* bands as shown in the lower panels of Fig. 3 for $r_h = 33$.

The stronger hard X-ray flux for $r_h = 33$ means that a factor 2 change gives a stronger response compared to the $r_h = 12$ for each

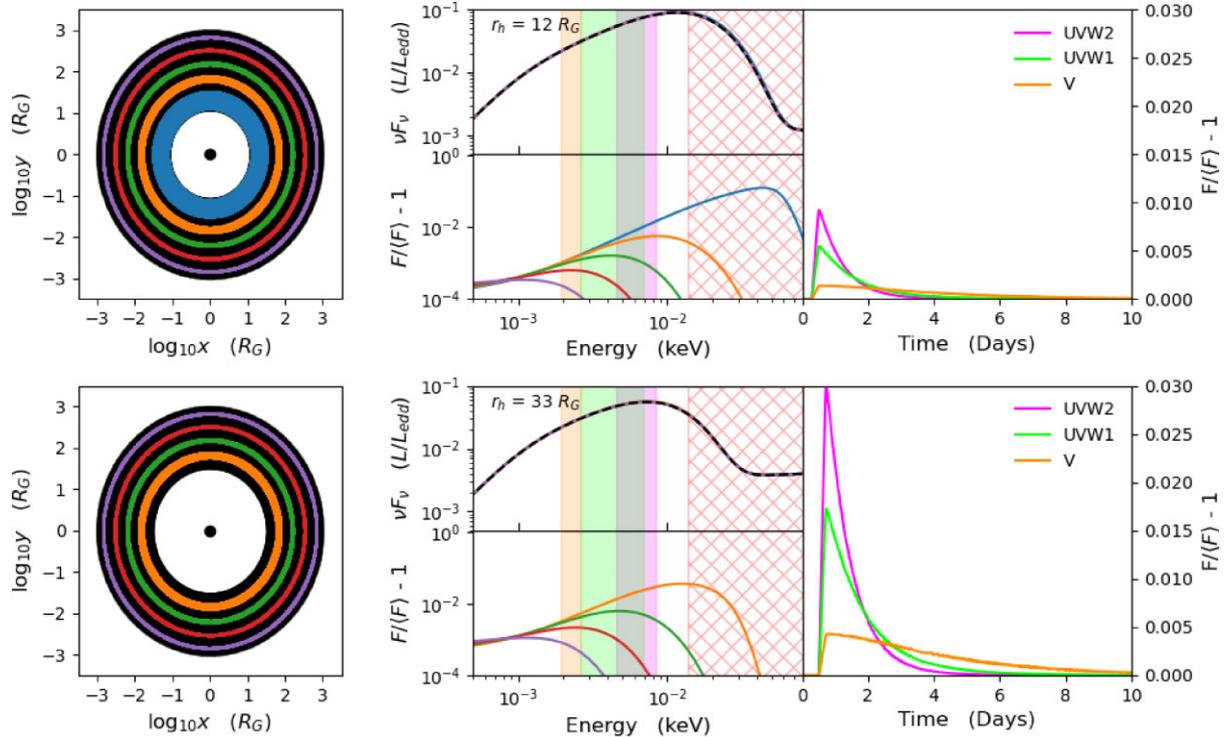


Figure 3. SED snapshots as an X-ray flash propagates across an accretion disc, with truncation radius $r_h = 12$ and 33 (top and bottom rows, respectively). The left-hand column shows the position of the X-ray flash, as seen by the disc, at each time (after the initial flash) we extract SED snapshots for. These times are 0.5 d (blue), 1 d (orange), 2 d (green), 4 d (red), and 8 d (purple). The middle column shows (top panel) each SED snapshot overlaid on top of the mean SED (black dashed line) and (bottom panel) the residual of the SED snapshot with respect to the mean. The red crosshatch region in this column indicates the unobservable part of the spectrum. The right-most column shows the responses, extracted for *UVW2* (magenta), *UVW1* (lime green), and *V* (orange) bands (assuming the same width as the *Swift-UVOT* filters). These bandpasses are also highlighted in the middle column as the shaded coloured regions, with the colours corresponding to the responses in the right-hand column.

time segment. While the underlying disc temperature is the same, the stronger illuminating flux gives a higher temperature fluctuation at 1-d delay (orange line, middle lower panels, compare between $r_h = 12$ and 33).

The position of the truncated disc edge also depends on the black hole mass and accretion rate as well as r_h , but generally the *UVW2* continuum does not sample the SED peak, even with a moderately truncated disc as assumed here. However, the light-curve response does. The light curve in any wavelength band on the disc is dominated by the contribution at that wavelength from the inner edge. The light curve in any wavelength band has a peak response lagged by the light traveltime to the inner disc edge, and then has an exponential decay whose time-scale encodes the expected $\lambda^{4/3}$ dependence.

Inclination increases the light-traveltime smearing as expected from $\Delta\tau(r) \approx R/c(1 - \sin(i))$ for the near side of the disc, to $\approx R/c(1 + \sin(i))$ for the far side. Fig. 4 compares the *UVW2*-, *UVW1*-, and *V*-band response for a face on disc with that for $i = 0, 30$, and 60 deg, with $r_h = 12$. As expected, we see that an increased inclination puts an additional smearing on the response function, increasing its width. Additionally it is also seen that as we increase the inclination the beginning of the response is shifted to earlier times by $\Delta\tau = (R_g/c)[h_x(1 - \cos(i)) + r\sin(i)]$ with respect to $i = 0$ deg, corresponding to $\Delta\tau \approx 0.08, 0.17$ d for $i = 30, 60$ deg, respectively. This is simply a reduction in the light traveltime to the inner edge of the disc on the side of the observer.

Since the responses peak at the same lag for varying energy bands, we would expect their cross-correlation to peak at zero.

Additionally, since they have a roughly exponential shape, we would also expect the cross-correlation to be an asymmetric function, where the asymmetry comes from the difference in decay times. This is shown in Fig. 5a for a cross-correlation between the light curves in *UVW2* and *UVW1* with $r_h = 12$ (black) and 33 (blue), at inclinations 0 (solid) and 30 deg (dashed). These show asymmetric functions as expected, with a reduction in the decay at higher inclination, arising from an increase in the width of the response function.

Clearly then the observed inter-band lags in AGN cannot arise from differences in the peak response time, as there is no difference in peak response time. Instead the increased response width at lower energies will lead to an increased mean delay, which gives the observed $\tau \propto \lambda^{4/3}$ lag relation (e.g. Edelson et al. 2019; Cackett et al. 2020; Vincentelli et al. 2021). Fig. 5b shows the mean centroid delay predicted by our response functions in solid green. These mostly follow the analytic $\tau \propto \lambda^{4/3}$ relation (black dashed line), apart from at higher energies originating close to the peak of the SED.

Our model can predict both the amount and the shape of the response at a given energy, through explicitly considering the energetics and geometry of the system and using these to calculate a set of time-dependent SEDs. These reproduce the analytically predicted lag-energy relation when the energies considered are away from the optically thick disc peak. We now use this model to generate the light curves expected in any band, with smoothing caused by both the distribution of light traveltimes to any given disc annulus and the continuum nature of the response.

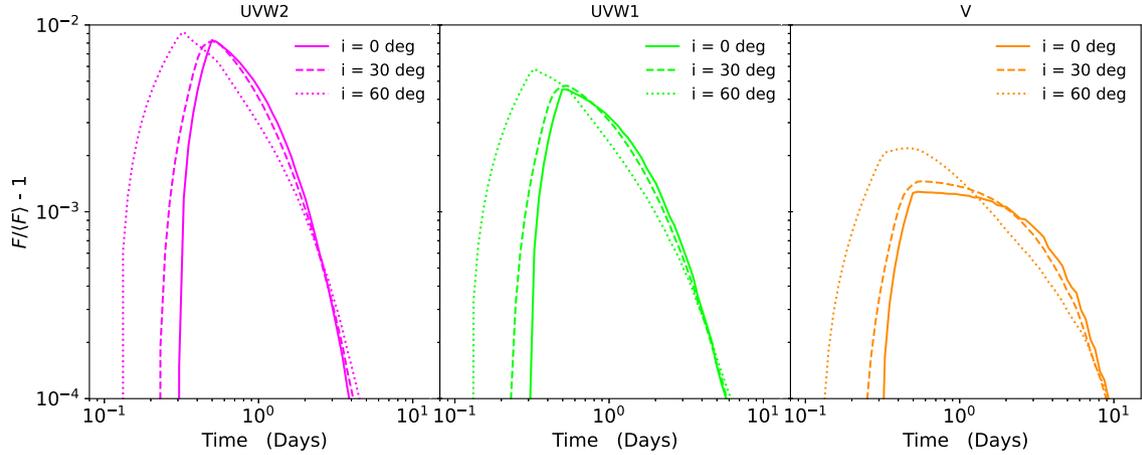


Figure 4. Response functions extracted for *UVW2* (left, magenta), *UVW1* (middle, lime green), and *V* (right, orange), as seen by an observer at inclination 0 deg (solid), 30 deg (dashed), and 60 deg (dotted). These have been calculated for a disc with $r_h = 12$.

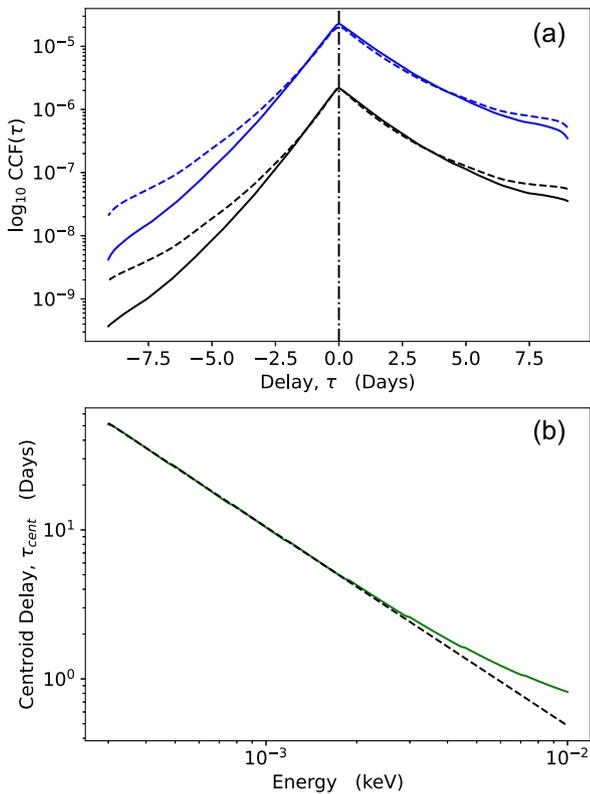


Figure 5. Top panel: The cross-correlation between the *UVW2* and *UVW1* response functions. Blue lines show the cross-correlation function (CCF) for $r_h = 33$, while black lines show for $r_h = 12$. Solid lines show inclination 0 deg, while dashed show inclination 30 deg. The dashed-dotted vertical line indicates 0 delay, where we clearly see that all the CCFs peak. Bottom panel: The centroid delay predicted by our model (green), and the analytic $\tau \propto \lambda^{4/3}$ relation (black, dashed line). Note, to avoid the model lag being affected by the outer edge of the disc, at low energies, we have increased the disc size to 10^5 for the purpose of calculating τ_{cent} .

3. THE DATA

3.1 The light curves

F9 has recently been the subject of an intensive monitoring campaign, using *Swift* and Las Cumbres Observatory (Hernández Santisteban et al. 2020). We use year 1 light curves obtained by *Swift* (provided by Jaun V. Hernández Santisteban and Rick Edelson through private communication). These light curves cover the *Swift*-*XRT* (Burrows et al. 2005) hard X-ray band (1.5–10 keV, henceforth *HX*), soft X-ray band (0.3–1.5 keV), and the *Swift*-*UVOT* (Roming et al. 2005) broad-band filters *UVW2*, *UVM2*, *UVW1*, *U*, *B*, and *V*. They have a cadence of ~ 1 d, and a duration of 300 d (from MJD 58250 to 58550). Since we expect the disc emission to peak around the *UVW2* band, and since this has the cleanest data, we will limit our analysis to the *HX* and *UVW2* light curves. A detailed description of the data reduction is given in Hernández Santisteban et al. (2020).

The light curves begin to rise beyond ~ 58500 MJD, which we speculate might be due to a change in the accretion structure. In order to simplify our analysis we therefore discard the section of the light curves beyond 58500 MJD, as a change in accretion structure would significantly complicate our results. Instead this will be the focus of a future paper. (Note, the full light curves can be found in Hernández Santisteban et al. 2020.)

In later sections, we will compare light curves using cross-correlation methods. We will also be evaluating their Fourier transforms. These techniques all rely on evenly sampled data (Gaskell & Peterson 1987; Uttley et al. 2014, GD17). Since real data will not be exactly evenly sampled, we linearly interpolate the raw light curves on to identical grids with width $\Delta t = 0.5$ d, and then re-bin on to a grid with $\Delta t = 1$ d.

3.2 Extracting the SED

We now extract the time-averaged SED for F9, using the spectral fitting package *XSPEC* v.12.12.0 (Arnaud 1996), and model the data using *AGNSED* (KD18); as described in Section 2. We stress that we use a slightly modified version of *AGNSED*, which includes heating from both sides of the disc when determining the reprocessed temperature. This then gives us constraints on both the energetics and physical parameters of the system, which we will use as our base model when analysing the light curves.

We start by extracting the time-averaged X-ray spectrum. *Swift-XRT* has limited effective area for spectroscopy, but *NICER* was also monitoring the source at this time. However, *NICER* has its own challenges regarding the background estimation. Hence instead we use *XMM-Newton* for a detailed spectral description. The archival observation on 2014 May 9 by Lohfink et al. (2016) has soft X-ray spectrum which is compatible with *NICER* and harder X-ray spectrum compatible with *Swift-XRT*. It also has *UVW1* flux from the OM within ~ 5 per cent of the *Swift-UVOT UVW1* flux, confirming that the source was in similar state at this time. We give more details in Appendix D.

We also use the UV continuum in our spectral modelling, as it is in this energy range where the disc should contribute most. We start by considering the time-averaged, host galaxy subtracted flux from each *UVOT* filter during the campaign, given in Hernández Santisteban et al. (2020). We use the conversion factors given in table 10 in Poole et al. (2008) in order to convert to a count-rate, allowing the use of XSPEC in the fitting process.

Finally, we model the intrinsic SED following Section 2, by using an updated version of AGNSED (KD18). In addition we include a reflection component, modelled with PEXMON (Magdziarz & Zdziarski 1995; Nandra et al. 2007), to model the Fe-K α line and Compton hump, along with RDBLUR (Fabian et al. 1989) to account for any smearing in the reflection spectrum. The detailed fits of Yaqoob et al. (2016) show that the iron line and reflection hump in this source are consistent with neutral material, corroborating our choice of reflection model. We also include a global photoelectric absorption component, PHABS, to account for galactic absorption. The final model is PHABS*(AGNSED + RDBLUR*PEXMON). Fig. 6 shows the final SED model after correcting for absorption.

While fitting the SED we freeze $kT_{e,h}$ to 100 keV, as we do not have sufficiently high energy coverage to clearly determine the electron temperature. We also find that the data suggest strong preference to a large warm Comptonized region, leading to a negligible contribution from the standard disc region. Hence, we simply set $r_w = r_{out}$, as this does not alter the fit statistic, and allows us to eliminate a free parameter. We also fix the galactic absorption column at $N_H = 0.035 \times 10^{22} \text{ cm}^{-2}$. The best-fitting parameters are shown in Table 1, and the SED is shown in Fig. 6. This forms our baseline model for the following analysis.

While the soft Comptonization gives a different spectrum from each disc annulus than the blackbody assumed in Section 3, its seed photons are assumed to come from reprocessing on an underlying passive disc structure. The soft spectral index means that the warm Comptonized emission peaks at an energy which tracks the seed photon energy, so it has the same $\tau \propto \lambda^{4/3}$ behaviour as a pure blackbody disc. Fig. C1 in the Appendix shows the equivalent of Fig. 3 for this specific model.

4. COMPARING MODEL AND OBSERVED LIGHT CURVES

4.1 Comparison to the unfiltered variability

We now construct model *UVW2* light curves for F9, and compare to the observations, working on the assumption that the observed *HX* light curve drives all of the variability in the UV. The variations in coronal luminosity are then the same as variations in the *HX* light curve, as described in Section 2.3. To generate model light curves from the resulting time-dependent SEDs, we use the *Swift-UVOT UVW2* response matrix (Roming et al. 2005) to extract the part of

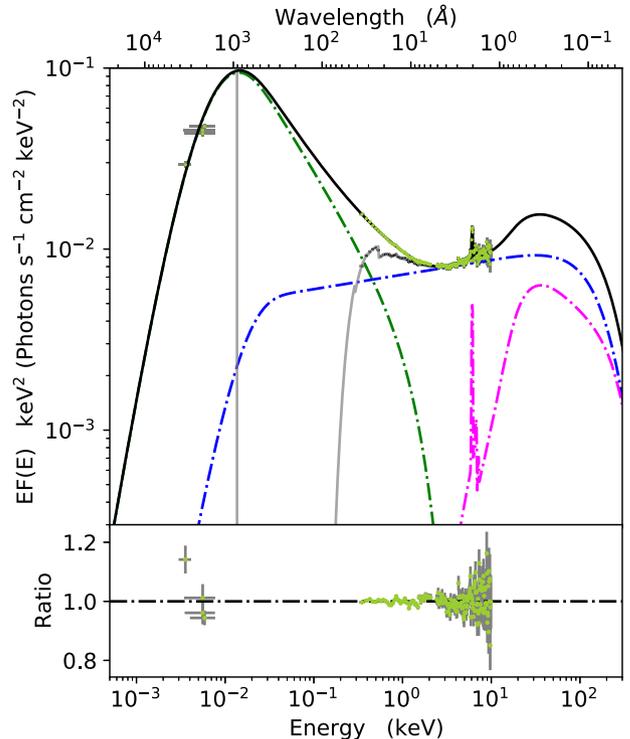


Figure 6. Time-averaged absorption corrected SED of F9. The solid black line shows the best-fitting spectrum, the dashed lines show the model components; which are: warm Comptonized region (green), hot Comptonized corona (blue), and neutral reflection (magenta); and the solid grey line shows the model before correcting for absorption. The residual shows the ratio between the best-fitting model and the data.

the SED in the correct energy range and account for the energy-dependent sensitivity across the filter. It is important to note that the fluxes in our driving X-ray light curve have errors, which need to be propagated through the model. To do this we take inspiration from the flux-randomization method, often used in determining the uncertainty on cross-correlation lags (Peterson et al. 1998). For each data point in the X-ray light curve we assign a Gaussian probability distribution, centred on the measured flux value and with a standard deviation set by the error bar. We then draw 5000 realizations of the X-ray light curve, with fluxes sampled from their probability distribution, and evaluate our model light curve for each realization. For each time-step within the input light curve we then have a distribution of 5000 model evaluations. Our set of model light-curve fluxes are then defined by the 50th percentile in each of these distributions, with the 16th and 84th percentiles defining the standard deviation for each model point.

Our resulting model light curve is shown in the bottom panel of Fig. 7, along with the observed ones. Clearly the model is not consistent with observations. The amplitude of variability in *UVW2* is dramatically underestimated.

The SED we have derived for F9 is clearly disc-dominated, with the UV disc component responsible for ~ 77 per cent of the total power, compared to the X-ray corona only contributing ~ 23 per cent. This tells us that when we calculate the effective temperature across the disc we are strongly dominated by the intrinsic disc emission, and so any changes in X-ray illumination will have a minimal effect on the SED. Essentially, there is not enough power in the variable X-ray light curve to reproduce the observed variability amplitude in the UV through reprocessing alone. In fact, it is clear that the majority of the

Table 1. Best-fitting parameters for our SED model. Values with no error were frozen in the fitting process. Note that the inner radii given in AGNSED and RDBLUR are not the same. This is because we only include RDBLUR in order to fit the iron line profile. However, since the main focus of this paper is continuum reverberation this is only a convenience component. Hence, we will only consider the AGNSED parameters in our modelling.

Component	Parameter (unit)	Value
PHABS	N_{H} (10^{20} cm^{-2})	3.5
AGNSED	M (M_{\odot})	2×10^8
	dist (Mpc)	200
	$\log(\dot{m})$ ($\dot{M}/\dot{M}_{\text{edd}}$)	$-1.215^{+0.024}_{-0.027}$
	a_*	0
	$\cos(i)$	0.9
	$kT_{\text{e,h}}$ (keV)	100
	$kT_{\text{e,w}}$ (keV)	$0.331^{+0.042}_{-0.035}$
	Γ_{h}	$1.921^{+0.026}_{-0.027}$
	Γ_{w}	$2.782^{+0.030}_{-0.032}$
	r_{h} (R_{G})	$18.8^{+1.1}_{-0.9}$
	r_{w} (R_{G})	$= r_{\text{out}}$
	r_{out} (R_{G})	$2.091^{+0.075}_{-0.058}$
	h_{x} (R_{G})	10
redshift	0.045	
RDBLUR	Index	-3
	r_{in} (R_{G})	382^{+492}_{-160}
	r_{out} (R_{G})	10^6
	Inc (deg)	25
PEXMON	Γ	$= \Gamma_{\text{h}}$
	E_{c} (keV)	10^4
	redshift	0.045
	Inc (deg)	25
	Norm (10^{-3})	$4.51^{+0.76}_{-0.68}$
$\chi^2/\text{degree of freedom}$	242.77/168 = 1.45	

variability must be driven by some other process than reprocessing of the assumed isotropic X-ray emission. The case for an alternative process is reinforced when we examine the long-term trend in the observed light curves. It is clear that the long-time-scale variations in the *UVW2* are not present in the X-ray. In fact, this was pointed out by Hernández Santisteban et al. (2020), who de-trended their light curves by fitting a parabola. If all of the *UVW2* variability was driven by reprocessing from the X-ray corona, then we would expect to see the long-term trend in the *HX* light curve too. There is simply no way to create this trend purely through reprocessing.

The inability of reprocessing to match observations becomes even more apparent when we examine the ACFs of the light curves. The *HX* ACF appears to contain two distinct components, one narrow due to the rapid variability, and one broad arising from longer term variations. *UVW2*, on the other hand, has a much broader ACF, indicating that the majority of the *UVW2* variability comes from longer term fluctuation than seen even in the longest time-scale in the hard X-rays. If disc reprocessing was the sole driver of the *UVW2* variability, then we would expect our model ACF to match the *UVW2* ACF. Clearly this is not the case. The model ACF is almost identical to that of the *HX* ACF, albeit with a slight broadening at the bottom of the narrow component. If disc reprocessing was giving a significant contribution to the total *UVW2* variability we could expect that the model ACF should at least be somewhere in between that of *HX* and *UVW2*. Again, this is not the case.

To really highlight the lack of impact of reprocessing in making the *UVW2* light curve, we show the CCFs. The *HX* and *UVW2* are poorly correlated. More interestingly, perhaps, is the correlation between our model and the *HX* light curve. The similarity between this CCF and their respective ACFs suggests that the time-scales over which the model light curve is smeared are far too small to make a significant impact. This is unsurprising when we examine the width of the narrow component in the CCF, and their ACFs, which we expect would be the first thing that would be smoothed out. We already know, from Section 2.4, that the majority of the response will come from the inner edge of the disc, which for the SED fit derived in Section 3 is at roughly ~ 5 –6 light-hour. This is considerably shorter than the width of the narrow component in the CCF/ACFs, which is on the time-scale of a few days. Although we also expect an increase in smoothing due to the continuum nature of the response, it is not expected that this would increase the time-scale sufficiently to wipe out the rapid variations. In other words, the rapid variability seen in the X-ray light curve is on time-scales longer than the smoothing imposed by the reprocessing model. Hence, this rapid variability must also be present in our model light curve, which explains the near identical nature of the model and *HX* ACFs, along with the strong similarity between these ACFs and their CCF.

These results clearly show that AGN continuum variability is somewhat more complex than can be described by reprocessor models. However, that is not to say that reprocessing does not take place, just that it cannot be the dominant source of variability. Hence, if we wish to study the continuum reprocessor, we need a way of extracting and disentangling the different sources of variations from observed light curves; the focus of the next subsection.

4.2 Isolating and exploring the short-term variability

When trying to isolate the reprocessed variability within the data, we first need some idea of what time-scales we expect the reprocessing to occur over. The clear choice, from examining the light curves, is to filter out the long-term variations in the *UVW2*, as these are not present in either the *HX* light curve or its ACF. If all variability was driven by reprocessing of X-rays, then we would expect to see the large trends within the *UVW2* light curve in the *HX* one too. However, that is clearly not the case, as we see no indication of the long ~ 100 -d trend, which is so clearly present in *UVW2*, in *HX*. This also fits nicely into our issue with the energetics not producing enough response. The *UVW2* variability is clearly dominated by this trend, and so filtering it out should reduce the variability amplitude.

Previous studies have removed long-term trends by fitting a function that roughly matches the observed shape of the trend-line, e.g. the parabola of Hernández Santisteban et al. (2020). This is useful in the sense that it provides a simple method for estimating the long-term time-scales involved, and for recovering the shape of the short-term variability. However, it is dependent on the function chosen fit to the trend, and does not give an insight into the variability power produced on different time-scales. The latter point is very important, as we have seen in the previous sections that the system energetics determine the amplitude of the reprocessed variability. So in order to constrain both the variable power and variability time-scales within the data we use a Fourier-based approach to de-trend our light curves.

Extracting information about variability on different time-scales through Fourier analysis is already commonly done in accretion studies, mostly for rapidly varying objects such as accreting white dwarfs, neutron stars, or black hole binaries (see Uttley et al. 2014 for a review); but are also increasingly common in AGN studies (e.g.

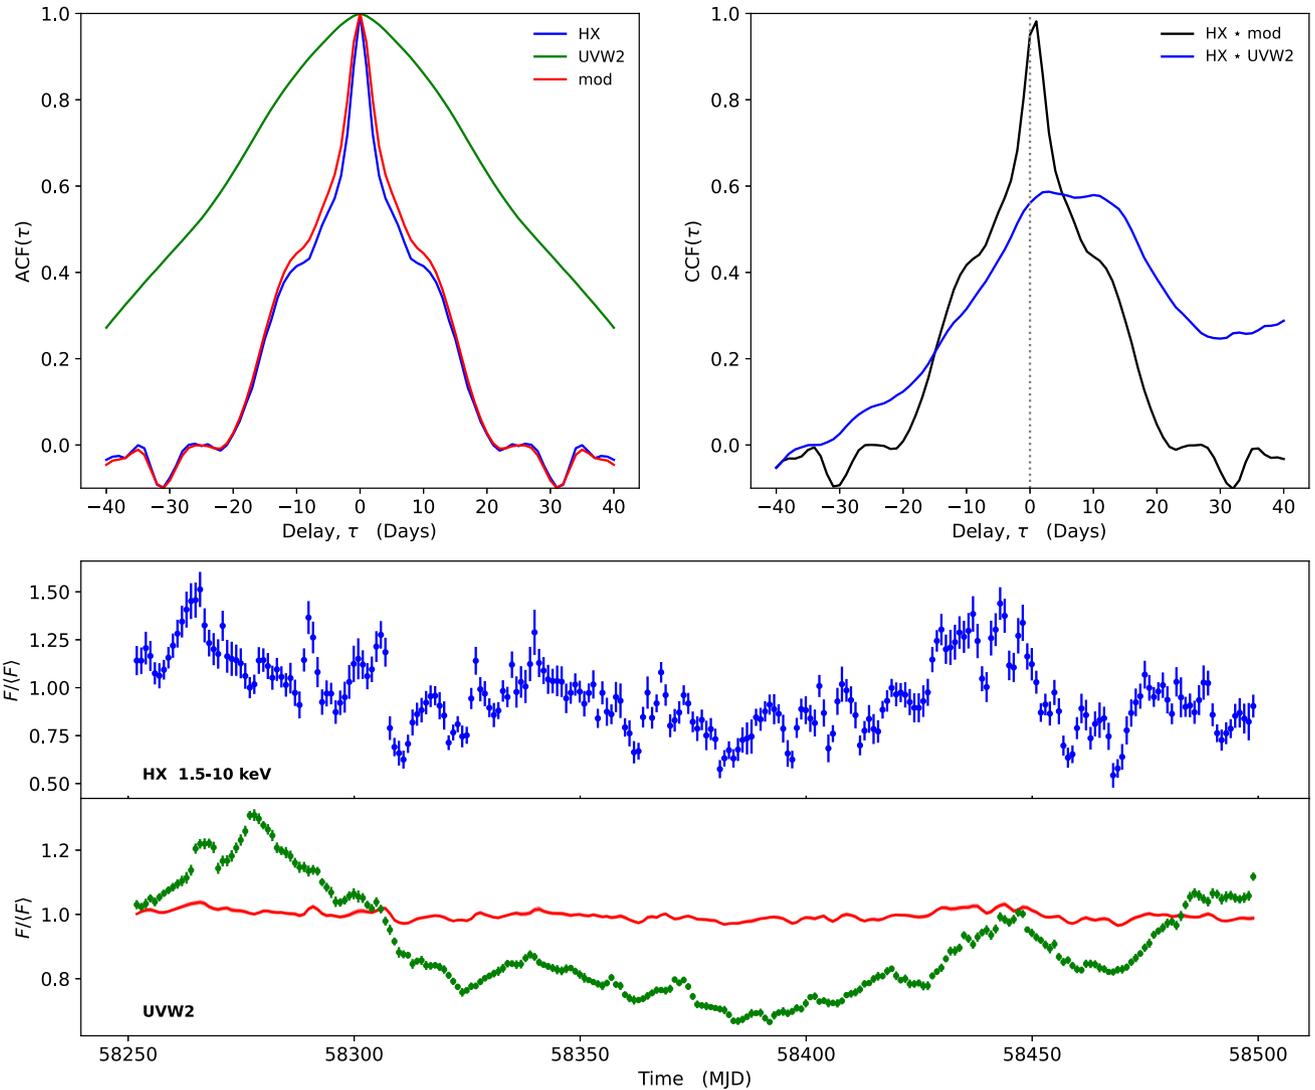


Figure 7. Top-left panel: The auto-correlation functions (ACFs) for each light curve. The colours correspond to the light curves in the bottom two panel. It is clear the *UVW2* is dominated by long-time-scale variations, giving the broad ACF, whereas *HX* appears to have at least two different variability time-scales; one rapid giving the narrow peak, and one slightly slower giving the broad base. These variations are not smoothed out by disc reprocessing, clearly indicated by the similarity between the model ACF and the *HX* ACF. Top-right panel: The CCFs between *HX* and the model (black), and between *HX* and *UVW2* (blue). The vertical grey dotted line shows $\tau = 0$ d. Bottom panels: The *HX* light curve (blue, top panel), *UVW2* light curve (green, bottom panel), and resulting model light curve (red, bottom panel). Evidently, the disc model completely underpredicts the response, and is in no way shape or form able to reproduce the observed *UVW2* light curve. Note that the $\pm 1\sigma$ error on the model light curve is too narrow to be seen clearly in this plot.

McHardy et al. 2005, 2006; Kelly, Sobolewska & Siemiginowska 2011). Previous studies focus on modelling the PSD, and use this to understand what processes are occurring on different time-scales. Instead, we will simply use the PSD to estimate the time-scales we need to filter out of our light curves.

We start by calculating the PSD for both the *UVW2* and *HX* light curves, using the PYTHON spectral timing analysis package STINGRAY² (Huppenkothen et al. 2019). Unlike the black hole binaries, our light curve is not long enough to calculate the averaged PSD from multiple segments, so we increase the samples at each frequency by geometric re-binning of $d\log f = 0.1 \text{ d}^{-1}$. This is shown in Fig. 8. None the less, the PSD at the lowest frequencies/longest time-scales consists of only a few (or even single) point, where the

uncertainty on the power is equal to the power, but the trend is clear. There is more power in the *UVW2* variability at low frequencies, $f \lesssim 0.01 \text{ d}^{-1}$, than in the *HX*. This shows that these long-time-scale variations cannot be driven by reprocessing. There is no clear point at which the *UVW2* power spectra change, so we place a cutoff frequency at $f_{\text{cut}} = 0.01 \text{ d}^{-1}$, below which we consider all UV variability to belong to the long-term trend. This will probably be an underestimate of the intrinsic UV variability.

We filter the *UVW2* light curve on this frequency by taking its discrete Fourier transform using the SCIPY.FFT library (Virtanen et al. 2020) in PYTHON. Discarding everything below f_{cut} , and taking the inverse-Fourier transform back again, gives a light curve only displaying variability on time-scales shorter than 100 d. The bottom panel in Fig. 9 shows the resulting *UVW2* light curve in green. As designed, the long-term trend that was previously present has been completely removed, while preserving the short-term variations.

²DOI: 10.5281/zenodo.6290078.

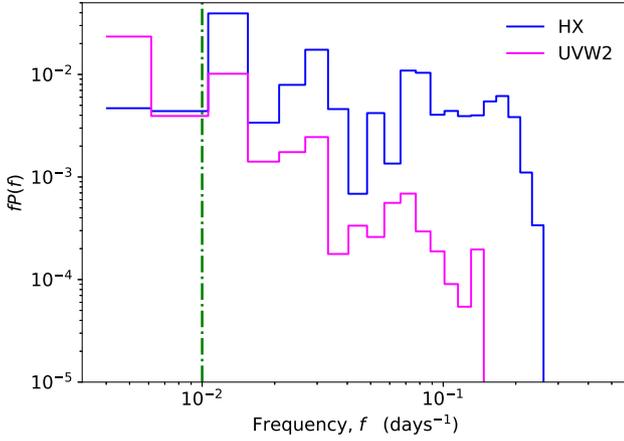


Figure 8. Power spectrum of *HX* (blue) and *UVW2* (magenta) light curves. These have been re-binned, with $d\log f = 0.1 \text{ d}^{-1}$. The green dashed line indicates the frequency cut, below which we set the *UVW2* power to 0 when performing the Fourier filtering. Note that due to the limited number of points in the light curve we have not averaged over multiple segments; as this would remove the long-term trend. Hence the errors on the power spectral density (PSD) will be equal to the power within each frequency bin. We have chosen not to plot the error bars, as these would make the figure incredibly unclear. Also, we highlight that this PSD is only used as a guide for the filtering process.

Additionally, the variability amplitude has been reduced, which is as expected considering the variability power is dominated by the low-frequency variations. We stress that we have only filtered the *UVW2* light curve. Since our model assumes all the X-rays originate from the central region, then if reprocessing occurs it should include all the X-ray variability. Hence, the *HX* light curve should remain unfiltered.

We now perform the same analysis as in Section 4.1, and show the results in Fig. 9. Immediately we see that, although closer, our model light curve still does not match the data. The variability amplitude for the fast fluctuations is still underestimated, albeit by a much smaller factor than previously. In fact, by eye, there are clear similarities, particularly the dip at ~ 58310 MJD, and the rise ~ 58430 MJD. Increasing the response though would not solve some of the underlying issues which is that our model predicts a much faster *UVW2* response to the *HX* fluctuations. We highlight this in the ACF, where it is clear that removing the long-term *UVW2* trend by filtering leads to a significant narrowing of the ACF. However, it is still not as narrow as the narrow core seen in the *HX* ACF which is imprinted on to the model *UVW2* ACFs. Additionally, we can see in the CCFs that filtering has not led to a better correlation between the *HX* and *UVW2* light curves. The CCFs are narrower, a result of removing all long-term trends, but the maximum correlation coefficient is still roughly the same, at just less than ~ 0.6 . This tells us that the main driver in the poor correlation between *HX* and *UVW2* is the presence of the strong rapid variability in the X-ray, which is not at all present in *UVW2*. The model needs to include additional smoothing occurring on time-scales greater than the most rapid *HX* variability.

4.3 Exploring the possibility for an additional reprocessor

Similar issues in the previous multiwavelength campaigns have highlighted an additional contribution from BLR scales. The BLR must have substantial scale height in order to intercept enough UV flux to produce the line luminosities observed. This structure has a larger scale height than the flat disc, so has more solid angle as seen from the X-ray source, so can be an important contributor

to reprocessing. Additionally an increased distance from the X-ray source will lead to an increased smoothing effect, potentially providing the mechanism we need to remove the fast-variability. However, increasing the distance also leads to longer lags, and we can see from the light curves in Fig. 9 that much longer lags may not work with the data either.

The argument for reverberation from a diffuse continuum produced by the BLR or an inner disc wind is that it gives a mechanism for the longer lags seen in the data. This is most clearly established in the *U* band, where the Balmer jump ($\sim 3600 \text{ \AA}$, commonly associated with the BLR) gives sufficiently strong emission such that the observed lag is dominated by the diffuse emission (Korista & Goad 2001, 2019; Cackett et al. 2018; Lawther et al. 2018), but will naturally also extend to other wavebands. More recent analysis have also given results consistent with reverberation off the BLR/inner wind (Dehghanian et al. 2019b; Netzer 2022; Vincentelli et al. 2022). Thus, we expect that the reprocessing structures are more complex. The disc is closest, so responds first, but then there is diffuse continuum response from the wind/BLR on longer time-scales.

To test this we construct a simple model, containing the same disc structure as AGNSED, but with the addition of some outflow launched from radius r_1 at an angle α_1 with respect to the disc. We assume this outflow contributes both to the diffuse and reflected emission, as sketched in Fig. 10. The issue now is that we do not have a clear model of the expected emission from this component. Instead, we approximate this diffuse emission as a blackbody. Our model then, consists of the intrinsic emission dissipated in the accretion flow, calculated as in AGNSED (KD18; see Section 2), followed by a blackbody component located in the UV, used to approximate the diffuse contribution to the SED, and a reflected component, modelled with PEXMON (Nandra et al. 2007). We assume that the luminosity of both the reflected and diffuse emission is set by the fraction of the X-ray luminosity absorbed/reflected by the outflow. This is simply set by the covering fraction, f_{cov} , of the outflow as seen by the X-ray corona, and the wind albedo, A_w , such that

$$L_{\text{diff}} = \frac{1}{2} L_x f_{\text{cov}} (1 - A_w) \quad (5)$$

$$L_{\text{ref}} = \frac{1}{2} L_x f_{\text{cov}} A_w \quad (6)$$

where L_{diff} is the diffuse luminosity, and L_{ref} is the reflected luminosity. The covering fraction is related to the solid angle of the outflow, as seen by the X-ray corona, through $f_{\text{cov}} = \Omega/(4\pi)$. The factor $1/2$ comes from the fact that the outflow should be launched from both sides of the disc, hence the covering fraction is the total fraction of the sky covered by the outflow (this is illustrated in Fig. A1 in Appendix A). However, the disc will block the emission from the outflow on the underside of the disc, resulting in the observer only seeing the top-side emission; i.e. half the emission. The SED model can then be described in XSPEC as AGNSED + $N_{\text{DIFF}} * \text{BBODY} + N_{\text{REF}} * \text{RDBLUR} * \text{PEXMON}$ where N_{diff} and N_{ref} are normalization constants set to satisfy equations (5) and (6). We have also included RDBLUR to account for any smearing within the Fe-K α line originating in the reflected spectrum. For simplicity we will refer to this model as AGNREF, and make it publicly available as an XSPEC model.³

We use the covering fraction of the wind to constrain L_{diff} and L_{ref} in our SED model; however, this does not set the absolute size scale for the wind. Hence, the SED is unable to constrain r_1 and α_1 .

³<https://github.com/scotthgn/AGNREF>

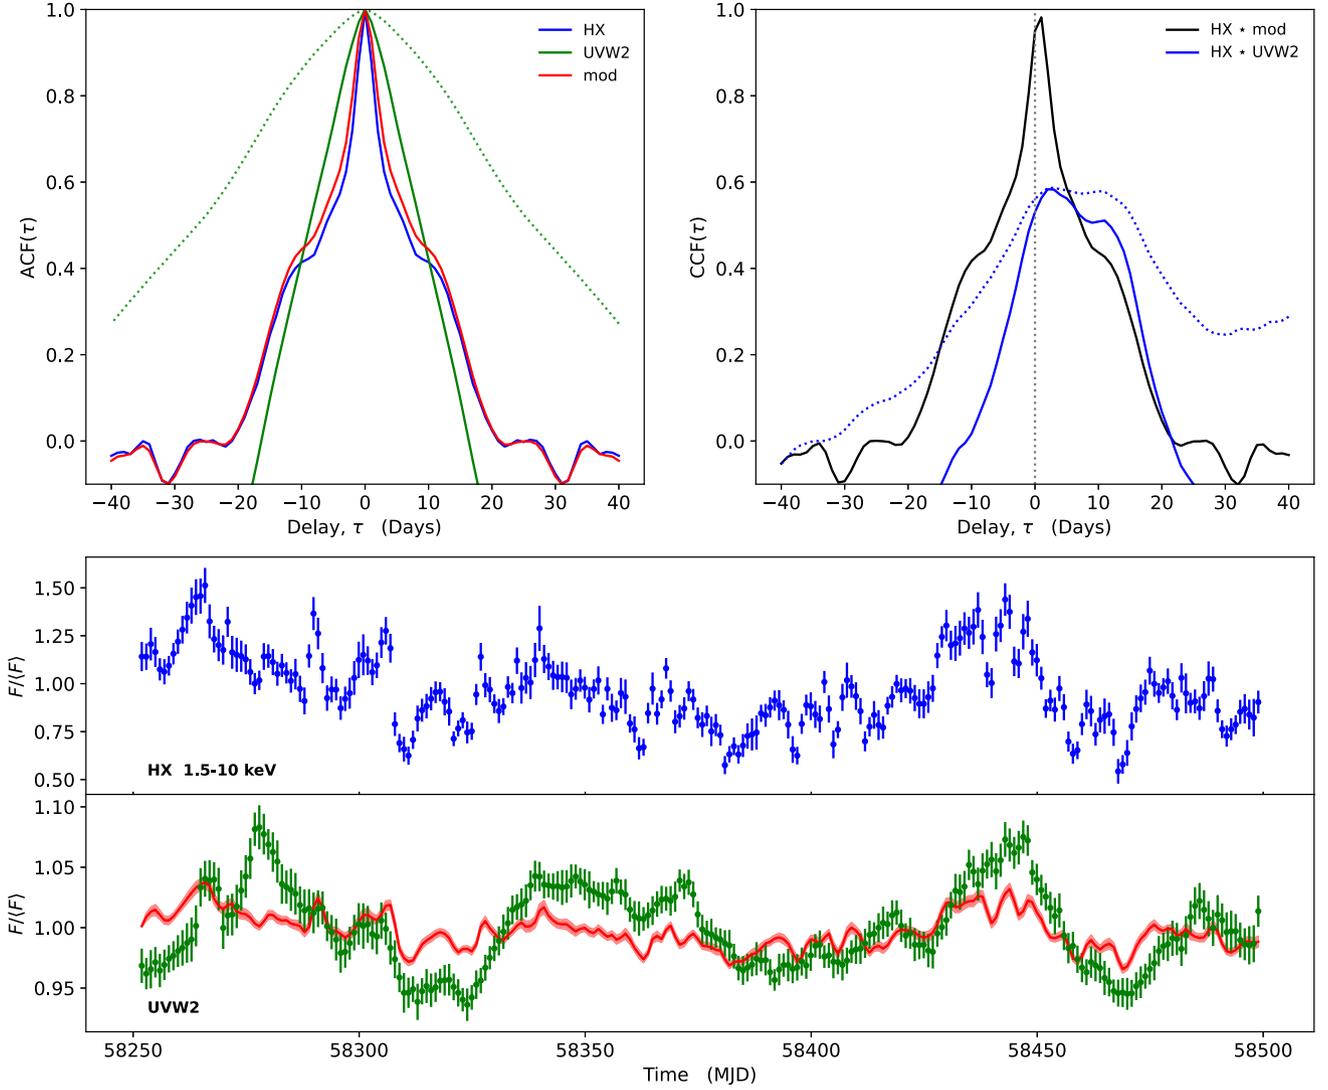


Figure 9. Same as Fig. 7, but now using the filtered *UVW2* light curve. The dotted lines in the top-left and top-right panels show the ACF and CCF, respectively, for the unfiltered *UVW2* light curves in order to highlight the effect the filtering process has on the correlation functions.

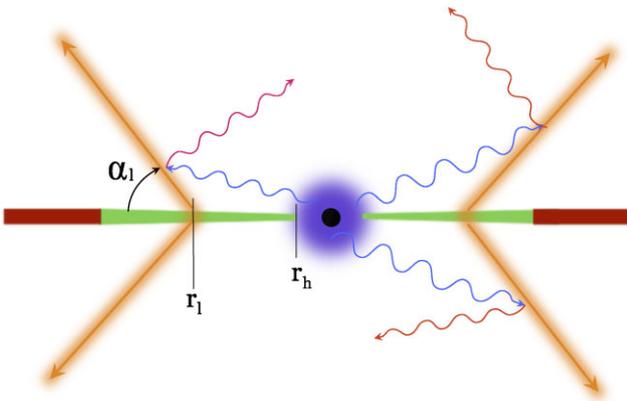


Figure 10. Sketch of the bi-conical geometry used for our additional reprocessor. The disc and coronal regions are identical to AGNSED, while the bi-cone (in orange) illustrates the outflow. We assume the X-rays (blue) induce both reflected (magenta) and reprocessed (orange) emission off the wind.

Instead we treat these as free parameters and marginalize over them in the timing analysis. Additionally, the assumed temperature of the wind will also play a role in the model, determining the position of the blackbody component in the SED and affecting the variability amplitude (see Appendix B).

As in Section 2.3, we calculate our model light curve by varying the input X-ray luminosity according to an observed light curve and creating a series of SEDs based on what X-ray luminosity each grid point in our model geometry sees at any given time. For details on how we do this for our outflow geometry, see Appendix A. However, unlike previous sections, we let the light curves play a role in determining the time-averaged SED to determine the thermal-component temperature, kT_{out} as this cannot be reliably constrained through SED fits as the blackbody shape is only an approximation to the full diffuse emission.

We start by defining an upper and lower limit on kT_{out} . As we have *Swift* light-curve data extending below *UVW2*, down to the *V* band, we do this by performing a grid-based search in kT_{out} and comparing to all available light-curve data. We set $\Delta kT_{\text{out}} = 1 \times 10^{-3}$ keV and search from $kT_{\text{out}} = 1 \times 10^{-3}$ to $kT_{\text{out}} = 1 \times 10^{-2}$ keV.

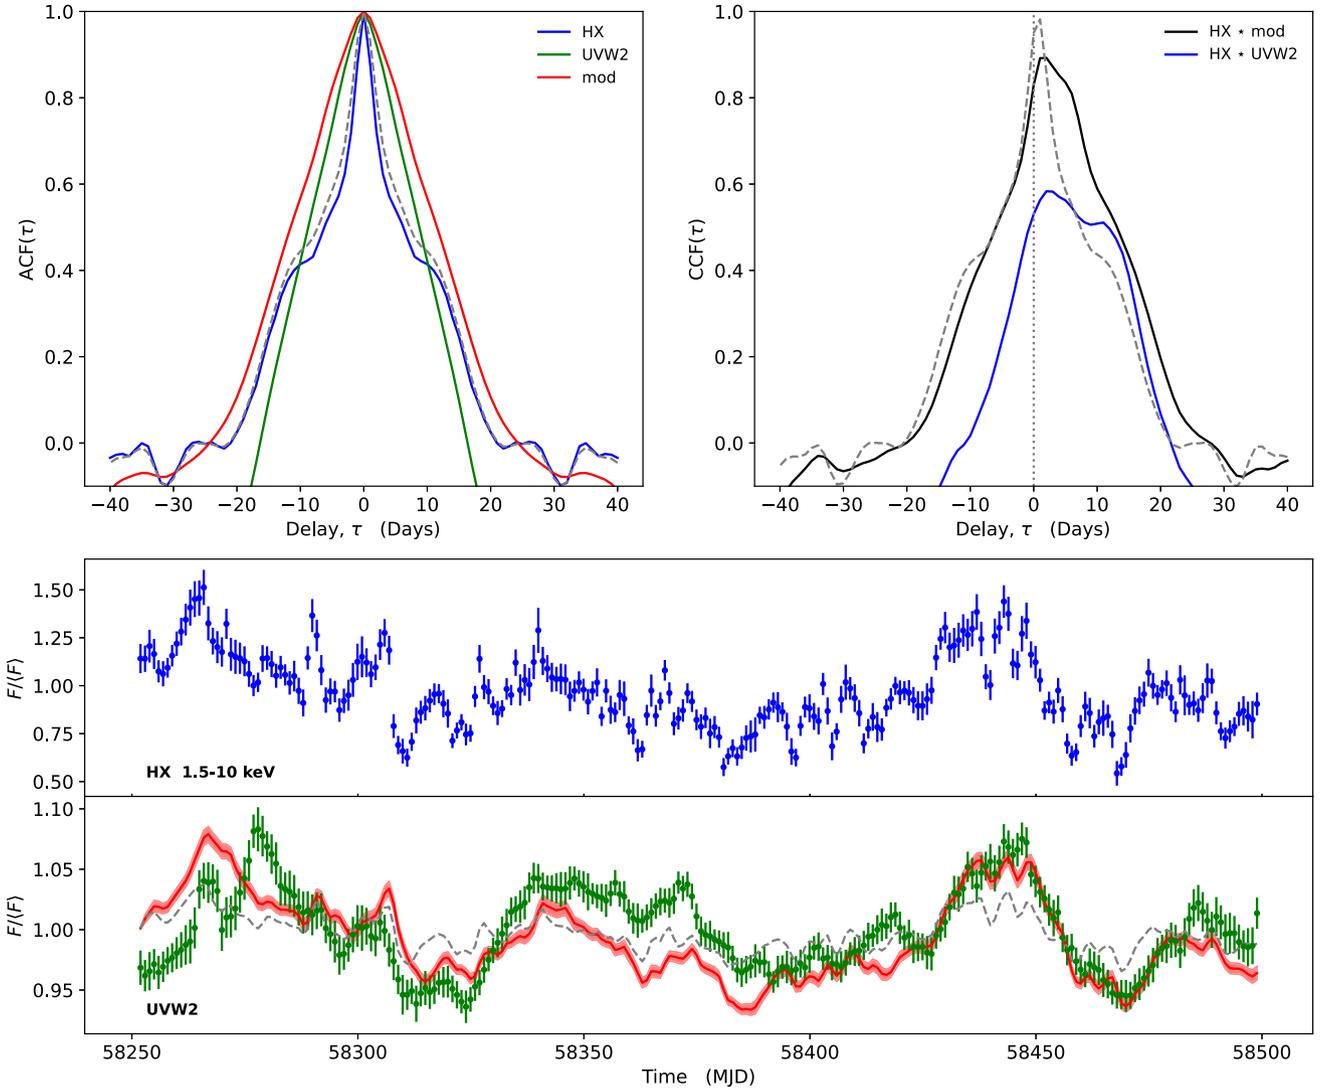


Figure 11. As in Figs 7 and 9, but for the AGNREF model. The dashed grey lines show the results from the AGNSED model, in order to highlight the changes induced by including the additional reprocessor.

From Appendix B, we know that kT_{out} will only affect the variability amplitude of the model light curves, and not the timing properties. Hence, for each point in our kT_{out} grid we fit the model SED, calculate the corresponding model light curves, and compare the variability amplitude to all observed bands, and check if it is overestimated or underestimated. This leads to a lower limit $kT_{\text{out,low}} = 2 \times 10^{-3}$ keV and an upper limit of $kT_{\text{out,up}} = 3 \times 10^{-3}$ keV. We now refine our grid to $\Delta kT_{\text{out}} = 1 \times 10^{-4}$ keV and search between our lower and upper limit. However, we now also construct grids in r_1 and α_1 , in order to attempt to constrain the geometry of the system. To determine the radial limits we take inspiration from the HX to UVW2 CCF in Fig. 9. Although the correlation is poor, the CCF peak tentatively suggests a lag of no more than $\sim 2\text{--}4$ d. Incidentally, this is also suggested by the ACFs, as 2–4 d would provide the smoothing necessary to remove the narrow core in the HX and model ACFs. Hence, we set our launch radial grid to be $200 \leq r_1 \leq 400$ (i.e. ~ 2.2 light-day and ~ 4.5 light-day), with a spacing of $\Delta r_1 = 10$. The lower grid limit in α_1 is set by the minimum angle that still allows the outflow to obtain the required covering fraction derived in the SED fit for each value of kT_{out} , which is $\alpha_{1,\text{low}} \approx 65$ deg. The maximum is simply $\alpha_{1,\text{up}} =$

90 deg, which would give a cylindrical geometry to the outflow. We set the grid spacing in α_1 to $\Delta\alpha_1 = 2.5$ deg. This provides us with a 3D grid in parameter space, over which we perform a parameter scan, providing 2541 potential light curves. For each of these we calculate the CCF between the model and the UVW2 light curve, and naively let our best-fitting parameters be those which give the best correlation with UVW2. This provides $kT_{\text{out}} = 2.5 \times 10^{-3}$ keV, $r_1 = 400$, and $\alpha_1 = 65$ deg. The light curve is shown in Fig. 11.

It is worth pointing out that both r_1 and α_1 are at their respective grid boundaries, indicating that either we have not let our grids extend far enough or the model does not match the light curve and so is unable to constrain the parameters. The latter seems more likely as Fig. 11 shows that there is still a clear mismatch to the observed UVW2 light curve, especially in the first half of the data.

Table 2 give the SED parameters resulting from the above analysis. It is interesting to note how this model deviates from our original AGNSED fit. Firstly, to compensate for the additional thermal component the inner disc radius is moved slightly inwards. The model also derives a large value for the covering fraction, $f_{\text{cov}} \sim 0.87$, so that most of the reprocessing is from the wind/BLR rather

Table 2. Best-fitting parameters for the AGNREF SED model. Parameters with no errors were kept frozen during the fitting process, while parameters highlighted in red are hardwired into the AGNREF model code; so are only included here for completeness. Similar to the fit in Table 1, r_{in} for RDBLUR is only used to set the width of the iron line, and carries no other physical meaning in our model. We have also hardwired all abundances in PEXMON to solar values.

Component	Parameter (Unit)	Value
PHABS	N_{H} (10^{20} cm $^{-2}$)	3.5
AGNREF		
AGNSED	M (M_{\odot})	2×10^8
	Distance (Mpc)	200
	$\log(\dot{m})$ ($\dot{M}/\dot{M}_{\text{edd}}$)	$-1.159^{+0.027}_{-0.031}$
	a_*	0
	$\cos(i)$	0.9
	$kT_{\text{e,h}}$ (keV)	100
	$kT_{\text{e,w}}$ (keV)	$0.347^{+0.045}_{-0.039}$
	Γ_{h}	$1.918^{+0.026}_{-0.027}$
	Γ_{w}	$2.781^{+0.032}_{-0.036}$
	r_{h} (R_{G})	$18.1^{+1.2}_{-0.9}$
	r_{w} (R_{G})	$= r_{\text{out}}$
	$\log(r_{\text{out}})$ (R_{G})	$1.940^{+0.064}_{-0.051}$
	h_{X} (R_{G})	10
	f_{cov} ($\Omega/4\pi$)	$0.852^{+0.081}_{-0.080}$
	A_{w}	0.5
BBODY	kT_{out} (keV)	2.5×10^{-3}
RDBLUR	Index	-3
	r_{in} (R_{G})	393^{+568}_{-167}
	r_{out} (R_{G})	10^6
PEXMON	Γ	$=\Gamma_{\text{h}}$
	E_{c} (keV)	10^5
	Redshift	0.045
$\chi^2/\text{degree of freedom}$	$265.43/166 = 1.60$	

than from the disc. This is most likely driven by the X-ray data, since the magnitude of iron line must be satisfied by the reflected component; while the UV points can easily compensate for changes in the thermal component by adjusting mass accretion rate and outer disc radius. The implications of this is that the diffuse reprocessed component, whose power is most probably set by the fit to the reflected component, becomes a significant factor in the SED. In terms of the light curves this would imply that a significant portion of the variability will originate from the outflow.

The significance of the outflow reprocessed variability in the model light curves becomes exceptionally clear in the ACF, seen in Fig. 11. Here we see that the model ACF no longer resembles that of *HX*, unlike the case where we only considered disc reprocessing. In fact, the majority of the fast variability has been completely removed, an effect induced by the smoothing. The complete removal of the narrow component in the ACF would also indicate that the outflow might even dominate the variability, unsurprisingly since the outflow will see considerably more of the X-ray power than the disc, due to the larger solid angle.

However, we note that there is still some variability which is not captured by the model (especially before 58425 MJD), likely due to both the simplistic nature of the outflow model and some further intrinsic *UVW2* variability which was not filtered out by our simple Fourier filter approach. We also note that this model gives a significantly worse fit to the SED than the original AGNSED fit.

We interpret this as being due to our approximation that the diffuse emission can be modelled as a blackbody. We will use more realistic models for this in future work.

5. CONCLUSIONS

We have developed a full spectral timing code to calculate model light curves in any band given a mean SED and X-ray light curve. Our approach assumes that the X-rays are from an isotropic central source, and that there are no extrinsic sources of variability (e.g. absorption events or changing source/disc geometry). Our approach predicts the amplitude and shape of the light curves, not just a mean lag time, allowing a point by point comparison to the observed light curves.

We apply our model to intensive multiwavelength campaign data set on F9. We fit the SED with a warm Comptonized accretion disc, plus hard X-rays produced by Comptonization from hot plasma heated by the remaining gravitational power from within $r_{\text{h}} \sim 20$. The predicted *UVW2* light curve is entirely inconsistent with the observations, completely underpredicting the observed amplitude of *UVW2* variability (see Fig. 7). Reprocessing of the observed hard X-ray emission cannot be the origin of the *UVW2* variability in this geometry.

None the less, there are features in the UV light curve which look like the X-ray light curve, so we try to isolate the contribution of reprocessing by filtering the *UVW2* light curve to remove the long-time-scale variability. This gives a closer match to our model light curve; however, the amplitude is still under predicted, and our model still varies too fast in comparison to the data.

Recent progress in understanding the observed lags in AGN have focussed on an additional reprocessor from the BLR/inner disc wind. We approximate this emission as a blackbody and fit to the observed light curves to constrain the contribution of the diffuse emission from this reprocessor in the *UVW2* waveband. This gives a better match to the observed variability amplitude, but it is clear that the simple filtering process did not remove all the intrinsic *UVW2* variability, and that a blackbody is not a good description of the diffuse reprocessed emission.

The intensive multiwavelength reverberation campaigns were designed to measure the size scale of the disc in AGN. Instead, the reverberation signal in the UV mostly comes from the BLR/inner wind. These campaigns also assumed that X-ray reprocessing was the main driver of variability in the UV light curves. This is clearly not the case in F9, as can be seen on simple energetic arguments. Unlike the lower Eddington fraction AGN (NGC5548 and NGC4151), the variable hard X-ray power is not sufficient to drive the variable UV luminosity. Instead, most of the variability in the UV comes from some intrinsic process in the disc itself (see also Mahmoud et al. 2022 for a similar conclusion in the similar Eddington fraction source Akn120). This highlights the disconnect between the time-scales predicted by the standard Shakura & Sunyaev (1973) disc models, where intrinsic mass accretion rate fluctuations can only propagate on the viscous time-scale which is extremely long. Instead, this could favour some of the newer radiation magnetohydrodynamic simulations, which indicate a disc structure that is quite unlike these classic models, where there is considerable variability on the sound speed (Jiang & Blaes 2020).

ACKNOWLEDGEMENTS

We would like to thank Juan Hernández Santisteban and Rick Edelson for providing the excellent data on Fairall 9. We would

also like to thank the anonymous referee for helpful comments regarding the manuscript. This work made use of data supplied by the UK Swift Science Data Centre at the University of Leicester. SH acknowledges support from the Science and Technology Facilities Council (STFC) through the studentship grants ST/S505365/1 and ST/V506643/1. CD acknowledges the Science and Technology Facilities Council (STFC) through grant ST/T000244/1 for support. Part of this work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/P002293/1, ST/R002371/1, and ST/S002502/1; Durham University; and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The *Swift* data used in this paper were provided by Jaun V. Hernández Santisteban and Rick Edelson through private communication. These are described in detail in their paper Hernández Santisteban et al. (2020), and are available through the *Swift* archive https://www.swift.ac.uk/swift_live/index.php. The *XMM-Newton* data used in the SED fits are archival, and can be directly accessed through HEASARCH (<https://heasarc.gsfc.nasa.gov/db-perl/W3Browse/w3browse.pl>).

REFERENCES

- Arnaud K. A., 1996, in Jacoby G. H., Barnes J., eds, ASP Conf. Ser. Vol. 101, Astronomical Data Analysis Software and Systems V. Astron. Soc. Pac., San Francisco, p. 17
- Bentz M. C., Katz S., 2015, *PASP*, 127, 67
- Blandford R. D., McKee C. F., 1982, *ApJ*, 255, 419
- Burrows D. N. et al., 2005, *Space Sci. Rev.*, 120, 165
- Cackett E. M., Horne K., Winkler H., 2007, *MNRAS*, 380, 669
- Cackett E. M., Zoghbi A., Reynolds C., Fabian A. C., Kara E., Uttley P., Wilkins D. R., 2014, *MNRAS*, 438, 2980
- Cackett E. M., Chiang C.-Y., McHardy I., Edelson R., Goad M. R., Horne K., Korista K. T., 2018, *ApJ*, 857, 53
- Cackett E. M. et al., 2020, *ApJ*, 896, 1
- Chelouche D., Pozo Nuñez F., Kaspi S., 2019, *Nature Astron.*, 3, 251
- Collier S., Horne K., Wanders I., Peterson B. M., 1999, *MNRAS*, 302, L24
- Dehghanian M. et al., 2019a, *ApJ*, 877, 119
- Dehghanian M. et al., 2019b, *ApJ*, 882, L30
- Done C., Davis S. W., Jin C., Blaes O., Ward M., 2012, *MNRAS*, 420, 1848
- Dovčiak M., Papadakis I. E., Kammoun E. S., Zhang W., 2022, *A&A*, 661, A135
- Edelson R. et al., 2015, *ApJ*, 806, 129
- Edelson R. et al., 2019, *ApJ*, 870, 123
- Elvis M. et al., 1994, *ApJS*, 95, 1
- Fabian A. C., Rees M. J., Stella L., White N. E., 1989, *MNRAS*, 238, 729
- Gardner E., Done C., 2017, *MNRAS*, 470, 3591(GD17)
- Gaskell C. M., Peterson B. M., 1987, *ApJS*, 65, 1
- Gierliński M., Done C., 2004, *MNRAS*, 349, L7
- Hartnoll S. A., Blackman E. G., 2000, *MNRAS*, 317, 880
- Hernández Santisteban J. V. et al., 2020, *MNRAS*, 498, 5399
- Horne K., Peterson B. M., Collier S. J., Netzer H., 2004, *PASP*, 116, 465
- Huppenkothen D. et al., 2019, *ApJ*, 881, 39
- Jiang Y.-F., Blaes O., 2020, *ApJ*, 900, 25
- Kaastra J. S. et al., 2014, *Science*, 345, 64
- Kammoun E. S., Papadakis I. E., Dovčiak M., 2019, *ApJ*, 879, L24
- Kammoun E. S., Dovčiak M., Papadakis I. E., Caballero-García M. D., Karas V., 2021, *ApJ*, 907, 20
- Keck M. L. et al., 2015, *ApJ*, 806, 149
- Kelly B. C., Sobolewska M., Siemiginowska A., 2011, *ApJ*, 730, 52
- Korista K. T., Goad M. R., 2001, *ApJ*, 553, 695
- Korista K. T., Goad M. R., 2019, *MNRAS*, 489, 5284
- Kubota A., Done C., 2018, *MNRAS*, 480, 1247(KD18)
- Laor A., Davis S. W., 2014, *MNRAS*, 438, 3024
- Lawrence A., 2018, *Nature Astron.*, 2, 102
- Lawther D., Goad M. R., Korista K. T., Ulrich O., Vestergaard M., 2018, *MNRAS*, 481, 533
- Liu B. F., Yuan W., Meyer F., Meyer-Hofmeister E., Xie G. Z., 1999, *ApJ*, 527, L17
- Lohfink A. M. et al., 2016, *ApJ*, 821, 11
- McHardy I. M., Gunn K. F., Uttley P., Goad M. R., 2005, *MNRAS*, 359, 1469
- McHardy I. M., Koerding E., Knigge C., Uttley P., Fender R. P., 2006, *Nature*, 444, 730
- McHardy I. M. et al., 2014, *MNRAS*, 444, 1469
- McHardy I. M. et al., 2018, *MNRAS*, 480, 2881
- Magdziarz P., Zdziarski A. A., 1995, *MNRAS*, 273, 837
- Mahmoud R. D., Done C., 2020, *MNRAS*, 491, 5126
- Mahmoud R. D., Done C., Porquet D., Lobban A., 2023, *MNRAS*, available at: <https://doi.org/10.1093/mnras/stac3809>
- Mehdipour M. et al., 2016, *A&A*, 588, A139
- 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
- Narayan R., Yi I., 1995, *ApJ*, 452, 710
- Netzer H., 2022, *MNRAS*, 509, 2637
- 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 and Breach, NY, p. 343
- Page D. N., Thorne K. S., 1974, *ApJ*, 191, 499
- Patrick A. R., Reeves J. N., Porquet D., Markowitz A. G., Lobban A. P., Terashima Y., 2011, *MNRAS*, 411, 2353
- Peterson B. M., 1993, *PASP*, 105, 247
- Peterson B. M., Wanders I., Horne K., Collier S., Alexander T., Kaspi S., Maoz D., 1998, *PASP*, 110, 660
- Peterson B. M. et al., 2004, *ApJ*, 613, 682
- Petrucci P. O., Ursini F., De Rosa A., Bianchi S., Cappi M., Matt G., Dadina M., Malzac J., 2018, *A&A*, 611, A59
- Poole T. S. et al., 2008, *MNRAS*, 383, 627
- Porquet D., Reeves J. N., O’Brien P., Brinkmann W., 2004, *A&A*, 422, 85
- Roming P. W. A. et al., 2005, *Space Sci. Rev.*, 120, 95
- Rózańska A., Czerny B., 2000, *A&A*, 360, 1170
- 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
- Uttley P., Cackett E. M., Fabian A. C., Kara E., Wilkins D. R., 2014, *A&AR*, 22, 72
- Vincentelli F. M. et al., 2021, *MNRAS*, 504, 4337
- Vincentelli F. M., McHardy I., Hernández Santisteban J. V., Cackett E. M., Gelbord J., Horne K., Miller J. A., Lobban A., 2022, *MNRAS*, 512, L33
- Virtanen P. et al., 2020, *Nature Methods*, 17, 261
- Welsh W. F., Horne K., 1991, *ApJ*, 379, 586
- Yaqoob T., Turner T. J., Tatum M. M., Trevor M., Scholtes A., 2016, *MNRAS*, 462, 4038
- Zdziarski A. A., Johnson W. N., Magdziarz P., 1996, *MNRAS*, 283, 193
- Zdziarski A. A., Lubiński P., Smith D. A., 1999, *MNRAS*, 303, L11
- Zoghbi A., Fabian A. C., Reynolds C. S., Cackett E. M., 2012, *MNRAS*, 422, 129
- Zoghbi A., Miller J. M., Cackett E., 2019, *ApJ*, 884, 26
- Zycki P. T., Done C., Smith D. A., 1999, *MNRAS*, 305, 231

APPENDIX A: DERIVING THE BI-CONICAL GEOMETRY AND VARIABILITY

To understand how our model outflow responds to changes in X-ray illumination we need to be able to describe the overall geometry of the wind as seen by the black hole (BH), such that we can calculate the time delay, and how much X-ray flux each grid point sees, such that our normalization is correct.

Our bi-conical model takes three parameter inputs to describe the global geometry: the covering fraction f_{cov} , the launch angle α_l , and the launch radius r_l . To determine the wind variability, we need to first define a grid across the wind surface and secondly determine the time delay at each grid point on this surface. We will start by defining our grid in terms of the polar angles $\cos(\theta)$ and ϕ , as these can be easily related to the solid angle of each grid point, which in turn tells us the X-ray power seen by the grid point. Throughout we place the BH in the centre of our coordinate system.

Firstly, we need to determine the limits of our wind surface. Since we assume it takes the shape of a bi-cone, it will be launched from all azimuths. Hence ϕ ranges from 0 to 2π . As the wind is launched from the disc in the x - y plane $\cos(\theta)$ will range from 0 to $\cos(\theta_m)$, where θ_m is the polar angle for the top of the wind (i.e. the maximal extent of the outflow). The wind will region in the sky (as seen from the BH) illustrated by the red band in Fig. A1. Hence, we can relate θ_m to f_{cov} through the solid angle of this red band, since $f_{\text{cov}} = \Omega_{\text{band}}/4\pi$. From the definition of the solid angle we have $d\Omega = \sin(\theta)d\theta d\phi$. Since θ is measured from the z -axis and down it is simpler to calculate the solid angle of the top conical section, Ω_{con} (non-shaded regions in Fig. A1), and relate it to that of the red band through $\Omega_{\text{band}} = 4\pi - 2\Omega_{\text{con}}$. This gives us the following expression for the band solid angle:

$$\frac{\Omega_{\text{band}}}{4\pi} = \cos(\theta_m) = f_{\text{cov}}. \quad (\text{A1})$$

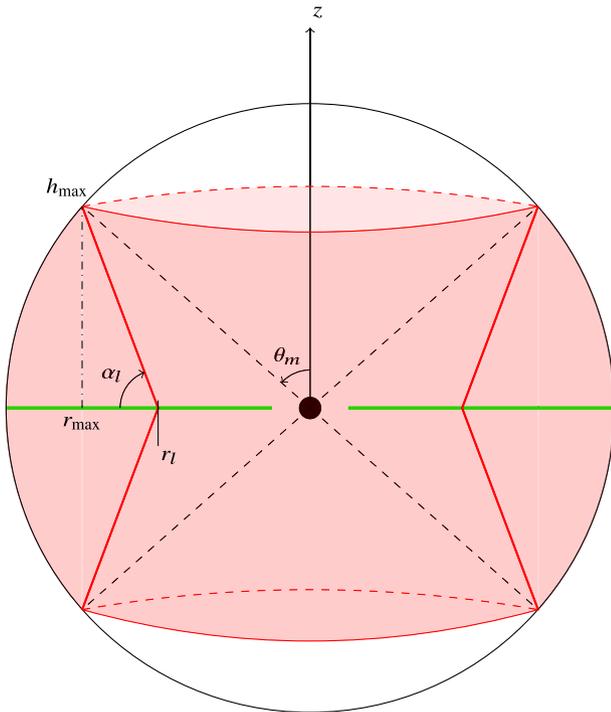


Figure A1. A sketch of the geometry used for the bi-conical outflow model. The outflows are shown as the solid red lines, travelling from the disc (green) to the outer edge of the sphere. The transparent red region shows the area of the sky this outflow subtends, as seen by the BH. α_l is the wind launch angle, r_l is the wind launch radius, while r_{max} and h_{max} are the radius and height of the top of the wind, while θ_m is the corresponding polar angle. The covering fraction given as a model parameter is then simply the solid angle of the red band divided by 4π .

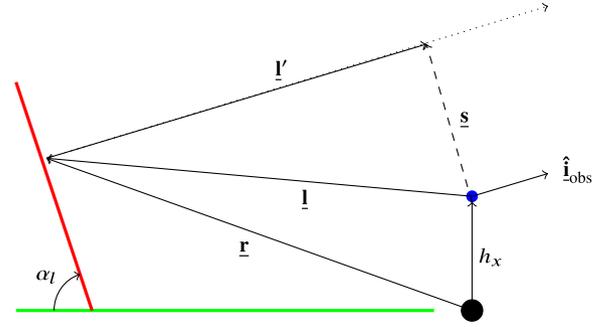


Figure A2. A sketch of the geometry used to determine the time delay for each point on the wind surface. The blue circle indicates the position of the X-ray source in the lamppost approximation, and the vector $\hat{\mathbf{i}}_{\text{obs}}$ is in the direction of the observer.

We now divide our wind area into a polar grid in $\cos(\theta)$ and ϕ , where ϕ is linearly spaced between 0 and 2π with spacing $d\phi$, and $\cos(\theta)$ is linearly spaced between 0 and $\cos(\theta_m)$ with spacing $d\cos(\theta)$.

Since our wind is now fully described by the polar coordinates θ and ϕ , the solid angle subtended by each grid-point is $d\Omega_{\text{grid}} = d\cos(\theta)d\phi$. The X-ray luminosity seen at each grid-point is then $L_x(d\Omega_{\text{grid}}/4\pi)$, giving the reprocessed luminosity at each grid point as

$$dL_{\text{rep}} = (1 - A_w)L_x \frac{d\cos(\theta)d\phi}{4\pi} \quad (\text{A2})$$

where A_w is the wind albedo.

However, to understand the response we still need to know the time delay at each grid-point. Again, we follow the method of Welsh & Horne (1991), but for the geometry sketched in Fig. A2. The time delay is then simply the light traveltime over the path difference between the direct emission, and the emission passing via the wind: $\tau = (R_G/c)(|\hat{\mathbf{l}} + \hat{\mathbf{l}}'|)$. We can relate $\hat{\mathbf{l}}$ to the positional coordinates (r, ϕ, h) through:

$$\hat{\mathbf{l}} = \begin{pmatrix} r \cos(\phi) \\ r \sin(\phi) \\ h - h_x \end{pmatrix} \quad (\text{A3})$$

where h_x is the height of the lamppost corona, and ϕ is measured from the x -axis within the x - y plane. For $\hat{\mathbf{l}}'$ we note that the path difference only extends to the point where $\hat{\mathbf{l}}'$ is tangential to the vector from the X-ray source to its tip, $\hat{\mathbf{s}}$. Since $\hat{\mathbf{l}}'$ must necessarily travel in the direction of the observer, we can also write $\hat{\mathbf{l}}' = d\hat{\mathbf{i}}_{\text{obs}}$, where

$$\hat{\mathbf{i}}_{\text{obs}} = \begin{pmatrix} \sin(i) \\ 0 \\ \cos(i) \end{pmatrix} \quad (\text{A4})$$

is the inclination vector. Writing $\hat{\mathbf{s}} = \hat{\mathbf{l}} + \hat{\mathbf{l}}'$, setting $\hat{\mathbf{l}}' \cdot \hat{\mathbf{s}} = 0$, and solving for d , we have

$$d' = (h_x - h) \cos(i) - r \sin(i) \cos(\phi) \quad (\text{A5})$$

which gives us a time delay

$$\tau(r, \phi) = \frac{R_G}{c} = \left\{ \sqrt{r^2 + (h - h_x)^2} + (h_x - h) \cos(i) - r \sin(i) \cos(\phi) \right\}. \quad (\text{A6})$$

However, as we have defined our grid in terms of $\cos(\theta)$ and ϕ , we need to transform our radial and vertical coordinates such that

$$r(\theta) = \frac{r_1 \tan(\alpha_1)}{\tan(\alpha_1) - \tan\left(\frac{\pi}{2} - \theta\right)} \quad (\text{A7})$$

$$h(\theta) = r(\theta) \tan\left(\frac{\pi}{2} - \theta\right). \quad (\text{A8})$$

Finally, putting this all together, and making the assumption that h_x is sufficiently small that the solid angle of a grid point seen from the BH is the same as that seen from the X-ray source, we have that the luminosity of each point must vary as

$$dL_{\text{rep}}(t, \theta, \phi) = (1 - A_w)L_x(t - \tau(r(\theta), \phi)) \frac{d\cos(\theta)d\phi}{4\pi} \quad (\text{A9})$$

and so the total outflow luminosity varies as

$$L_{\text{rep}}(t) = \sum_{\theta=\frac{\pi}{2}}^{\theta_m} \sum_{\phi=0}^{2\pi} dL_{\text{rep}}(t, \theta, \phi). \quad (\text{A10})$$

APPENDIX B: VARYING THE OUTFLOW PARAMETERS

Here, we assess the impact of the outflow blackbody temperature, kT_{out} , launch radius, r_1 , and launch angle, α_1 (measured from the disc), on the output model light curves. We vary kT_{out} from 2×10^{-3} to 3×10^{-3} keV, r_1 from 200 to 400 R_G , and α_1 from 65 to 90 deg. For each temperature kT_{out} , we construct an SED using our AGNREF model, in order to constrain the energetics and outflow covering fraction, before we calculate the model light curves. Hence, for each

value of kT_{out} , we use a slightly different SED to generate the light curves. These are shown in Fig. B1. Additionally, we only consider launch angles $\alpha_1 \geq 65$ deg, since for certain combinations of r_1 and f_{cov} (predicted from the SED) it is not possible to reach the correct solid angle for $\alpha_1 \lesssim 65$ deg. These results are shown in Fig. B2.

It is clear from Fig. B1 that varying kT_{out} has only a marginal effect on the total SED. It does, however, have a notable effect on the response within the light curve, where we can see that reducing kT_{out} within the limits we have set gives an increase in the variability amplitude. This is simply explained that as we decrease the temperature of the thermal component, it becomes more dominant in the SED at energies associated with the *UVW2* emission. Additionally, changing the outflow temperature makes no difference in how much X-ray flux the outflow sees, and hence no difference in the total variability of the thermal component. Hence, when we shift the thermal emission to a temperature where it makes a greater contribution to the band we are observing in we will see a higher variability amplitude; even though the overall variability of the component itself is not changing. This also explains why the ACF does not change when we vary kT_{out} , since the time-scales the variations occur on are also not changing.

Moving on to the launch radius we see that changing r_1 has an effect on both the light curves and the ACF. This is important, since r_1 cannot be constrained by the SED, hence we can only use temporal information (i.e. the light curves, ACF, or CCF) to estimate this parameter. Firstly, we note that increasing r_1 increases the smoothing effect in the light curve, as highlighted by the widening of the narrow component in the ACF. This is entirely expected, as increasing r_1 will increase the light traveltime to the outflow, which in turn increases the time-scale over which smoothing occurs. We also see slight reductions in amplitude around some of the sharper peaks in the light curve, as r_1 increases. This is an effect of smoothing, not geometry; since the covering fraction is kept constant for each value of r_1 , and so the outflow sees the same fraction of X-ray luminosity no matter the launch radius. Smoothing causes this apparent reduction in amplitude because an increase in smoothing leads to a marginalization over a greater range of X-ray luminosities within the input light curve for each time-step in the model light curve. This explains why we only see this reduction around the sharp peaks, as these will be most strongly affected by the increased range in X-ray luminosities. In other words, the smoothing works exactly as one would expect it to.

Finally, we examine the effect of varying the launch angle α_1 . Firstly, we note that there appears to be a narrower range in output light curves and ACFs when varying α_1 compared to r_1 and kT_{out} . This is most likely due to the limit range in α_1 that we are exploring. We can see, however, that decreasing the launch angle does have an effect on the smoothing of the light curve, as we can see a slight widening of the narrow core in the ACF when α_1 is reduced. This is an effect arising from the way our model hard-wires the outflow solid angle. f_{cov} is set by the SED, and so remains constant under variations in α_1 . This leads to the radial extent of the outflow increasing as α_1 decreases, since the outflow needs to extend to sufficient radii such that it reaches a height large enough to satisfy the solid angle set by f_{cov} . Clearly increasing the radial extent of the outflow will increase the light traveltime, and so increase the range of time delays across the outflow grid; finally leading to an increased smoothing effect.

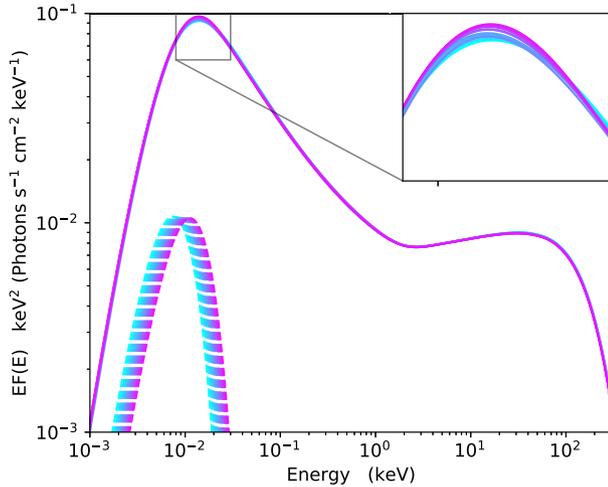


Figure B1. The set of SEDs used for the parameter scan in the outflow model. These have been calculated by fitting AGNREF to the spectral data over the range of kT_{out} . The colour scheme is ascending from blue to purple as we increase kT_{out} . The solid lines show the total SEDs, while the dashed lines show the thermal, diffuse, components. Note that we have removed the reflected component for clarity, as we are mostly interested in the intrinsic SED and thermal component when considering the UV variability. The inset shows a zoomed version of the SED peaks, as this is where we see the greatest effect of changing kT_{out} .

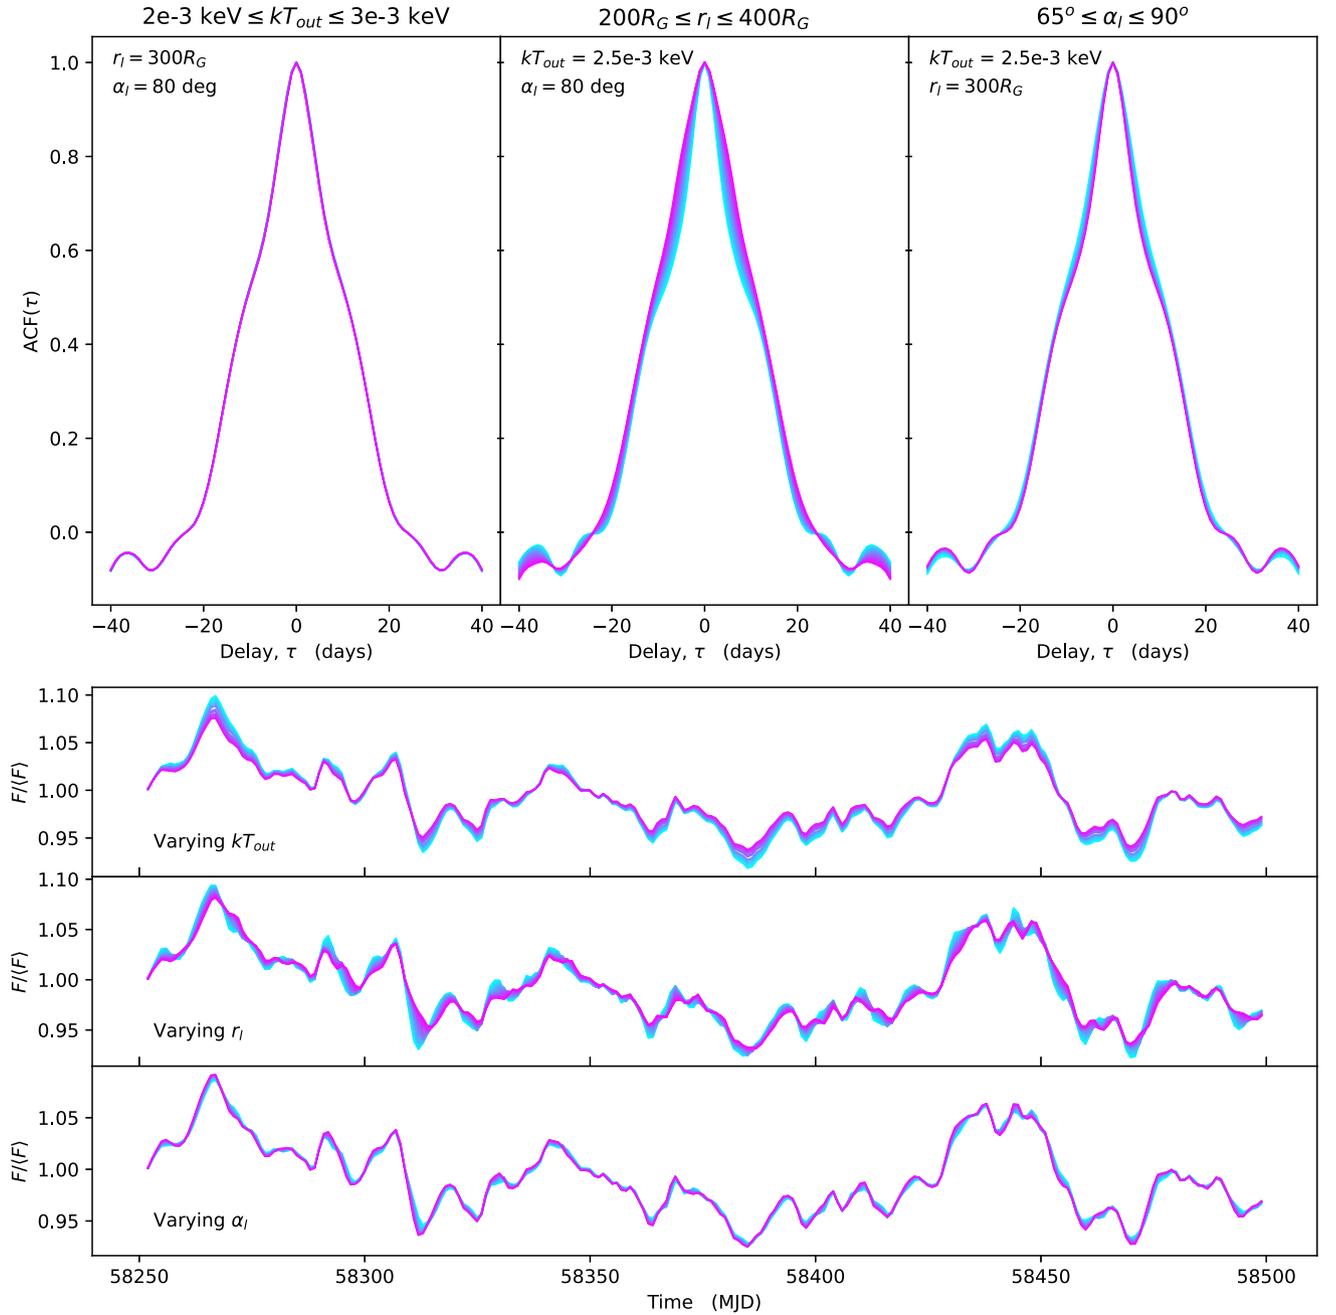


Figure B2. Model light curves from the AGNREF model. The top row shows the ACFs, while varying kT_{out} (left-hand panel), r_l (middle panel), and α_1 (right-hand panel). The bottom three rows show the model light curves, again while varying kT_{out} (top panel), r_l (middle panel), and α_1 (bottom panel). The colour scheme is set such that it ascends from blue to purple as the relevant parameter increases in value (i.e. the lowest parameter value gives a light blue line, while the highest gives a deep purple line). While not being varied, each parameter is kept constant at $kT_{out} = 2.5 \times 10^{-3}$ keV, $r_l = 300 R_G$, and $\alpha_1 = 80$ deg.

APPENDIX C: THE RESPONSE FROM A COMPTONIZED DISC

Repeating the experiment from Section 2.4, but for a Comptonized disc, we find that although there is a change in spectral shape,

the response functions remain almost identical. This is because the spectrum still strongly peaked at an energy linearly related to the seed photon temperature from the disc as the Comptonization is steep. The results are shown in Fig. C1. Clearly, like the case for a standard disc, the response at all energies is dominated by the inner disc response.

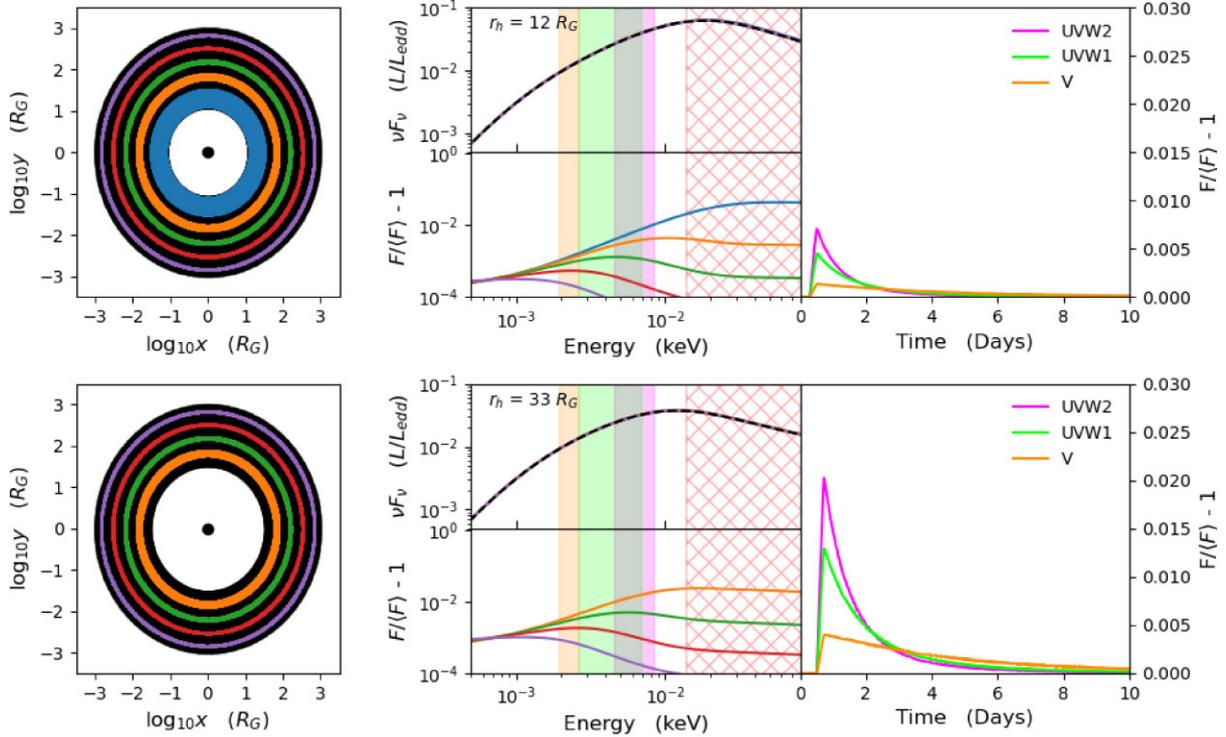


Figure C1. Same as Fig. 3, but for a Comptonized disc.

APPENDIX D: *Swift*, *NICER*, OR *XMM*

In Section 3.2, we made the choice of using the archival *XMM-EPN* data instead of the campaign data from *NICER-XTI* or *Swift-XRT*. Fig. D1 shows a comparison of these spectral data. The *Swift* (black) and *NICER* (red) data are surprisingly quite different below 2 keV, with *Swift* being ~ 30 per cent dimmer at the lowest

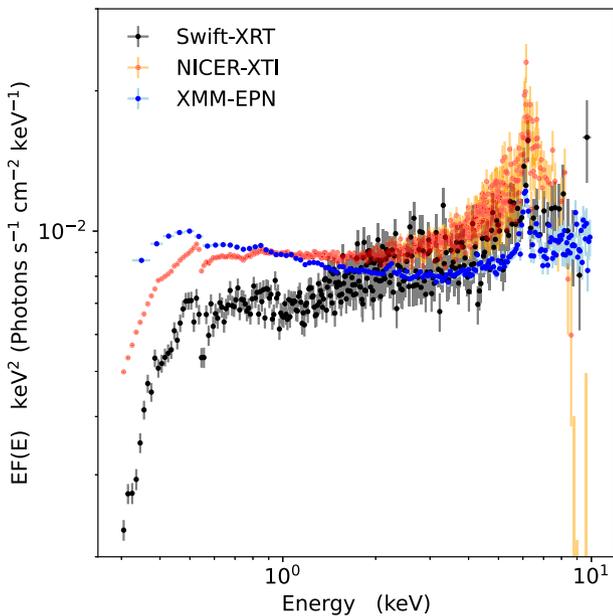


Figure D1. A comparison of *Swift-XRT* data (black), *NICER-XTI* data (red) taken during the observation campaign, and archival *XMM-EPN* data (blue). energies. We checked that this was not due to the different time

sampling. The mean count rate in the *Swift-XRT* light curve across the entire monitoring period is ~ 0.86 counts per second which is actually slightly higher than the count rate in *Swift-XRT* for periods corresponding to just the *NICER* observation periods (~ 0.83 counts per second). Thus, the difference must be due to cross-calibration uncertainties rather than intrinsic to F9. It is not clear which of *NICER* or *Swift* is closer to the ‘real’ spectrum, though with *Swift* being the older (and less sensitive) instrument then it is perhaps more likely that this has been subject to degradation/contamination at low energies. Using *NICER* would then be the obvious choice, except that this instrument has systematic uncertainties at higher energies due to complexity of background subtraction. The derived spectrum is very dependent on the background subtraction for $E \gtrsim 5$ keV, yet even our best attempt, using the newly released SCORPEON background model (released with v.6.31 of HEASOFT), oversubtracts background at the highest energies, as the counts go negative for $E \gtrsim 8$ keV.

In essence, we do not trust *NICER* at high energies and we do not trust *Swift* at low energies. The archival *XMM-Newton* data (blue) give a better match to the *NICER* data than *Swift* at low energies, and a better match to *Swift* than *NICER* at high energies, so we choose this to set the time-averaged spectral shape.

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