





Continuum reverberation in bright quasars using NASA/ATLAS

Zachary Steyn ^{1,★}, Christian Wolf ^{1,2}, Christopher Onken ^{1,2}, Ken Smith ^{3,4}, Ji-Jia Tang ^{1,5},
 Andjelka B. Kovačević ⁶, John Tonry ⁷ and Alejandro Clocchiatti ^{8,9}

¹Research School of Astronomy and Astrophysics, Australian National University, Cotter Road, Weston Creek ACT 2611, Australia

²Centre for Gravitational Astrophysics (CGA), Australian National University, Building 38 Science Road, Acton ACT 2601, Australia

³Department of Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

⁴Astrophysics Research Centre, School of Mathematics and Physics, Queen's University Belfast, Belfast BT7 1NN, UK

⁵Research Center for Space and Cosmic Evolution, Ehime University, Matsuyama, Ehime 790-8577, Japan

⁶Department of Astronomy, University of Belgrade-Faculty of Mathematics, Studentski trg 16, Belgrade, Serbia

⁷Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822-1897, USA

⁸Millennium Institute of Astrophysics (MAS), Nuncio Monseñor Sotero Sanz 100, Of. 104, Providencia, Santiago, Chile

⁹Instituto de Astrofísica, Facultad de Física, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, Santiago, Chile

Accepted 2026 April 10. Received 2026 April 9; in original form 2025 November 23

ABSTRACT

Variable continuum emission from active galactic nuclei (AGNs) can be used to probe the structure of their accretion discs via reverberation mapping. Assuming a variable, hot inner light source irradiates the surrounding accretion disc, time delays between different continuum band light curves reveal light-travel times between their respective emission regions. Inter-band delays measured in several low-luminosity AGNs are ubiquitously ~ 3 times longer than expected from standard disc theory, with evidence that this size discrepancy may decrease in more luminous AGNs. We have analysed high-cadence light curves of 9498 of the brightest quasars between redshift 0.3–2.5 in the largest continuum reverberation study to date. Given the large sample size, we construct bins and fit delays jointly to combine inference across the parameter space and improve lag detections. We find that the size discrepancy persists in our high-luminosity sample, and that the previously seen anticorrelation with luminosity is likely driven by wavelength effects. The complex, non-monotonic wavelength dependence of delay amplitudes strongly suggests that contamination of inter-band delays by variable diffuse emission is widespread in the AGN population. We test delay behaviour against a variety of quasar properties finding longer lags in quasars with: higher Eddington ratios, redder colours, larger optical Fe II EWs, higher iron ratios (both UV Fe II/Mg II and optical Fe II/H β), C IV broad absorption troughs, and lower C IV blueshift.

Key words: accretion, accretion discs – galaxies: active – quasars: general – quasars: supermassive black holes.

1 INTRODUCTION

Investigations of the geometry and kinematics of active galactic nuclei (AGNs) are challenged by our inability to resolve the compact sizes of their central regions (for notable exceptions see: Gravity Collaboration 2018; Event Horizon Telescope Collaboration 2019). In lieu of spatial resolution, the wavelength-dependent aperiodic variability and distinct spectral shape of AGN emission can encode information about the underlying physical geometry and processes in said systems. Reverberation mapping (RM; R. D. Blandford & C. F. McKee 1982; S. J. Collier et al. 1998) aims to decode this hidden spatial structure by cross-correlating the observed variability of different physical components. From the resulting time delays, we can then infer geometric distances provided that the coherent variability between each emission region is linked irradiatively.

Although first implemented within the context of the broad-line region (BLR; J. Clavel et al. 1991; B. M. Peterson 1993; I. Wanders et al. 1997), RM has now been used to measure linear sizes for almost all AGN components from X-ray to infrared (see E. M. Cackett, M. C. Bentz & E. Kara 2021, for a review), achieving equivalent microarcsecond spatial resolution. In this work, we focus on continuum RM in AGN accretion discs, where the irradiating source is frequently ascribed to be the X-ray corona. Under this model of a central X-ray ‘lamp-post’, delay times between continuum bands are a measure of the light travel time between the respective emission regions, as the surrounding disc instantaneously reprocesses the incident X-ray stimulus. We thus expect delays to increase with wavelength, as the variable X-ray emission first reaches the hotter inner disc before the cooler outer regions.

Intensive campaigns of local type 1 AGNs with high cadences and wide wavelength coverage report a lag-wavelength dependence consistent with $\tau \propto \lambda^{4/3}$ (S. J. Collier et al. 1998; M. M. Fausnaugh et al. 2017), as is predicted for the widely adopted

* E-mail: zachary.steyn@anu.edu.au

non-relativistic thin-disc model (geometrically thin and optically thick; N. I. Shakura & R. A. Sunyaev 1973, hereafter **SSD** for Shakura and Sunyaev Disc). The normalizations of lags found in these campaigns are almost universally larger than expected by a factor ~ 3 (M. M. Fausnaugh et al. 2016; E. M. Cackett et al. 2018; I. M. McHardy et al. 2018; R. Edelson et al. 2019; F. M. Vincenzelli et al. 2021). This size discrepancy has been independently reported in microlensing observations of quasar optical emission regions (D. Pooley et al. 2007; C. W. Morgan et al. 2010; A. M. Mosquera & C. S. Kochanek 2011).

The literature has generated many potential explanations for the observed size discrepancy, including complex disc geometries (E. Gardner & C. Done 2017; D. A. Starkey et al. 2023), internal AGN reddening (C. M. Gaskell 2017), disc winds (Y.-P. Li, F. Yuan & X. Dai 2019; M. Sun et al. 2019), and non-blackbody emission (P. B. Hall, G. T. Sarrouh & K. Horne 2018). We further describe here three popular models that we are poised to examine in our data set:

(i) **Diffuse BLR model:** Diffuse continuum emission from the BLR is expected to respond to changes in the ionizing luminosity from the central engine (K. T. Korista & M. R. Goad 2001; D. Lawther et al. 2018; K. T. Korista & M. R. Goad 2019), contributing variable emission across the entire UV–optical spectrum and lengthening delays (as the BLR is more distant than the accretion disc). This introduces a complex wavelength dependence for measured lags (hereafter lag spectra) owing to the wavelength-dependent contribution to continuum bands. Evidence for the BLR model is seen through excess lag signal surrounding the Balmer jump (M. M. Fausnaugh et al. 2016; E. M. Cackett et al. 2018); however, the degree to which this effect is present is not universal across all AGNs (R. Edelson et al. 2019; E. S. Kammoun, I. E. Papadakis & M. Dovčiak 2021a; E. Kara et al. 2023). Frequency-resolved lags provide further evidence for a more distant reprocessor as standard cross-correlation delays are dominated by longer time-scale variability consistent with the BLR (E. M. Cackett, A. Zoghbi & O. Ulrich 2022; C. Lewin et al. 2023). Diffuse BLR continuum emission may even be the dominant signal in inter-band delays, with BLR models matching several AGN lag spectra with only minor contributions from an irradiated disc (H. Netzer 2022; H. Netzer et al. 2024).

(ii) **K21 model:** By considering relativistic coronal irradiation, local disc ionization and substantial corona heights, E. S. Kammoun et al. (2021b) (K21) find the X-ray reprocessing scenario is consistent with observations, predicting wavelength-dependent lags (E. S. Kammoun et al. 2021a; E. S. Kammoun et al. 2023), power spectral densities (PSD; C. Panagiotou et al. 2022a), and spectral energy distributions (SEDs; E. Kammoun et al. 2024) in several well-studied AGNs. Using the best-fitting parameters for NGC 5548 (C. Panagiotou et al. 2025), the K21 model self-consistently predicts the frequency-resolved lags of E. M. Cackett et al. (2022) on all but the longest time-scales where UV–optical variability has been known to exceed the X-ray (P. Uttley et al. 2003).

(iii) **CHAR model:** M. Sun et al. (2020) propose a model where the central X-ray corona and the surrounding accretion disc are magnetically coupled, with coronal variability instead thermalized in the disc through the dissipation of the induced magnetic turbulence. Coherent fluctuations propagate from the corona to the disc at the Alfvén velocity, with the stochastic disc emission response being damped and delayed on the order of the thermal time-scale as the system re-establishes thermal equilibrium. The

model is applied to light curves from NGC 5548 (M. M. Fausnaugh et al. 2016), simultaneously fitting both the measured time delays and the structure functions of all 18 bands. The CHAR model further predicts the frequency-resolved lags of several AGNs (J. Chen, M. Sun & S. Zhou 2024a; J. Chen, M. Sun & Z.-X. Zhang 2024b).

While the UV–optical emission of AGNs shows clear correlated behaviour (J. Clavel et al. 1991; J. H. Krolik et al. 1991) with wavelength-dependent variability (D. E. Vanden Berk et al. 2004; E. Morganson et al. 2014) and delay structure (M. M. Fausnaugh et al. 2016; E. M. Cackett et al. 2018), it remains unclear whether X-ray reprocessing is the dominant mechanism driving these phenomena. Some systems show only low-to-moderate correlation between X-ray and UV–optical light curves (D. J. K. Buisson et al. 2018; R. Edelson et al. 2019; A. M. Morales et al. 2019; E. M. Cackett et al. 2023), particularly on short time-scales (E. R. Partington et al. 2024). C. Panagiotou, E. Kara & M. Dovčiak (2022b) suggest this lack of correlation arises naturally when considering a dynamical X-ray corona, where the physical properties and geometric configuration that determine reprocessing are variable, as may be the case in several AGNs (W. N. Alston et al. 2020; M. D. Caballero-García et al. 2020; A. Zoghbi, J. M. Miller & E. Cackett 2021; C. Panagiotou et al. 2022a). Strong X-ray-UV–optical correlation is only expected if X-ray emission is sufficiently luminous and disc albedo is sufficiently low (A. Secunda, Y.-F. Jiang & J. E. Greene 2025). X-ray luminosities can fall well below this threshold, sometimes requiring X-rays to be a factor $\gtrsim 10$ more luminous to explain changes in UV–optical emission (J. Dexter et al. 2019; M. Marculewicz et al. 2023). Evidence for excess variable emission beyond X-ray reprocessing is predominantly present on longer time-scales (P. Uttley et al. 2003; E. Breedt et al. 2009; M. W. J. Beard et al. 2025), while short-time-scale variations can plausibly be accounted for (C. Panagiotou et al. 2022a; M. Papoutsis et al. 2024; E. R. Partington et al. 2024).

Our picture of coherent AGN variability is further complicated by the detection of inward-travelling intrinsic disc signals. J. V. Hernández Santisteban et al. (2020) first found evidence of this in Fairall 9, where the low-frequency component (determined by a quadratic fit) experienced its minima at earlier times in redder wavelengths. Later, P. Z. Yao et al. (2023) quantitatively measured Fairall 9’s inward travelling lags, positing a steeper radial disc scale height profile as the cause for the expected long viscous time-scales falling within the observed ranges. Inward travelling fluctuations have also been detected in inferred temperature variations (J. M. M. Neustadt & C. S. Kochanek 2022; Z. Stone & Y. Shen 2023; J. M. M. Neustadt et al. 2024), while flux light curves show the commonly seen lamp-post behaviour on short time-scales.

While intensive studies have provided a wealth of high-quality data to test interesting and complex accretion theories on, other works have foregone the intense monitoring of single objects, exchanging higher wavelength and time resolution for more modest specifications and greatly increased sample size to study a broader population of AGNs. These studies typically cover more massive and more luminous AGNs with larger anticipated delays, owing to an insufficient observational cadence to resolve lags in the smaller, individually studied AGNs.

In large time-domain survey studies, the presence of a disc size discrepancy is much less universal. Y.-F. Jiang et al. (2017) present evidence for the ‘disc too big’ problem in their sample of 240 quasars using Pan-STARRS1 survey data (E. F. Schlafly

et al. 2012), although estimates may be biased by making a significance cut which preferentially selects objects with larger lags given the moderate cadence. Using light curves from the Dark Energy Survey (DES; R. Kessler et al. 2015), D. Mudd et al. (2018) derive continuum lag measurements for 15 quasars, finding agreement with both the reported size issue and the standard theory with an elevated (but still sub-Eddington; $L/L_{\text{Edd}} < 1$) accretion rate due to their large uncertainties. Y. Homayouni et al. (2019) take a Bayesian approach to estimating a disc size normalization parameter for their sample of 95 well-defined lags from the first year of the SDSS-RM campaign (Y. Shen et al. 2015), with results consistent with SSD predictions within $\pm 1.5\sigma$. Z. Yu et al. (2020a) present accretion disc sizes for 22 quasars in the DES standard star and supernova C fields, finding agreement with SSD predictions if the radii probed are weighted by their response to X-ray stimulus rather than their steady-state flux (this predicts lags a factor ~ 1.5 larger than the standard luminosity-weighted assumption).

The larger sample sizes of continuum RM survey studies have enabled correlating lags with a range of physical parameters, giving further clues to the origin of inter-band delays. Meta-analysis combining delays from Y. Homayouni et al. (2019) and Z. Yu et al. (2020a) with several intensively studied Seyfert galaxies reports an anticorrelation between the ratio of observed-to-SSD-predicted lags with luminosity (T. Li et al. 2021). The authors find that both BLR contamination (if diffuse BLR continuum equivalent width is anticorrelated with luminosity as is the case for UV broad-lines; J. A. Baldwin 1977) and the CHAR model can explain their results. For a sample of 94 AGNs, H. Guo, A. J. Barth & S. Wang (2022b) produce lags from Zwicky Transient Facility (ZTF; F. J. Masci et al. 2019) light curves, reporting the presence of both elevated disc size measurements and an anticorrelation in the observed-to-SSD-predicted lag ratio with luminosity. The estimated radius at 5100 Å also follows a tight scatter with the luminosity at 5100 Å ($R_{5100} \propto L_{5100}^\alpha$ with $\alpha = 0.48 \pm 0.04$), suggesting an intrinsic relation between the accretion disc and BLR sizes (which scale similarly with continuum luminosity; M. C. Bentz et al. 2013; S. Kaspi et al. 2021; Y. Shen et al. 2024). This conclusion is supported by S. Wang, H. Guo & J.-H. Woo (2023), who collate both continuum disc (R_{5100}) and BLR ($H\beta$) size estimates for several objects in the literature (half of which come from H. Guo et al. 2022b), finding a tight scatter (~ 0.28 dex) between the two. H. W. Sharp et al. (2024) further examine the quasars present in Y. Homayouni et al. (2019), also finding an anticorrelation of observed-to-SSD-predicted lag ratio with luminosity. However, they do not find evidence of diffuse BLR or Fe II emission in the rms spectra and are unable to recreate the tight scatter seen in S. Wang et al. (2023) when using SDSS-RM derived $H\beta$ radii (C. J. Grier et al. 2017), instead citing the CHAR model as a more likely explanation.

It is clear that the current state of continuum reverberation mapping is a complex topic of open discussion. To break degeneracies between literature theories, we need to examine lags for a broad population of AGNs. We tackle this problem by conducting the largest continuum RM study to date with ~ 9500 of the most UV-optical luminous spectroscopically confirmed type 1 quasars, using ~ 3 -d-cadence photometry from the Asteroid Terrestrial-impact Last Alert System (ATLAS; J. L. Tonry et al. 2018). With this large and luminous sample we hope to verify whether the previously seen anticorrelation between the disc size discrepancy and luminosity persists, and if so, try to discern between the BLR, K21, and CHAR models.

In Section 2, we discuss our sample selection process and describe the structure and cleaning of the ATLAS light curves. Section 3 outlines the calculation of theoretical lags and the generation and verification of simulated light curves for independently testing our ability to obtain expected results with the ATLAS data structure. A description of literature-standard RM delay algorithms and our specific implementation of them is provided in Section 4. Results from the ATLAS data are presented in Section 5 with discussion therein. We adopt a flat Λ CDM cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $\Omega_\Lambda = 0.7$. All magnitudes presented are in AB units.

2 SAMPLE AND DATA

Reverberation mapping replaces spatial resolution with temporal resolution, and like any telescope, there is a minimum necessary resolution to probe certain structures. While there is no concrete agreement on how fine temporal sampling must be in irregularly sampled light curves, successful campaigns usually have a finer cadence than the expected delay signal (E. M. Cackett et al. 2021). Continuum RM at its minimum requires light curves in two passbands that probe different rest-frame UV-optical emission regions in the AGN’s accretion disc. While not strictly necessary, an increased wavelength coverage is desirable to compare delays between similar physical radii across redshifts and to investigate the radial profile that sets wavelength dependency.

2.1 Quasar samples

2.1.1 Main sample selection

We construct our parent sample by selecting all spectroscopically confirmed type 1 quasars in the SDSS DR14 quasar catalogue (S. Rakshit, C. S. Stalin & J. Kotilainen 2020, hereafter R20) that lie within the ATLAS field of view (see Section 2.2). From here, we retain quasars with good-quality spectral decomposition estimates of the λL_{3000} continuum available. This ensures that we have the best estimates of accretion disc luminosity for our SSD size predictions, free from iron continuum and line emission contamination. The sample is further restricted to quasars that are sufficiently bright (*Gaia* magnitude $R_p < 18$), unobscured by Milky Way dust ($E(B - V)_{\text{SFD}} < 0.1$), and persistently low air-mass (Dec. $< +70$), retaining quasars with the highest average signal-to-noise in both ATLAS bands. We further require a BpRp Excess Factor < 1.4 (implying a coherent point source) to exclude multiply-imaged quasars. This leaves a sample of 12 055 quasars with bolometric luminosities between 10^{45} and $10^{48} \text{ erg s}^{-1}$ that are strongly correlated with redshift (Fig. 1), as is expected in a flux-limited sample.

We reject quasars with strong radio counterparts to avoid potential contaminating variability from jet processes or non-standard disc structure. Following K. I. Kellermann et al. (1989), objects with a ratio of radio-to-optical flux (R ; 6 cm to B band) greater than 100 are classified as radio loud. The original sample of K. I. Kellermann et al. (1989) sits at a median redshift $z \sim 0.5$, corresponding to approximate rest-frame wavelengths of 4 cm and 3000 Å for radio and optical emission, respectively. The optical luminosity (L_{3000}) is taken from the R20 catalogue. We estimate radio emission by cross-matching with the *NRAO VLA Sky Survey* (NVSS; J. J. Condon et al. 1998) and the *Faint Images of the Radio Sky at Twenty-centimeters* survey (FIRST; R. H. Becker, R. L. White & D. J. Helfand 1995), taking the brighter




Figure 1. Parent sample: luminosities and black hole masses versus redshift compared to the wider SDSS quasar population illustrating our chosen sample is far among the most luminous. Bolometric luminosities are estimated from the R20 λL_{3000} values using a bolometric correction of 5.15 (G. T. Richards et al. 2006). A weaker correlation of black hole mass with both quantities is present.

1.4 GHz flux of the two within 30 arcsec. We then extrapolate a rest-frame 4 cm luminosity assuming a mean radio spectral slope of -0.3 (K. I. Kellermann & G. L. Verschuur 1988). This slope choice is conservative for steep-spectrum quasars at $z \lesssim 4$, cautiously rejecting a greater number of sources. We make no further attempts to distinguish between radio-quiet ($R < 3$) and radio-intermediate ($3 < R < 100$) as NVSS and FIRST are not sensitive enough to make a complete distinction in our sample. This leaves a remaining 10 950 quasars.

2.1.2 Removing objects with contaminating variability

We remove objects with prominent variability from external sources that would dilute the disc RM signal. This includes variable PSF contamination from bright neighbours due to atmospheric seeing and possible microlensing events. As quasars have characteristic variability properties, we can examine objects that greatly depart from the norm to identify candidates for removal. Characterization and comparison of quasar variability is performed on the cleaned ATLAS light curve data (refer to Section 2.2 for details).

To prevent contamination from bright nearby sources, objects are required to be isolated within a 15 arcsec radius in the *Gaia* catalogue. To test whether this criterion is sufficient, we compare the ATLAS light curve rms values of quasars with bright nearby stars (*Gaia* $R_p < 12.5$, < 10.0 , < 7.5 , and < 5.0 mag within 30, 60, 120, and 240 arcsec, respectively) to the distribution of rms values for similar quasars (in λL_{3000} and z) without bright nearby sources. We find no bias in rms values for quasars with bright nearby sources, finding rms values spread uniformly across the comparison rms cumulative distribution function (CDF) independent of angular separation. This suggests that the ATLAS photometry is uncontaminated for our remaining sample.

To search for potential microlensing events, we visually inspect all 10 950 light curve pairs. We find and remove two conspicuous candidate microlensing events based on this visual inspection, leaving the detailed analysis to a future paper. Due to the stochastic nature of AGN emission, weak microlensing contamination




Figure 2. Left: ATLAS colour dependence on redshift. Right: distribution of colour residuals standardized by redshift-dependent median and MAD. The median ($\pm 3\text{MAD}$) is shown as the solid (dashed) black line in both panels. Grey points and grey distribution tails represent removed objects.

can be degenerate with intrinsic fluctuations, even by eye. Thus, we only remove objects where microlensing contributions appear significant. These candidate microlensing events are seen in quasars with the highest fractional changes in flux emission (90th percentile and above) for their respective luminosity. We further remove 11 objects whose light curve pairs show notably strange behaviour or questionable quality.

2.1.3 Removing the dust-reddened tail

Quasars reddened by host galaxy dust will have fainter continuum luminosity estimates, and thus smaller anticipated accretion disc sizes (C. M. Gaskell 2017; C. M. Gaskell et al. 2023). Quasars that are systematically reddened compared to the widespread population are thus removed from our main sample. To determine this, we derived median $c - o$ colour estimates from each object's cleaned NASA/ATLAS light curve pair (see Section 2.2) using a linear interpolation of a 50-d running median to obtain time coincident magnitude estimations for all epochs. To correct for Galactic extinction, we follow R20, choosing the D. J. Schlegel, D. P. Finkbeiner & M. Davis (1998) dust map and the Milky Way extinction law of E. L. Fitzpatrick (1999) with $R_V = 3.1$. This ensures that derived colours are as indicative of uncorrected dust extinction in the R20 λL_{3000} estimates as possible. We further assume that the underlying quasar SEDs are well approximated by the J. Selsing et al. (2016) composite spectral template when correcting the ATLAS magnitudes.

The corrected median $c - o$ colour for each object shows a strong dependence on redshift as prominent line emission migrates through our filters (Fig. 2). Thus, we fit a running median and median absolute deviation (MAD) with width $\Delta z = 0.1$ and remove objects whose median colour lies outside $\pm 3\text{MAD}$ from the running median. This process is repeated a second time as the proportion of dust reddened quasars was large enough to influence the initial running median and MAD. Two iterations were sufficient to produce a convergent, near-symmetric distribution of median colour that contains our main quasar sample of 9498 objects. Of those removed, 1304 objects belonged to dust-reddened tail (a similar fraction to that seen in G. T. Richards et al. 2003) and 135 objects exhibited substantially bluer colour for their corresponding redshift.




Figure 3. Redshift evolution of Fe II EW and its ratio to the relevant low-ionization line in each regime (separated by the black dashed line $z = 0.72$). A running median and MAD trend is displayed in grey. The marginal distributions for each regime are given in the right-hand panels, along with the distribution medians.

2.1.4 Fe II subsamples

We consider the role iron emission may play in influencing continuum RM behaviour. While not typically considered in the BLR contamination framework due to its weaker equivalent width (EW; Y. Homayouni et al. 2022) and difficulty to characterize through modelling (H. Netzer 2022), the presence of Fe II emission is expected to increase measured delays (H. Netzer et al. 2024) and potentially provide a better description of Fairall 9 lags when interpreted with an additional iron BLR reprocessor (R. Edelson et al. 2024). We also wish to explore delay behaviour on $R_{\text{Fe}} = \text{Fe II}/\text{H}\beta$ (and its UV analogue Fe II/Mg II) as one of the principal variables by which quasar spectral diversity can be explained (J. W. Sulentic et al. 2000, 2007). The latter of these two measures may also serve as a metallicity indicator (G. De Rosa et al. 2011). We thus define subsamples partitioned by Fe II EW and the ratio of Fe II EW to a prominent low-ionization line EW (Mg II or H β) to search for potential delay dependences.

As our sample covers a significant redshift range, we are unable to compare a single iron complex across the entire sample. We take Fe II estimates from the UV (optical) complex above (below) a redshift of $z = 0.72$ (panels 7 and above in later binning). Similarly, we choose Mg II (H β) for our low-ionization line above (below) this value. This boundary achieves the greatest number of valid estimates ($\text{EW} > 0$) while aligning with our binning strategy (see Section 4.3). To be confident in the amount of iron emission in our sample, we require a median signal-to-noise per pixel > 10 around the relevant $3000 \text{ \AA}/5100 \text{ \AA}$ continuum. We also require uncertainties in UV Fe II EW, optical Fe II EW, UV Fe II-to-Mg II ratio, and optical Fe II-to-H β ratio to be less than 10 \AA , 6 \AA , 0.4 , and 0.25 , respectively. Fig. 3 shows the redshift evolution and overall distribution of the iron emission derived quantities. We see no evidence of significant redshift evolution with these parameters, motivating a single value cut-off to separate our subsamples. This value is chosen as the median given the approximately symmetric, unimodal distributions. This leaves a remaining 4140 (4138) quasars with low (high) Fe II EW and 4074 (4074) quasars with low (high) iron ratios.

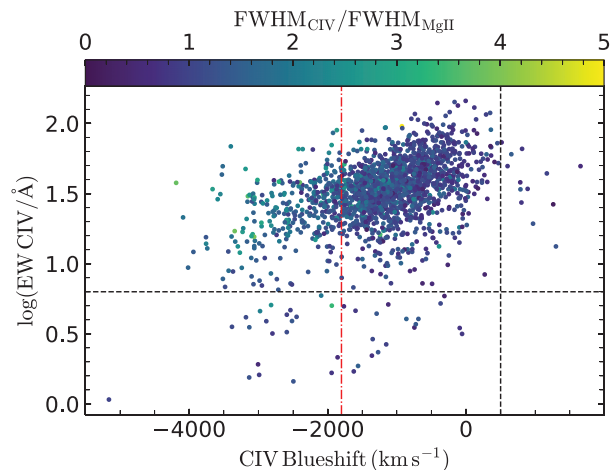


Figure 4. C IV blueshift anticorrelation with equivalent width for the non-BAL quasars with sufficient line SNR, luminosity and FWHM. Black dashed lines represent sample cuts for equivalent width and blueshift. Red dash-dotted line represents the $\leq -1800 \text{ km s}^{-1}$ wind velocity cut-off.

2.1.5 Outflow subsample

As we extend the redshift range probed past previous studies, the high-ionization C IV line becomes visible in the SDSS spectra, with C IV blueshift serving as a diagnostic for strong wind outflows (K. M. Leighly 2004; G. T. Richards et al. 2011). We can thus investigate whether quasars with strong winds exhibit different lag behaviour. To ensure C IV line properties can be accurately determined, we require a median S/N ratio > 15 around the C IV ($1500\text{--}1600 \text{ \AA}$) and Mg II ($2700\text{--}2900 \text{ \AA}$) lines, with both line luminosity and FWHM being greater than zero. Objects that show broad UV absorption line features ($\text{BAL_FLAG} \geq 1$) are also removed. We further require lines to be sufficiently prominent ($\log(\text{EW}_{\text{C IV}}/\text{Å}) \geq 0.8$) and not excessively broad ($\text{FWHM}_{\text{C IV}}/\text{FWHM}_{\text{Mg II}} \leq 3.5$) to be confident fitted line properties represent real line structure (these criteria are visualized in Fig. 4). We calculate C IV blueshift as the blueshift of the peak C IV emission relative to the blueshift of the peak Mg II emission (as Mg II is the best indicator of the systematic redshift available in the SDSS spectral range at our redshifts; Y. Shen et al. 2016).

As we wish to test if the presence of a strong outflow affects lag measurements, we divide our subsample into two bins. Quasars with C IV blueshifts $\leq -1800 \text{ km s}^{-1}$ are labelled as having strong C IV outflows. This threshold is chosen as quasars with these outflow velocities have systematically lower variability amplitudes compared to the general quasar population (J.-J. Tang et al. 2023), although similar behaviour is seen when separating at -550 km s^{-1} (M. Sun et al. 2018b). Objects with peak C IV emission redshifted more than 500 km/s relative to Mg II velocities are removed from the subsample as these may be a sign of an intrinsically different phenomenon rather than having ‘no wind’. This leaves a remaining 1512 (347) quasars with low (high) C IV blueshift.

2.2 ATLAS light curves

ATLAS was developed in partnership between NASA and the University of Hawaii to identify near-Earth and potentially haz-

ardous asteroids (NEAs and PHAs, respectively) through the use of time domain observations. Beginning its program in June 2015, ATLAS started with a single 0.5 m telescope at Haleakala (ATLAS-HKO) in Hawaii. Shortly after, in February 2017, a second identical telescope was added at Mauna Loa (ATLAS-MLO), also in Hawaii. Both the ATLAS-HKO and ATLAS-MLO units can view a declination range of $-45^\circ < \delta < +90^\circ$ which, after accounting for the unobservable sky within 60° of the Sun, leaves an accessible sky of about $24\,500 \text{ deg}^2$ on any given night. With a pixel scale of 1.86 arcsec and a 10560×10560 CCD, each telescope can capture 28.9 deg^2 of sky in a single exposure (excluding the edge mosaic). This means a single telescope can image the entire accessible sky in less than 900 exposures, or at 30 s per exposure (+10 s for readout and re-pointing), within one night. The installation of an additional two units at Sutherland Observatory in South Africa and El Sauce Observatory in Chile in early 2022 has further increased ATLAS coverage in the southern hemisphere and improved the median cadence.

The ATLAS system uses two broadband filters, cyan (*c*: 4157 – 6556 Å) and orange (*o*: 5582 – 8249 Å), with a limiting 5σ magnitude of 19.8 in both bands. Although much broader than the typical filters used in continuum RM, the observed-frame wavelengths probed are similar to the widely used *gri* filters (Y.-F. Jiang et al. 2017; D. Mudd et al. 2018; Y. Homayouni et al. 2019; Z. Yu et al. 2020a; H. Guo et al. 2022b), with approximate conversions given by $c \sim 0.49g + 0.51r$ and $o \sim 0.55r + 0.45i$ (J. L. Tonry et al. 2018). The ATLAS system also utilizes a unique asymmetric cadence structure between the bands with cyan measurements taken only during dark sky periods. Orange band observations are taken irrespective of lunar phase.

2.2.1 Bespoke ATLAS data set

This work uses a bespoke data set created by the ATLAS team from their proprietary data products. Light curves are constructed by first obtaining astrometric and photometric measurements for sources within every ATLAS science exposure using a specially modified version of DOPHOT (P. L. Schechter, M. Mateo & A. Saha 1993; J. Alonso-García et al. 2012). For a detailed description of the custom DOPHOT implementation and its version-specific modifications, see section 2 of A. N. Heinze et al. (2018). Summarized briefly, a smooth, spatially varying PSF model of modified elliptical Gaussians (appropriate for point sources like quasars) is fit first to the brightest sources. Bright sources are progressively subtracted and the PSF model is refit to allow for better determination and measurement of fainter objects. Significant detections in each image, and the associated telescope boresight pointing, are catalogued allowing us to iteratively crossmatch and aggregate any measurements within 3.6 arcsec (~ 2 ATLAS pixels) of our target positions.

The resulting light curves differ from those available on the ATLAS Forced Photometry Server (Shingles et al. 2021)¹, which instead uses the specialized ATLAS photometry software TPHOT to re-measure the photometry at the user-provided coordinates. TPHOT achieves superior performance when measuring photometry from differenced images and is ideal for extracting the light curves of variable sources, such as supernovae, that coincide with extended static sources. However, photometry produced by TPHOT from raw images is susceptible to contamination from

Table 1. Median per-season values with 16th and 84th percentiles. Δt values are provided in observed frame in units of days.

Filter	Δt_{median}	Δt_{min}	N_{epoch}	$(\sigma_f/f)_{\text{median}}$
<i>c</i>	$4.02^{15.96}_{2.92}$	$1.86^{3.97}_{0.94}$	15^{26}_9	$0.03^{0.04}_{0.02}$
<i>o</i>	$2.96^{3.98}_{1.65}$	$0.88^{1.73}_{0.01}$	56^{84}_{43}	$0.04^{0.05}_{0.03}$

neighbouring objects and variations in the background flux. As information about the relative flux variability and total emission differences between passbands can provide important context for continuum RM, we prefer using DOPHOT to obtain photometry from unsubtracted science exposures.

2.2.2 Cleaning individual measurements

To clean the ATLAS light curves, we remove individual photometric observations that are likely not representative of the true flux. Observations taken before June 2017 are excluded due to degraded image quality, with field corrector lenses replaced in May 2017. This reduces our light curve length to a typical 8 seasons per object. We remove observations around the edge mosaic of the detector (within 100 pixels of the *x* or *y* boundary). Measurements that deviate from point source behaviour (not classified as either image type 1: point source or type 7: undetermined) or show large disagreement between aperture photometry and PSF fit ($\text{abs}(\text{AP} - \text{FIT}) < 1.5$) are removed. Images with unusual zeropoints ($m_{\text{zeropoint}} \notin [25, 26]$), poor seeing (PSF major axis > 4 arcsec), and high background brightness (cyan and orange sky counts > 1000 and 3000 , respectively) are also removed. We exclude observations with relatively large errors (more than twice the median error in the light curve of a given object). Any remaining outlying measurements are determined by subtracting a running median with a 50-d width and removing observations outside ± 4 median absolute deviations ($\pm 2.7\sigma$ for a Gaussian distribution).

2.2.3 Aggregating observations within a night

As the core function of the ATLAS project is to track and constrain trajectories for moving objects, ATLAS reduces its nightly coverage to image each field (up to) four times over the course of an hour with an individual telescope. Given the steep drop-off in AGN variability amplitude towards higher frequencies (R. F. Mushotzky et al. 2011; K. L. Smith et al. 2018), we expect to see no meaningful variability over the course of an hour relative to the flux uncertainties, and thus choose to aggregate measurements over these one-hour windows through inverse variance weighting. We do not aggregate measurements between telescopes on the same night as observing conditions on any given night may be dissimilar. Light curve sampling properties after aggregating within a night are shown in Table 1.

As having fewer than four observations per night could be an indicator of degraded viewing conditions (the presence of structured cloud), we wish to determine whether a minimum number of observations for aggregation is required. Such conditions may mean that our uncertainties are underestimated, presenting as non-physical variability that dilutes the desired RM signal. We compare having one, two, and three observations per night to our closest source of ground truth: four observations in an aggregated epoch. The ‘ground truth’ fluxes for comparison are derived from linear interpolation of a running median with a width of 50 d.

¹Available at <https://fallingstar-data.com/forcedphot/>.

The difference in flux between aggregating < 4 observations per night and the estimated ‘ground truth’ is normalized by the associated flux uncertainty to compile a set of standardized residuals across a random 10 per cent subsample of light curves. We do not find strong evidence of underestimated flux uncertainties with reduced χ^2 values falling between 0.9 and 1.6 and 0.8 and 1.1 for the cyan and orange bands, respectively. The slight elevation in reduced χ^2 is more likely due to the imperfect imputation strategy employed that is insufficient to mimic the intrinsic short-time-scale variability present.

Using this same framework, we perform a test to identify the presence of any large discrepancies between data from different telescopes. Unlike comparing the number of observations per night, there is not an obvious choice of which telescope closest resembles the underlying truth. Given our interpolation strategy, we instead opt to use the telescope with the largest number of observations in each filter, with this being ATLAS-MLO for cyan and ATLAS-HKO for orange. We do not find strong evidence for a measurement offset between telescopes with the mean of each residual distribution consistent with zero. We do note larger reduced χ^2 values in this test (between 0.9 and 2.6 for cyan and 1.0 and 1.6 for orange) but again argue this is due to an imperfect imputation method that is sensitive to time-sampling density and separation as, unlike the previous comparison, observations with the lowest flux uncertainties (four-observations-per-night) are no longer necessitated to be time coincident with the input ‘ground truth’ points. This point is further supported by considering that the input points to the ‘ground truth’ approximation achieve the lowest reduced χ^2 in both bands and that the cyan ‘ground truth’ has approximately half the number of contributing epochs as for orange. Even if small biases remain present, we argue this would have a minor effect on measuring lags, given the statistical nature of our analysis where epochs are all considered, rather than isolating a particular feature.

3 SIMULATIONS

Cadence structure and observed photometric noise have a complex relationship with resolving RM lags. To verify that we can recover the expected X-ray reprocessing signal in the ATLAS light curves, we simulate comparable light curves for which we know the underlying truth. As we have observational data to anchor our simulated light curves, we need only predict a lag for each object (Section 3.1), and then generate a light curve that mimics the expected variability structure (Section 3.2). We verify whether our simulated data suitably characterizes the real data behaviour in Section 3.3. We also note here that we use the model defined assuming X-ray irradiation, but it is easily generalized to assume a different, centrally emitted wavelength like the EUV/FUV, which is more strongly correlated with the variable UV–optical emission (E. R. Partington et al. 2024).

3.1 Theoretical lags

With several competing literature models available, we default to the SSD model as our null hypothesis. We refer back to our introduction where we motivated this paper, in part, by aiming to explore whether the anticorrelation of the ratio of observed-to-expected lags with luminosity extends to the most luminous quasars. Under the SSD model, intrinsic disc viscosity facilitates the outward transport of angular momentum, converting gravitational potential energy into heat as the material is accreted in-

wards. Assuming Keplerian dynamics, a massless disc, and local blackbody emission, the effective temperature profile of the disc due to viscous heating is given by (N. I. Shakura & R. A. Sunyaev 1973)

$$T_{\text{visc}}^4(R) = \frac{3GM_{\text{BH}}\dot{M}}{8\pi\sigma R^3} \left[1 - \left(\frac{R_{\text{ISCO}}}{R} \right)^{1/2} \right], \quad (1)$$

where R_{ISCO} is the radius of the innermost stable circular orbit (ISCO; 6 gravitational radii R_g for a non-rotating black hole), σ (G) is the Stefan-Boltzmann (gravitational) constant, and M_{BH} and \dot{M} are the mass and accretion rate of the black hole, respectively.

The lamppost model assumes the surrounding disc is irradiated by an external X-ray source. While the precise geometry of this ‘lamppost’ is unknown, it is commonplace in continuum RM to approximate it as a point source sitting at height H above the central black hole. This simplification is appropriate as E. Gardner & C. Done (2017) show there is little difference between this approximation and considering a radially extended corona. The irradiative local temperature contribution at each annulus for a flat disc is given by (D. A. Starkey, K. Horne & C. Villforth 2016)

$$T_{\text{irr}}^4(R) = \frac{L_X(1-A)H}{4\pi\sigma(H^2 + R^2)^{3/2}}, \quad (2)$$

where L_X is the luminosity of the X-ray lamppost and A is the albedo of the disc (assumed to be radially independent). For the combined local temperature, we include a GR correction factor assuming a face-on disc inclination to account for gravitational redshift and time-dilation effects in the observed temperature profile (T. Hanawa 1989),

$$T_{\text{GR}}(R) = \sqrt{1 - \frac{3GM_{\text{BH}}}{Rc^2}} [T_{\text{visc}}^4(R) + T_{\text{irr}}^4(R)]^{1/4}. \quad (3)$$

As the radii discussed here in the X-ray reprocessing scenario are far larger than the black hole’s ergosphere, we do not correct for Lense-Thirring effects. Relativistic beaming is also neglected as this effect is more dominant in edge-on viewing scenarios not expected with the near face-on assumption of type 1 quasar geometry.

To estimate expected delays across our discs, we translate our observed SDSS parameters of mass and continuum luminosity into a mass-accretion rate. Given a black hole mass estimate, we apply an iterative solver to determine mass-accretion rate such that the observed monochromatic rest-frame luminosity λL_{3000} is recovered using

$$L_\lambda = \int_{R_{\text{ISCO}}}^{R_{\text{out}}} \frac{16\pi^2 h c^2 \cos i}{\lambda^5} \frac{R dR}{e^{hc/\lambda k T(R)} - 1}, \quad (4)$$

where h and k are the Planck and Boltzmann coefficients, respectively, and the radial temperature profile $T(R)$ is given by equation (3). An average viewing inclination of $i \sim 30^\circ$ (P. D. Barthel 1989) is assumed to account for the isotropy assumption used in R20. For objects with poor black hole mass quality flags, we instead use the median black hole mass of quasars within ± 0.1 dex λL_{3000} and ± 0.1 redshift. This approximation is appropriate as black hole mass causes only a second-order effect on expected delays (see Fig. 6 and surrounding discussion). Temperature profiles are evaluated logarithmically in radius, using equation (3) with the interior radii set to $6R_g$. The exterior radii are determined by the $R \gg R_{\text{ISCO}}$ limit ($T \propto R^{-3/4}$) such that $T(R_{\text{out}}) \approx 1500\text{K}$, which is a typical dust sublimation temperature for silicate grains (R. Barvainis 1987). An estimate for the lamppost luminosity is

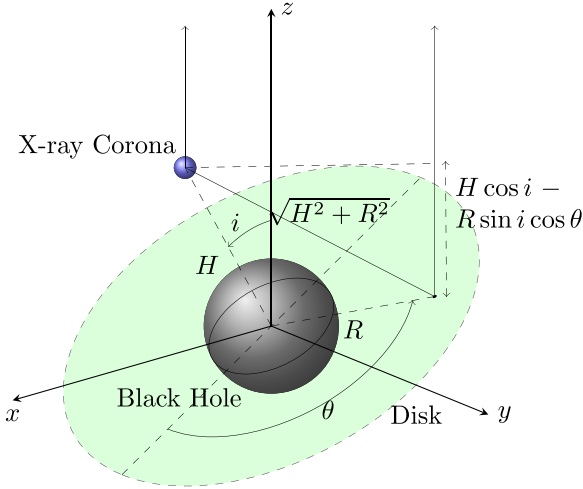


Figure 5. Time delay seen by a distant observer along the z -axis for an X-ray photon reprocessed by the disc where i is the system inclination and $R(\theta)$ is the radial (azimuthal) position.

calculated assuming that the monochromatic X-ray luminosity can be described by a photon index $\Gamma = 2$ over the 2–10 keV range with $\log(L_{2\text{keV}}) = (0.71 \pm 0.03) \log(L_{2500}) + (5.0 \pm 1.0)$ (H. Liu et al. 2021). Here, the L_{2500} is derived from λL_{3000} assuming a near ultraviolet (NUV) slope of $\alpha = -0.3$ ($f_\nu \propto \nu^\alpha$; X. Xie et al. 2016). The lamppost is assumed to be irradiating from a height $H = 10R_g$, as the X-ray emitting region is measured to be in the vicinity of the ISCO (C. W. Morgan et al. 2012; J. A. Blackburne et al. 2014). The albedo is set to zero, creating an upper bound for the lamppost contribution, although estimates are not sensitive to the particular value $A \in [0, 1)$ given the temperature profile in large discs is dominated by internal viscous heating for this model.

Reverberation mapping frequently assumes a linear relationship between driving $D(t)$ and response $X(t)$ lights curves, as given by

$$X_\lambda(t) = \int \psi(\tau|\lambda) D(t - \tau) d\tau, \quad (5)$$

where $\psi(\tau|\lambda)$ is the transfer function and is given by (S. J. Collier 1998)

$$\psi(\tau|\lambda) = \int_0^{2\pi} \int_{R_{\text{isco}}}^{R_{\text{out}}} \frac{\partial B_\nu}{\partial T} \frac{\partial T}{\partial L_X} \frac{R dR d\theta \cos i}{D^2} \times \delta(\tau - \tau(R, i, \theta)), \quad (6)$$

where B_ν is the Planck function, D is the distance to the AGN, and $c\tau(R, i, \theta) = \sqrt{H^2 + R^2} + H \cos i - R \sin i \cos \theta$ (see Fig. 5). The expected lag between bands is given as difference between their respective transfer function centroids. Here, we evaluate the transfer functions for the effective wavelength of each filter, finding little difference when accounting for the finite filter widths. We also note a $\lesssim 0.06$ dex difference comparing with $c\tau(R, i, \theta) = R(1 + \sin i \cos \theta)$, with deviation occurring primarily when $M_{\text{BH}} \gtrsim 10^9 M_\odot$. We do not include the first-order GR correction in the temperature profile when calculating the transfer function, with the responsivity of the disc requiring a more careful GR consideration not given here (K21).

As accretion rate is somewhat degenerate with black hole mass when constrained only by continuum luminosity, we test whether




Figure 6. Ratio of expected delays derived here to those estimated from equation (9) of M. M. Fausnaugh et al. (2016) for our main quasar sample in the λL_{3000} – M_{BH} plane. Black dashed lines represent quasars of fixed Eddington ratio L/L_{edd} .

black hole mass provides strong independent information in the SSD delay paradigm. We compare our expected lags to those derived using equation (9) of M. M. Fausnaugh et al. (2016) using $X = 5.04^{3/4}$ for converting temperature to radius assuming radii are weighted by their response to the incident radiation (S. S. Tie & C. S. Kochanek 2018). This analytical prescription does not include a dependence on the ISCO (and hence black hole mass), instead treating $M_{\text{BH}} \dot{M}$ as a scaling parameter (constrained only by luminosity). Fig. 6 shows the ratio of the two methods has a weak dependence on black hole mass, with delays that account for the ISCO geometry deviating by less than 0.1 dex from pure scaling behaviour (for all but the highest black hole masses). This value decreases to 0.05 dex if the corona instead sits at $20R_g$. This motivates continuum luminosity and redshift as the main driving parameters in the SSD X-ray reprocessing theory, setting the scale of, and the location probed, on the disc, respectively.

3.2 Generating simulated light curves

With the above transfer function machinery, we can generate simulated light curves for each of the ATLAS bands from a given X-ray light curve. We assume the null hypothesis that X-ray reprocessing is the dominant variability mechanism and choose to model X-ray light curves as a random walk ($\text{PSD} \propto f^{-2}$). Under linear reprocessing, this ensures the corresponding UV–optical light curves match the observed random walk variability seen on rest-frame time-scales of 10–250 d in a subset of our current quasar sample using ATLAS data (J.-J. Tang, C. Wolf & J. Tonry 2024). Generating light curves through transfer function convolution suppresses short-time-scale variability below the random walk assumption, especially for the larger discs. Such suppression is not inconsistent with results from J.-J. Tang et al. (2024), although the more simple interpretation of a pure random walk is favoured.

With a specified PSD shape, we use the J. Timmer & M. König (1995) algorithm to generate a random X-ray light curve to convolve with the relevant transfer functions for each passband. Light curves are generated to have an observed-frame sampling of 0.1 d, and have a length at least 10 times the observational baseline to account for red-noise leakage (the tendency for light curves to show power on time-scales longer than the observed baseline; P. Uttley, I. M. McHardy & I. E. Papadakis 2002). Quasars with

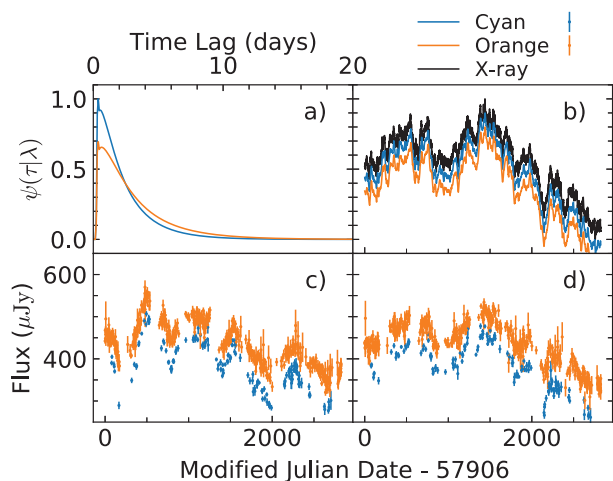


Figure 7. (a) Example transfer function with arbitrary normalization evaluated for an object with $\log(\lambda L_{3000}/\text{erg s}^{-1}) = 45.02$, $\log(M_{\text{BH}}/M_{\odot}) = 8.47$, $z = 0.309$, and $\dot{M} = 0.53M_{\odot} \text{ yr}^{-1}$. (b) Generated X-ray light curve (black) with associated cyan and orange fluxes. Flux values and offsets are arbitrary and chosen for display purposes. (c) Observed ATLAS light curves for the example object. (d) Calibrated simulated light curve based on panels b and c. Both panels c and d apply a $-50 \mu\text{Jy}$ offset to cyan epochs for clarity.

similar luminosity and redshift to our sample have shown random walk behaviour may extend out to at least the long time-scales assumed here (Z. Stone et al. 2022). Power on such long time-scales is inconsistent with X-ray reprocessing in real AGNs as X-rays exhibit flatter PSD slopes on long time-scales (M. Paoillo et al. 2023; A. Georgakakis et al. 2024; S. A. Prokhorenko et al. 2024), but their inclusion is necessary for calibrating against the real ATLAS data. We do later compare light curve behaviour with and without this long-time-scale variability through detrending.

To calibrate the generated UV–optical light curves we sample epochs through linear interpolation to match the observed sampling distribution before normalizing to the mean and variance of the real light curves. Each simulated epoch is given the corresponding observed flux uncertainty and is permuted with Gaussian noise accordingly. The resulting set of light curves have near-identical fractional variability to their observed counterparts with known delay and variability structure as specified by standard X-ray reprocessing theory. This process is illustrated for an example object in Fig. 7.

3.3 Inter-band correlation properties

We further verify whether our simulations mimic the true ATLAS data by comparing the peak inter-band correlation (r_{max} – calculated using ICCF; see Section 4.1) distribution between the two light curve sets. We make these comparisons on two different time-scales: on light curves with and without detrending (using independent quadratic fits in each band). As sample cuts based on observed light curve correlation are frequently made within the literature (Y. Homayouni et al. 2019; H. Guo et al. 2022b), we also examine whether the inter-band correlation of UV–optical emission may be related to the expected lag magnitude.

In the detrended light curves, the pairwise distribution of inter-band correlations is largely symmetric about the one-to-one line with dispersion increasing towards lower r_{max} values (Fig. 8). With a strong anticorrelation between r_{max} and the median

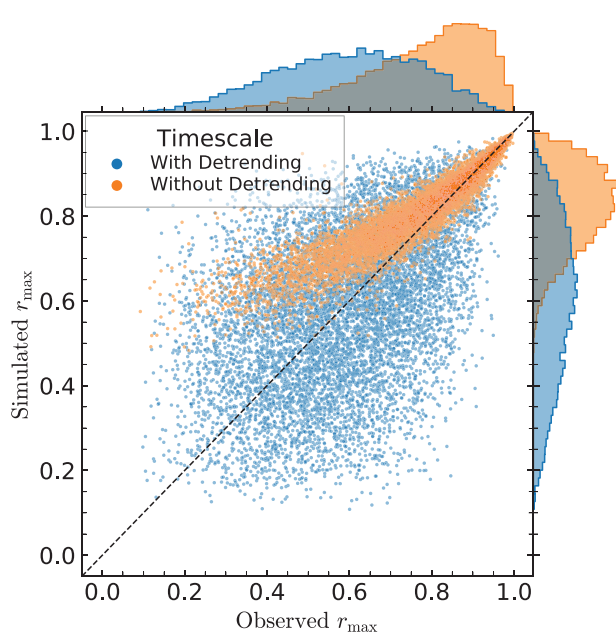


Figure 8. Pairwise distribution of ICCF-derived r_{max} values in the observed and simulated data sets. Values are shown for light curves both with (blue) and without (orange) quadratic detrending. Associated marginal distributions are shown as well.

cyan/orange relative flux uncertainty, the observed diversity and dispersion is plausibly explained by our photometric precision. Light curves with small flux uncertainties will show correlated behaviour irrespective of strong momentary stochastic variability. As quadratic detrending depends on the realized variability and does not filter particular frequencies directly, the inter-band correlation in noisier light curves is strongly subject to how much of the total attributed variance (normalizing on fractional variability) remains after detrending. We also note the real ATLAS data has a higher concentration of objects with $r_{\text{max}} \in [0.4, 0.7]$ and the simulated data having a higher concentration of objects in the $r_{\text{max}} < 0.4$ tail. At face value, this is unexpected as real AGN light curves may contain some level of non-correlated variability not associated with continuum reverberation, whereas our simulated light curves are intrinsically correlated. Normalizing simulated light curves to have near identical fractional variability preserves only the total amount of variability relative to the noise, but not on which time-scales that variability is attributed. Given we are comparing short-time-scale behaviour in the detrended light curves, it is likely the lower average inter-band correlation in our simulated data set comes from overestimating the power of very long time-scale variability (which is then removed). Additionally, the increased inter-band correlation in the real data may be influenced to the partial overlap of the ATLAS filters, whereas the simulated data considers the distinct effective wavelengths only. At this stage we have no reason to suspect either that the intrinsic correlations of UV–optical variability differ significantly from those simulated here on short time-scales or that this variety of correlation cannot be explained by light curve properties, plausibly allowing for sample cuts based on r_{max} (see further discussion in Section 4.3).

When comparing the longer time-scale behaviour of the original cleaned light curves, we see a much smaller level of intrinsic scatter between the simulated and real sets (Fig. 8). This

is expected as without detrending, inter-band correlation is less dependent on the realized variability and is more dependent on the light curve properties (i.e. fractional variability and cadence) which are conserved. We do note however that the inter-band correlation in this simulated data is higher than in the real, with the fraction of light curve pairs with a peak correlation $r_{\max} < 0.6$ being 1.51 per cent and 14.5 per cent, respectively. This bias towards stronger inter-band correlation in the simulated data is seen most strongly in objects with lower luminosities and at higher redshifts. As it remains unclear whether this is an unexplained dependence on magnitude or a genuine departure from correlated behaviour, we analyse light curves both with and without detrending (see Appendix A for comparison).

4 ESTIMATING TIME DELAYS

Most common RM algorithms in the literature aim to measure the centroid of the transfer function, either by calculating the cross-correlation function (CCF) empirically (e.g. ICCF; Section 4.1) or by assuming a functional form for the transfer function (e.g. JAVELIN; Section 4.2). All methods considered here assume a static geometry (lag behaviour is not a function of time), the relationship between light curves is linear (equation 5), and that variability is wide-sense stationary (variability structure is not a function of time). In this section we not only detail our implementation of delay algorithms but also demonstrate improved delay measurements through joint inference (Section 4.3) on our simulated data set (Section 3.2). Further discussion on creating sample cuts based on light curve properties is also provided.

4.1 ICCF

The interpolated cross-correlation function (ICCF; B. M. Peterson et al. 1998) attempts to solve our unevenly sampled light curve problem in a pairwise fashion, projecting sampled times from the reference light curve (with an added delay) onto the response. The response light curve is then linearly interpolated to estimate the flux that corresponds to the projected reference times. The Pearson correlation coefficient can then be applied, iterating over a number of delays to construct the CCF; the centroid of which (τ_{cent}) is ideally equivalent to the centroid of the transfer function (A. P. Koratkar & C. M. Gaskell 1991) and is thus taken as our delay estimate. Uncertainties of τ_{cent} are built up over many Monte Carlo (MC) simulations using random subset selection (RSS; sampling epochs with replacement) and flux randomization (FR; permuting fluxes by their uncertainties) during each MC realization. We use the Cython implementation of ICCF² developed by M. Sun, C. J. Grier & B. M. Peterson (2018a), which reduces flux uncertainties by a factor $n^{1/2}$ for each time the associated epochs are selected during RSS (W. F. Welsh 1999). Our final lag estimate is taken as the median of the resulting cross-correlation centroid distribution (CCCD), with errors given by the 16th and 84th percentiles.

We default to using detrended light curves in our ICCF analysis, finding this preprocessing is a necessary step for recovering delays, reflecting previous recommendations (W. F. Welsh 1999). We choose to globally detrend using an independent quadratic fit for each filter, as is common within the reverberation mapping sphere (I. M. McHardy et al. 2014; M. M. Fausnaugh et al. 2016;

J. V. Hernández Santisteban et al. 2020; R. Edelson et al. 2024). We implement the default two-interpolation strategy which calculates the CCF as the average of the individual CCFs produced when each passband is treated as the interpolated response curve. We do not alter the default thresholds for peak correlation significance ($r_{\max} \geq 0.2$; also the lower bound of our short-time-scale r_{\max} distribution) or the boundaries for calculating the CCF centroid value ($r > 0.8r_{\max}$). CCFs are calculated across ± 100 d in 1-d increments, building up 10 000 MC samples using FR+RSS. We see no difference in results when extending the search range to ± 200 d. We also do not consider any anti-aliasing weighting of the CCF (see C. J. Grier et al. 2017; Q. Yang et al. 2020, for examples) as the range of lags explored is far less than the length of our light curves. Such methods become important when exploring delays with minimal overlapping points between light curves by down-weighting correlations that arise spuriously from the low sample counts.

4.2 JAVELIN

JAVELIN³ (Y. Zu, C. S. Kochanek & B. M. Peterson 2011) is an RM algorithm that models noisy light curves with a Gaussian process (G. B. Rybicki & W. H. Press 1992). Light curve values $\mathbf{y} = \mathbf{s} + \mathbf{n} + L\mathbf{q}$ are represented as a combination of inherent signal \mathbf{s} , observational noise \mathbf{n} , and a general trend $L\mathbf{q}$ (by default the mean flux value in each passband) and are comprised of both the driving ($D(t)$) and response ($X(t)$) curve observations. Entries of the signal covariance matrix $S = \langle \mathbf{ss} \rangle$ are calculated assuming the driving curve follows a damped random walk (DRW),

$$\langle D(t_i), D(t_j) \rangle = \sigma_{\text{DRW}}^2 e^{-|t_i - t_j|/\tau_{\text{DRW}}},$$

where σ_{DRW} is the variability amplitude on short time-scales and τ_{DRW} is the decorrelation time-scale. The covariance between observations in different bands, or within the response band, can then be calculated through appropriate convolution of the DRW kernel and the transfer function as the light curves, and hence the covariances, are assumed to be linearly related (equation 5). JAVELIN assumes the true transfer function can be represented by a tophat function with parameters A , τ , and w representing the scale, temporal shift (lag), and smoothing width, respectively.

By assuming Gaussian likelihoods for the signal, noise, and linear terms, and marginalizing over \mathbf{s} and \mathbf{q} , the likelihood of observing the data given the DRW kernel and tophat function parameters \mathbf{p} is evaluated as

$$P(\mathbf{y}|\mathbf{p}) \propto \mathcal{L} = |S + N|^{-1/2} |L^T C^{-1} L|^{-1/2} \exp\left(-\frac{\mathbf{y}^T C_{\perp}^{-1} \mathbf{y}}{2}\right), \quad (7)$$

where N is the covariance matrix of the noise terms, $C = S + N$, and C_{\perp}^{-1} is the component of C orthogonal to the fitted linear parameters \mathbf{q} . Equation (7) is passed to a Markov Chain Monte Carlo (MCMC) sampler to build up posteriors of model parameters, of which the 16th, 50th, and 84th percentiles are taken as our lower bound, estimate, and upper bound. In practice, DRW priors are first obtained by running JAVELIN on just the reference curve.

Our JAVELIN analysis uses light curves without detrending (see Appendix A discussion), finding delays are better constrained without this preprocessing step. JAVELIN is also prone to aliasing (M. M. Fausnaugh et al. 2016; Z. Yu et al. 2020b) and will return spurious peaks at lags that prevent overlap between light

²Available at https://bitbucket.org/cgrier/python_ccf_code.

³Available at <https://github.com/nye17/javelin>.

curves as there is no data to reject the proposed parameters. For this reason, outlier exclusion is essential, particularly for those with small uncertainties, as JAVELIN will prefer aligning outliers present in both light curves, or placing the outlier in observational gaps if present in only one light curve. Given the seasonal gaps present in our data, we restrict lag searches to be between ± 50 d to prevent aliasing. We find little evidence that lags outside of this window are present in our sample (as shown by ICCF centroids). Unlike ICCF's centroid calculations, JAVELIN does not rely on neighbouring points to estimate lags and thus this smaller search width is chosen.

4.3 Binning and stacking

It is common for large, survey-based continuum RM studies to derive robust lags for only a fraction of their candidate light curves (see Table 2). Given that measuring a delay depends heavily on observed light curve properties such as sampling rate, flux uncertainty, and momentary stochastic variability, it is reasonable to assume that we are still in a regime limited by our data capabilities rather than probing two classes of AGNs, with and without inherent time delays. Previously, light curves that could not constrain a delay were discarded. We describe here a method to extract the information present in these discarded light curves to improve the measurement of the underlying continuum RM signals present in our data set.

In order to accurately constrain lag estimates as a function of our sampled parameter space, we combine inference through stacking. This allows us to marginalize over influential physical properties and states that are not well constrained (orientation, line/diffuse continuum emission, spin, etc). This includes time-dependent phenomena such as disc-wind column density (E. Kara et al. 2021; C. Lewin et al. 2024), coronal properties (W. N. Alston et al. 2020; M. D. Caballero-García et al. 2020; C. Panagiotou et al. 2022a), and momentary stochastic variability. Stacking is already implicit when light curves are analysed as a whole, as delays themselves can be time-varying (C. Lewin et al. 2024). Declination effects such as season length, airmass-dependent average seeing, and photometric quality are also mitigated. The remaining lags are then presented for the ‘average’ AGN as a function of known parameters.

Evidence of complex wavelength dependence in lag spectra (M. M. Fausnaugh et al. 2016; E. M. Cackett et al. 2018) prioritizes binning resolution in wavelength over luminosity, which has a monotonic relation with delay magnitude (H. Guo et al. 2022b; S. Wang et al. 2023). We thus divide our sample space into 10 uniform bins of $1/(1+z)$ before further dividing the space into 5 uniform bins of $\log \lambda L_{3000}$. After conditioning on $1/(1+z)$, each luminosity bin sees scatter of ± 0.2 dex in λL_{3000} . As discussed in Section 3.1, we do not control for black hole mass as its effect is subdominant for the majority of our sample in the SSD prescription. We also test a further 28 $1/(1+z)$ bins shifted one-quarter out of phase to verify trends are independent of Poisson noise from binning choices, although only the independent bins are presented throughout.

Stacking increases the amount of discriminatory information that may be lacking for a single object approach. To stack lag information we take an approach similar to that of previous BLR RM work (S. Fine et al. 2013; J. Li et al. 2017; U. Malik et al. 2024) by combining information at the inference level. Specifically, for ICCF we leverage the Pearson cross-correlation coefficient's

equivalent representation as the mean product of standard scores,

$$r_{xy}(\tau) = \frac{1}{N(\tau) - 1} \sum_{i=1}^{N(\tau)} \left(\frac{x_{i+\tau} - \bar{x}}{s_x} \right) \left(\frac{y_i - \bar{y}}{s_y} \right), \quad (8)$$

where x_{t_i} (y_{t_i}), \bar{x} (\bar{y}), and σ_x (σ_y) are the flux at time t_i , mean flux, and standard deviation of each light curve, respectively. With this representation, we can combine CCF estimates between similar quasars as

$$\begin{aligned} r_{\text{stack}}(\tau) &= \frac{\sum_{j=1}^M (N_j(\tau) - 1) r_j(\tau)}{\sum_{j=1}^M N_j(\tau) - 1} \\ &= \frac{1}{\sum_{j=1}^M N_j(\tau) - 1} \sum_{j=1}^M \sum_{i=1}^{N_j(\tau)} \left(\frac{x_{i+\tau,j} - \bar{x}_j}{s_{x,j}} \right) \left(\frac{y_{i,j} - \bar{y}_j}{s_{y,j}} \right), \end{aligned} \quad (9)$$

where i indexes the observations for each quasar and j indexes each quasar in the bin. This construction is equivalent to having a much longer observational baseline of an individual ‘class’ of quasar discs while enforcing a level of stationarity through detrending (see Section 4.1) and preventing interpolation between the constituent objects. Such an approach has the added benefit of addressing limitations in the ICCF centroid measurement due to the finite sampling of the underlying autocorrelation function (W. F. Welsh 1999). We do not choose to stack by season as well as by object, as the large anticipated lags in our discs require computing cross-correlations at lags that greatly diminish the number of overlapping epochs in a single season.

To verify whether stacking CCFs offers improved lag detection capability, we apply the method to the detrended data set simulated in Section 3. Fig. 9 (top) compares the stacked approach to several other implementations of ICCF, illustrating the following: Panel A shows that with the available quality of our light curves, we are unable to accurately constrain the underlying lags by applying ICCF to light curve pairs individually. For Panel B, if we apply quality cuts – requiring positive lags, $\geq 1\sigma$ significance, and low percentage errors ($|\tau_{84} - \tau_{16}|/\tau_{50} < 1$) – we preferentially select longer lags and bias our recovered lag distribution. Panel C instead restricts lags based on detrended light curve inter-band correlation, showing a distribution of lags that appears approximately unbiased, albeit with lower significance at the short delay end. Panel D makes a similar cut to Panel C, instead based on detrended orange band variability (measured as $\sqrt{S^2 - \langle \sigma^2 \rangle}$ where S^2 is the light curve variance and $\langle \sigma^2 \rangle$ is mean squared flux uncertainty). The orange band is chosen to probe the underlying short-time-scale variability as it is much more densely sampled than cyan. A similar effect is seen in the variability-selected lag distribution, as short-time-scale inter-band correlation and variability probe similar principles: the amount of signal present relative to the noise. In Panel E, combining posterior samples also alleviates the bias, but does little to further constrain the inherent lag distribution except for the longest lags. Panel F shows that to more accurately constrain the true lag distribution, we must stack at the inference level i.e. with cross-correlation functions. Slight improvement is again offered in Panel G as we restrict the constituent light curves to have detrended inter-band correlation greater than 0.5. This threshold is chosen as it retains $\sim 2/3$ of the sample, keeping sample statistics high. The resulting distribution appears to achieve slightly better performance for short delays when compared to Panel F stacking.

Table 2. Number of candidate objects versus the number of successful lags for different cadences (in units of days) and data sources from previous survey continuum RM studies. Bracketed counts are for a more stringent success definition.

Paper	N_{cand}	N_{lags}	Δt_{median}	Data source
This work	9 498		$\sim 3-4$	ATLAS
Y.-F. Jiang et al. (2017)	240	102	~ 3	PanSTARRS1
D. Mudd et al. (2018)	771	15 ^a	~ 7	DES
Y. Homayouni et al. (2019)	222	95(33)	~ 2.5	SDSS RM
Z. Yu et al. (2020a)	802	22	~ 1	DES
H. Guo et al. (2022b)	455	94(38)	$\sim 3-4$	ZTF

^aD. Mudd et al. (2018) select based on variability, not lag measurements.




Figure 9. (a) We apply ICCF (top) and JAVELIN (bottom) to each simulated light curve pair in our sample individually, plotting the recovered versus simulated underlying delay. Error bars are not displayed for clarity. (b) Similar to panel a, but restricting to recovered lags to be > 0 at the 1σ level and have $|(\tau_{84} - \tau_{16})/\tau_{50}| < 1$. (c) Similar to panel a, but requiring detrended light curve inter-band correlation to be in the top 1 per cent. (d) Similar to panel a, but keeping only objects with the 1 per cent most variable objects based on detrended orange-band light curves. (e) Binned posterior samples from the individual applications to each constituent object. (f) Identical binning to panel e, instead of applying the stacked ICCF approach. (g) Identical to panel f, requiring the constituent objects to have a detrended light curve inter-band correlation > 0.5 .

We can apply an analogous method to stacking with JAVELIN as we did with ICCF. Under the reasonable assumption that epochs in each object are independent from every other object, we can combine the likelihood functions of all objects within the bin to infer a joint parameter set

$$\log \mathcal{L}_{\text{stack}} = \sum_{j=1}^M \log \mathcal{L}_j. \quad (10)$$

Quasars with similar luminosity and redshift will have similar variability (σ_{DRW} and τ_{DRW}) and transfer function (A , τ and w) parameters, allowing these quantities to be jointly estimated. Any dispersion in estimating σ_{DRW} and τ_{DRW} due to the range of observed momentary variability is not expected to influence results as the importance of these parameters for recovering lags is minor (Z. Yu et al. 2020b).

We conduct near-identical tests on the simulated light curves without detrending to verify the benefit of stacking JAVELIN (Fig. 9, bottom). The behaviour seen here broadly mirrors that of the

ICCF comparisons. Bar a few outliers, the core distribution of uncertainties produced by JAVELIN is consistently smaller than for ICCF, as is noted in several other studies (e.g. R. Edelson et al. 2019; Z. Yu et al. 2020b, see also Appendix B). Interestingly, JAVELIN individual applications performed better when cut by short-time-scale inter-band correlation and fractional variability even though JAVELIN is applied to light curves without detrending.

Despite the (albeit minor) improvement in recovering simulated lags when cutting by short-time-scale inter-band correlation, we prefer the full stacked approach used in panels f of Fig. 9 for our further analysis. Fig. 10 shows that the short-time-scale inter-band correlation is correlated with luminosity at fixed redshift which is expected as brighter quasars will have smaller photometric uncertainties on average. Requiring that $r_{\max} \geq 0.5$ thus asymmetrically removes a larger number of quasars from the lower luminosity bins. Fig. 10 also displays the median colour residual relative to the redshift trend identified in Fig. 2 across the available parameter space. We see that cutting by inter-band correlation will also preferentially make bins bluer on average by removing red sources (particularly in low-luminosity bins) and thus altering our final results through a colour-dependent effect (see Section 5.5). We note that the correlation-dependent behaviour seen here, and in the following discussion, is also reflected in the long-time-scale inter-band correlation (which is not shown for brevity).

To investigate whether the link between short-time-scale inter-band correlation with quasar residual colour is driven by cyan light curve quality (the limiting band), we match blue quasars ($\Delta_{c-o} \in [-3, -1]$) and red quasars ($\Delta_{c-o} \in [1, 3]$) based on their nearest median cyan relative flux error (σ_c/f_c) counterpart with a maximum difference of 0.005. Quasars can be matched to multiple colour counterparts but duplicate matches are dropped. We divide this matched sample into the previously described $1/(1+z)$ bins and confirm with a one-sided z -test that the median orange percent flux error (σ_o/f_o) is larger among the bluer quasars when controlling for σ_c/f_c as expected. We repeat this test now for short-time-scale inter-band correlation and find that redder quasars still have significantly lower correlations (in all bar panels 3/4) even with comparable cyan, and more precise orange, light curves. This suggests that lower inter-band correlation is an intrinsic trait of redder quasars rather than a data-driven property. Interestingly, we see similar behaviour (albeit with smaller test statistics) in the simulated data set. As the simulated data is fundamentally correlated by design with identical sampling and flux uncertainties as the real light curves, this dependence must be encoded through the variability amplitude.

We test whether the fractional variability of the real light curves (to which our simulations are anchored) differs significantly with residual quasar colour. The fractional variability is significantly lower in redder quasars in both the cyan and orange bands (except panel 4). This is especially surprising for the orange passband where the relative errors are smaller in redder quasars (with comparable cyan errors to bluer counterparts). This behaviour is consistent with dust reddening if the redder quasars are intrinsically brighter and thus less variable (B. C. Kelly, J. Bechtold & A. Siemiginowska 2009; T. Simm et al. 2016; K. Chanchaiworawit & V. Sarajedini 2024). Alternatively, the reduced variability could be explained by a higher Eddington ratio (C. L. MacLeod et al. 2010; B. C. Kelly et al. 2013; S. Kozłowski 2016) as we find a correlation between residual colour and Eddington ratio in our sample. In

spite of controlling for relative flux error (i.e. magnitude for a uniform survey exposure time) and redshift (i.e. distance), we can still see a statistical difference in luminosity for some bins, with redder quasars appearing more (less) luminous in high (low) redshift bins. We retest binning by both $1/(1+z)$ and λL_{3000} as described for the delay analysis above. Of the 50 independent bins we use to compare, we drop 15 for having fewer than 10 unique quasar pairs. Of the remaining 35 bins, 19 bins have statistically less well correlated red quasars and 13 bins have statistically less inter-band correlation and fractional variability while not having significantly higher luminosity. There does not appear to be a pattern across the parameter space with significance (bar panels 3/4 containing no significant bins).

While stacking delay estimates at the inference level offers clear utility in constraining the underlying delay distribution, it suffers from a lack of flexibility. Delay experiments must be chosen prior to analysis. Testing for delay dependence on any new physical parametrization requires re-binning and stacked lags to be re-calculated. This becomes computationally expensive on the large data sets where the stacked methodology is most applicable. Individual lags only need to be derived once and correlations can be drawn easily thereafter and are thus preferred when data quality allows (unlike here; see Fig. 9).

4.4 Multi-peaked behaviour

Delay distributions possessing multiple peaks is a known issue in AGN time delay analysis. These multimodal distributions can be produced in the presence of semirepeating features, particular cadence structures, or poor sample statistics when exploring delays that cause minimal overlap between light curves. The delay distributions in this study are no exception. We wish to identify which measurements may contain artificial signal not associated with real reverberation.

A common approach to remove signal aliases is to smooth delay distributions with a Gaussian kernel before identifying a primary peak and adjacent troughs. A weighting scheme can also be applied to down-weight potentially spurious signals at lags that cause minimal overlap. New lag estimates can then be derived from the truncated distribution subject to possessing a minimum probability mass (C. J. Grier et al. 2017; Y. Homayouni et al. 2019; Q. Yang et al. 2020).

Given our use of a stacked delay analysis, it is entirely plausible to observe multiple peaks in the recovered lag distributions – particularly in bins with low counts where the inherent distribution may not be unimodal. Disentangling which peaks are artificial and which are real becomes difficult in all but the most obvious cases (many peaks separated at regular intervals) without a complete catalogue of predictors that can influence delays on the order of separations seen. As such, we choose not to remove bins that display multiple peaks in their lag distributions, but instead reduce the opacity (i.e. visual weighting) in all further plots. This way, bins that observe potentially real diverse lag behaviour are not discounted, but are not given the same standing as bins that are perhaps better marginalized and representative of the average quasar in that parameter space. The number of peaks here is classified from the number of local maxima in a smoothed posterior lag distribution that exceed 10 per cent of the peak density. Smoothed distributions are obtained through convolution with a Gaussian kernel with a two-day width. We provide discussion on what light curve properties influence the detection of multiple posterior peaks for each delay algorithm in Appendix B.

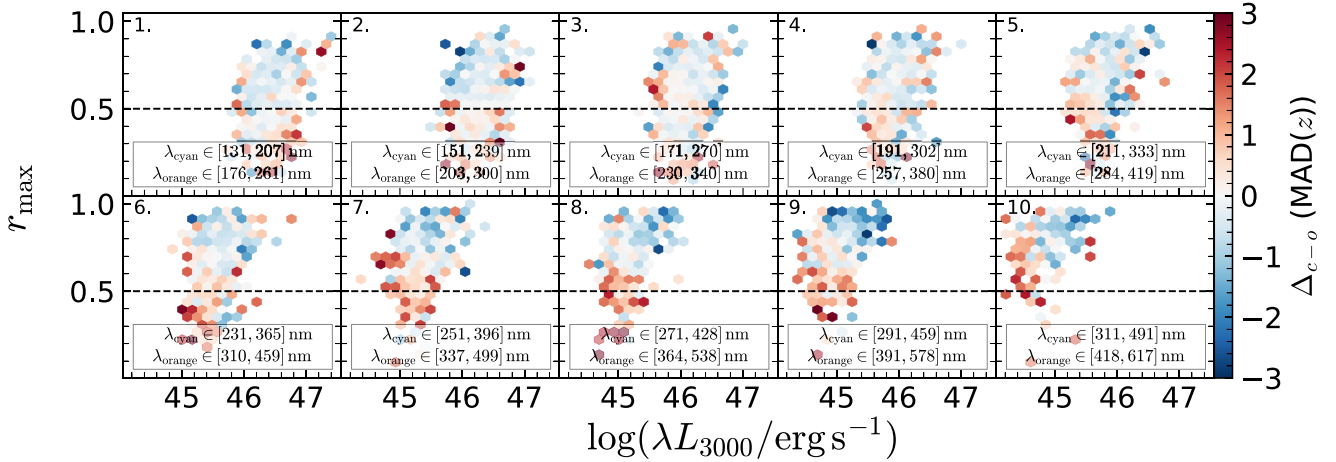


Figure 10. Short-time-scale inter-band correlation versus luminosity for each of the representative set of independent $1/(1+z)$ bins (with displayed wavelengths given by the filter edges at the $1/(1+z)$ mid-point). Each hexagon is coloured by the median colour residual in units of the redshift-dependent spread (MAD) shown in Fig. 2. Bluer quasars on average have stronger short-time-scale inter-band correlation. The black dashed line represents the $r_{\max} \geq 0.5$ cut imposed for the stacking.

5 RESULTS AND DISCUSSION

Armed with the relevant tools to predict (Section 3.1) and measure inter-band continuum delays (Section 4) in our sample, we can now explore what properties influence lags in real AGNs. Enabled by the SDSS spectral measurements (R20) and ATLAS-derived colours (Section 2.1.3) we can also use new and previously described (e.g. Sections 2.1.4 and 2.1.5) subsamples to further investigate time delay behaviour.

5.1 Main sample

We begin our continuum delay analysis with the main sample, where we have the largest number of statistics available and the best chance of constraining robust lag measurements. This sample also has the widest coverage in luminosity and redshift available to us. One of the defining goals of this work was to extend the dynamic range of AGN luminosities for which continuum delay measurements are available, and thus explore whether the anticorrelation of observed-to-expected delays with continuum luminosity persists in our highly luminous sample.

Given the inherent correlation between luminosity and redshift in our flux-limited sample, we choose to explore lag dependence on luminosity while controlling for rest-frame properties. Fig. 11 provides this comparison for the binning strategy described in Section 4.3. Here, luminosity is translated into units of the predicted SSD delay (with increasing luminosity predicting larger delays) for the given redshift range. This allows for easy interpretation of whether lag dependence on luminosity is steeper or flatter than that of the SSD paradigm. The inherent correlation between luminosity and redshift remains encoded as predicted lags migrate to larger values with decreasing rest-frame wavelength. The dynamic range of continuum luminosity within each $1/(1+z)$ bin sits between 1 and 1.5 dex as measured by the difference in median λL_{3000} between the most and least luminous bins.

From Fig. 11 we qualitatively observe that the measured delays scale as expected with luminosity in the SSD paradigm, and that when controlling for redshift, the previously reported anticorrelation is not observed. As most constituent light curves have

approximately the same length, and all objects within a panel have similar redshifts, direct comparisons between lags are valid as we have controlled for the rest-frame wavelengths and time-scales probed. These results disfavour the CHAR model for explaining the larger-than-expected lags as, for a fixed observing duration, it predicts delay times should decrease towards the SSD expectation as luminosity increases (M. Sun et al. 2020; T. Li et al. 2021; J. Chen et al. 2024b). Previous association of the BLR model with the apparent luminosity- $\tau_{\text{obs}}/\tau_{\text{SSD}}$ anticorrelation is predicated on a Baldwin effect for nebular emission (T. Li et al. 2021; H. Guo et al. 2022b). While broad Balmer lines do not show this behaviour across the general AGN population (Rakić et al. 2017; R20), we do detect a weak anticorrelation in our sample for quasars that have valid $H\beta$ measurements; however, the magnitude of this effect is small and on the order of the intrinsic scatter in $H\beta$ EW. It remains unclear how the expected diffuse contamination changes with luminosity, with the contribution to delays potentially growing with increasing (X-ray) luminosity (F. M. Vincentelli et al. 2022).

Our observations comparing the stacked ICCF and JAVELIN methods broadly reflect those made in the literature for single-object applications. Delay estimates agree under the quoted uncertainties in most bins, with JAVELIN often displaying smaller uncertainties (e.g. R. Edelson et al. 2019; Z. Yu et al. 2020b). We present only JAVELIN measurements for the remainder of this work, preferring this algorithm for its smaller scatter around the a priori assumption that delays should follow a power-law scaling. JAVELIN also exhibits multimodal delay posteriors less frequently than ICCF (see Appendix B for more discussion).

We dive further to examine an apparent amplitude trend as we move through our redshift bins. Fig. 12 shows the mean delay amplitude relative to the SSD prediction in each redshift bin. With no strong visual evidence that luminosity-scaling deviates from the SSD prediction, we choose this null hypothesis to normalize the delay amplitudes to compare between redshifts and remove the inherent luminosity correlation. The displayed uncertainties in the mean ratios are given by the scatter around the expected relation as there are clearly systematic uncertainties beyond quoted errors (with most panels possessing significant χ^2 values when compared to the mean ratio under the quoted uncertainties).

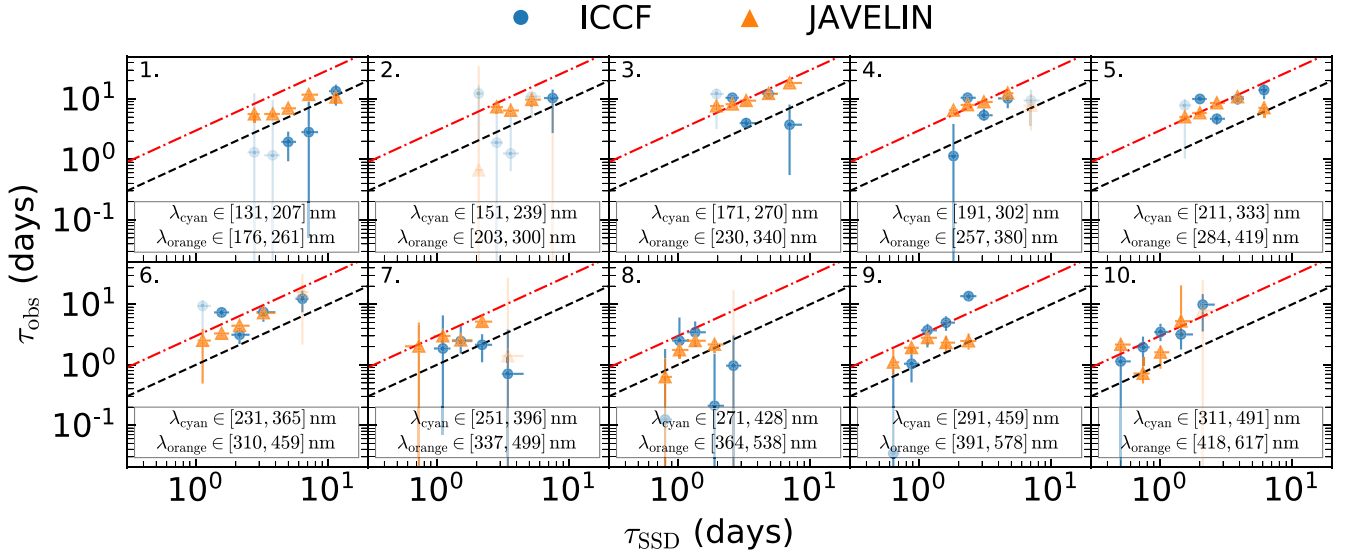


Figure 11. Observed-frame measured delays (symbols) versus SSD expectation (dashed lines are SSD sizes, red dash-dotted lines are $3 \times \tau_{\text{SSD}}$) for each luminosity bin at fixed redshift range (panels are list in order of $1/(1+z)$). The displayed wavelength ranges are the filter edges for the bin's $1/(1+z)$ mid-point. Low opacity indicates delayed posteriors with multiple peaks (Section 4.4).




Figure 12. The mean (and RMS) observed-to-predicted ratio of JAVELIN lags in each panel of Fig. 11 against each filter's effective wavelength at the $1/(1+z)$ mid-points. A non-monotonic amplitude behaviour with wavelength is observed. The global wavelength dependence is re-encoded with a $(1+z)^{-4/3}$ term in the bottom panel. The black dashed and red dash-dotted lines again represent one and three times the SSD prediction.

Such systematic uncertainty is likely caused by not controlling for all confounding parameters in our binning. The resulting mean delay amplitudes show a significant, non-monotonic dependence

on redshift. Such behaviour is likely a wavelength effect rather than a resonant time-scale. This disfavors the K21 model as the sole and primary explanation for unexpectedly large disc sizes as it predicts a monotonic relationship with wavelength. Given that Fig. 12 provides a differential estimate of the lag spectra, it is difficult to rule out whether the underlying disc sizes are inflated and whether the K21 model is still applicable. The BLR model, however, offers a natural way to explain both the overall size discrepancy and the non-monotonic amplitude behaviour with wavelength.

Panel 10 of Fig. 11 shows an elevated lag normalization consistent with those seen in intensively studied Seyfert galaxies (M. M. Fausnaugh et al. 2016; R. Edelson et al. 2019), but it is unclear whether direct comparison is appropriate given our much longer light curve lengths (R. Edelson et al. 2024). Photoionization models of the BLR predict that the Paschen continuum should contribute significantly to the variable optical-NIR emission (K. T. Korista & M. R. Goad 2019; H. Netzer 2022), and may thus explain the extended delay signals of panel 10. While there is tentative evidence for a Paschen jump in the lag spectrum of NGC 4593 (E. M. Cackett et al. 2018), there is little evidence for a Paschen jump in AGN SEDs. This can be reconciled if the Paschen jump is counterbalanced by a blending of high-order Paschen lines (H. Guo et al. 2022a), with diffuse emission from this series predicted to make a significant contribution to the total flux (F. M. Vincetelli et al. 2021; H. Guo et al. 2022a).

Photometric contamination from diffuse emission can explain the convex lag amplitude function in λ (Fig. 12), with a local minimum where the ATLAS filters straddle the Balmer jump. The observed delay is minimized when the Balmer continuum contributes significantly to the leading cyan filter and not the orange. Such excess contributions around the Balmer jump have been seen in local type 1 Seyferts through an excess lag signal in the U -band (M. M. Fausnaugh et al. 2016; E. M. Cackett et al. 2018; R. Edelson et al. 2019), although the degree of this effect is not universal (E. S. Kammoun et al. 2021a; E. Kara et al. 2023). Given its presence in our large population study, this suggests contamination from the diffuse BLR emission is widespread. We




Figure 13. Log ratio of measured-to-expected lags in each bin against continuum luminosity. Points are colour and opacity coded by their $1/(1+z)$ bin and multi-peaked status, respectively, with one and three times the SSD prediction shown as before. An anticorrelation between the two is seen either side of the Balmer jump.

do not test for delay behaviour as a function of Balmer continuum strength directly as the R20 spectral fitting does not take into account Balmer edge shifts from higher electron number densities and Doppler broadening which can be required in order to accurately measure Balmer continuum contributions (C. Jin et al. 2012).

The decrease in delay amplitude shortwards of 300 nm can be explained by the decreasing contribution to total flux from the diffuse Balmer continuum relative to the disc continuum emission. Ly α emission is not expected to contribute substantially to panel 1 (Fig. 11) delays as only the higher redshift objects in this bin will probe Ly α , and the effect will be washed out due to the broad nature of our photometric filters. Detailed photoionization calculations are beyond the scope of this paper but are necessary for verifying the interpretation given here. These findings are consistent with F. Pozo Nuñez et al. (2025) where similar rest-frame wavelength delays (to panel 1) are consistent with the SSD model in a similarly bright and massive quasar ($L_{\text{bol}} = 8.27 \times 10^{47} \text{ erg s}^{-1}$ and $M_{\text{BH}} = 8.9 \times 10^8 M_{\odot}$) using carefully chosen medium band filters that avoid BLR emission.

The complex delay amplitude dependence on wavelength also offers a way to reconcile the lack of luminosity- $\tau_{\text{obs}}/\tau_{\text{SSD}}$ anticorrelation seen here with previous studies. Fig. 13 displays the ratio of observed-to-expected delays in the traditional comparison with luminosity. We see an apparent anticorrelation with luminosity when we split before and after panel 6. This anticorrelation is significant for the higher redshift bins, while the lack of significant correlation detected in lower redshift bins is perhaps unsurprising given the lower number statistics and higher uncertainties. This two population behaviour is more evident in our analysis as the close separation of our filters makes us particularly sensitive to deviations from the generally observed $\tau \propto \lambda^{4/3}$ trend (S. J. Collier et al. 1998; M. M. Fausnaugh et al. 2016). Previous survey studies that do not explore as high redshifts, and either use more widely separated filters (Y. Homayouni et al. 2019; H. W. Sharp et al. 2024) or normalize over several passbands (Z. Yu et al. 2020a; H. Guo et al. 2022b), are more strongly dominated by the overall wavelength trend and are less sensitive to local deviations, plausibly smoothing out the dependence we see here into a general luminosity anticorrelation.

5.2 Eddington ratio dependence

The N. I. Shakura & R. A. Sunyaev (1973) disc is the prevailing paradigm for observers modelling accretion in AGNs. The model is, in part, predicated on the assumptions of a thin disc (disc scale height $H_{\text{D}} \ll R$) that is radiatively efficient (all viscous heat can be dissipated radiatively from the disc surface) and where gas follows Keplerian orbits (there is negligible interaction between neighbouring annuli). For large Eddington ratios ($L/L_{\text{Edd}} \geq 0.3$, the requirement to be geometrically thin is violated (A. Laor & H. Netzer 1989). In such highly accreting discs (dubbed slim discs), the local scale height is appreciable (but still $H_{\text{D}} \lesssim R$) and advective cooling along the equatorial plane becomes significant (M. A. Abramowicz et al. 1988; M. A. Abramowicz & P. C. Fragile 2013).

We aim to test for differences in accretion disc structure with Eddington ratio using inter-band time delays. To do this, we limit our subsample to those objects that have high-quality (QUALITY_MBH = 0) black hole measurements using either Mg II or H β lines. We prefer Mg II derived masses in our high-redshift bins as C IV derived masses can be overestimated, particularly in quasars with high C IV blueshift (L. Coatman et al. 2016). This choice also ensures that derived Eddington ratios are comparable across a wider range of bins. We transition to using H β -derived masses exclusively when $z \leq 0.72$ (panels 7–10) as more quasars have valid black hole measurements in this region (compared to Mg II). To separate our subsamples we chose a cut-off of $L/L_{\text{Edd}} = 0.2$ as significant changes in the ionizing continuum behaviour are seen above this value (M. J. Temple et al. 2023).

Fig. 14 shows a preference for quasars with higher Eddington ratios to exhibit longer delays. Given that we control for continuum luminosity in our binning, an alternative interpretation is that quasars with larger M_{BH} possess shorter delays. We disfavour this interpretation given the weak predicted dependence of the disc time delay on M_{BH} for fixed luminosity (Fig. 6).

Higher Eddington ratio quasars may have longer delays as a result of possessing slim (rather than thin) discs. When the photon diffusion time-scale in the vertical direction exceeds the viscous inflow time-scale, photons become ‘trapped’ with their energy being advected into the black hole. The photon trapping radius within which this effect is expected is given by (K.-y. Watarai 2006),

$$R_{\text{tr}} \approx 102 \left(\frac{\dot{m}}{100} \right) R_{\text{g}} \quad (11)$$

where $\dot{m} = \dot{M}/\dot{M}_{\text{crit}} = \dot{M}c^2/L_{\text{Edd}} = L/\eta L_{\text{Edd}}$. In this regime where advective cooling dominates ($R < R_{\text{tr}}$), the radiative efficiency decreases relative to a thin disc, with the effective temperature instead following a shallower $T \propto R^{-1/2}$ relation (J.-M. Wang & Y.-Y. Zhou 1999; K.-y. Watarai 2006). As a result, less light is emitted from the inner radii and the delay signal is weighted towards larger distances. The degree of this effect is thus dependent on how much of the relevant emitting region falls within the photon trapping radius. This picture is somewhat complicated by slim discs having sufficiently fast radial velocities to allow optically thick thermal emission within the ISCO (M. A. Abramowicz et al. 1988), which can emit beyond the peak temperature of the equivalent SSD description for large accretion rates. However, this interpretation is likely still valid as, due to the small spatial size of the region, the overall SED is still redder than that of the thin disc (S. Mineshige et al. 2000), requiring

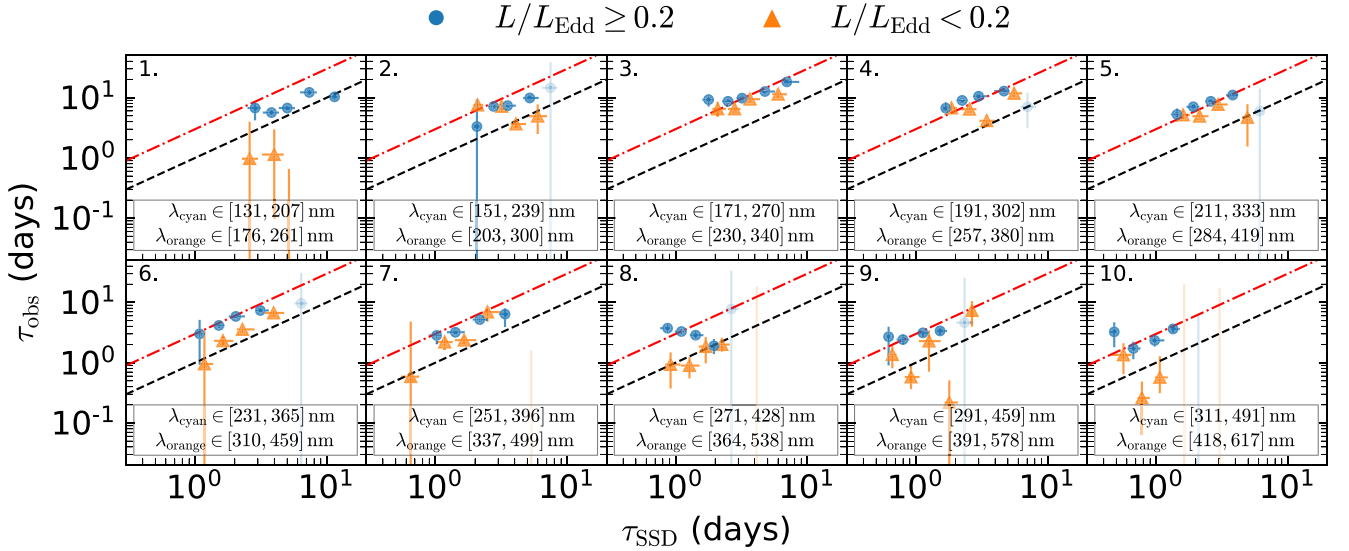


Figure 14. Observed-frame stacked JAVELIN delays versus one (black dashed line) and three (red dash-dotted line) times τ_{SSD} . Delays are split between high (blue circle) and low (orange triangles) Eddington ratios, with low opacity points reflecting multimodal posterior distributions (Section 4.4). The stated wavelength ranges are the filter extents at each panel’s $1/(1+z)$ mid-point. Higher Eddington ratio sources display longer lags.

some net reduction in flux in the inner region that produces the UV continuum.

To reason whether the reduced radiative efficiency of slim discs can explain the lag-Eddington ratio dependence we see, we first estimate the photon trapping radius using equation (11). To convert our previously estimated Eddington ratio to the relevant \dot{m} parameter, we use $\eta = 1/25$ following S. Mineshige et al. (2000) for $L/L_{\text{Edd}} < 2$. We then compare the estimated R_{tr} to the responsivity-weighted radius of the orange passband R_o (c times the orange transfer function centroid), finding only 42.4 per cent of the sample have $R_{\text{tr}}/R_o \geq 0.1$, and of those quasars, 89.8 per cent are in bins 1–3. It is difficult to quantify what effect on the expected delay this ratio will translate to without modelling the temperature profiles explicitly as analytic prescriptions often fail to describe the innermost regions when $\dot{m} < 100$ (K.-y. Watarai 2006). Irrespective, we see that significant fractions of the expected emitting region are only expected to be affected by photon trapping at larger redshifts where the ATLAS filters probe the inner radii. As we see examples of larger lags in higher Eddington ratio quasars across the entire redshift range, photon trapping is unlikely to be the sole driver of the effect seen.

It is unclear whether an alternative explanation lies in the relation between disc scale height and Eddington ratio. As Eddington ratio increases, so does the predicted scale height of the disc (M. A. Abramowicz et al. 1988; A. Laor & H. Netzer 1989) but with a concave radial profile ($\partial^2 H_D / \partial R^2 \leq 0$). Only in the case of a convex disc scale height profile ($\partial^2 H_D / \partial R^2 \geq 0$) do we expect delay times to grow (D. A. Starkey et al. 2023) as the incident angle of X-rays onto the disc becomes stronger at larger radii (as opposed to the concave case where the disc is self-shielding and has the steepest incidence angle in the inner regions). Clearly, more complex modelling is needed to rule out whether Eddington ratio-dependent delays can be explained by structural changes in the accretion disc over other potential confounding effects.

Given the large degree of diffuse BLR contamination in our continuum delays, the observed Eddington ratio dependence may only indirectly trace accretion disc changes and instead be primarily driven by changes in BLR structure. Eddington ratio has

been identified as the primary driver of spectroscopic diversity in AGNs as categorized by the four-dimensional eigenvector 1 (4DE1) space (T. A. Boroson & R. F. Green 1992; J. W. Sulentic et al. 2000, 2007; P. Marziani et al. 2010). This highlights a key dependence of BLR structure on L/L_{Edd} (possibly through BLR column density; X.-B. Dong et al. 2011). Detailed theoretical work modelling BLR structure is beyond the scope of this paper but is necessary to help disentangle how much of our measured inter-band delays belong to the BLR and how much residual accretion disc signals will depend on Eddington ratio.

5.3 Iron contamination

Motivated by the presence of broad line emission in our continuum delays (Section 5.1), we explore whether the relative strength of iron emission in our quasars influences the recovered inter-band delays. As reverberation mapping lags having been successfully measured for the optical Fe II complex (A. J. Barth et al. 2013), some fraction of iron emission must respond to variations in the ionizing continuum and contaminate our inter-band continuum delays. As one of the most distant BLR regions (A. J. Barth et al. 2013; D. Chelouche et al. 2014; M. Marinello et al. 2016), Fe II emission may have an appreciable effect on our measured delays even with smaller EWs ($< 200 \text{ \AA}$).

We measure lags using the Fe II EW divided subsets constructed in Section 2.1.4, finding longer delays in quasars with stronger optical Fe II emission and no discernible difference with UV Fe II (Fig. 15). The FWHM of the UV and optical iron complexes are comparable in magnitude and correlated across the AGN population (C. Hu et al. 2008; H. A. N. Le & J.-H. Woo 2019; J. Kovačević-Dojčinović & L. Č. Popović 2015), suggesting these two complexes originate from similar regions within the BLR. However, the emission behaviour between the two regimes differs, with the UV iron complex emitting more asymmetrically (emitted preferentially towards the centre) due to the higher optical depths and optical iron emission being more isotropic (G. J. Ferland et al. 2009), reconciling previous issues with the

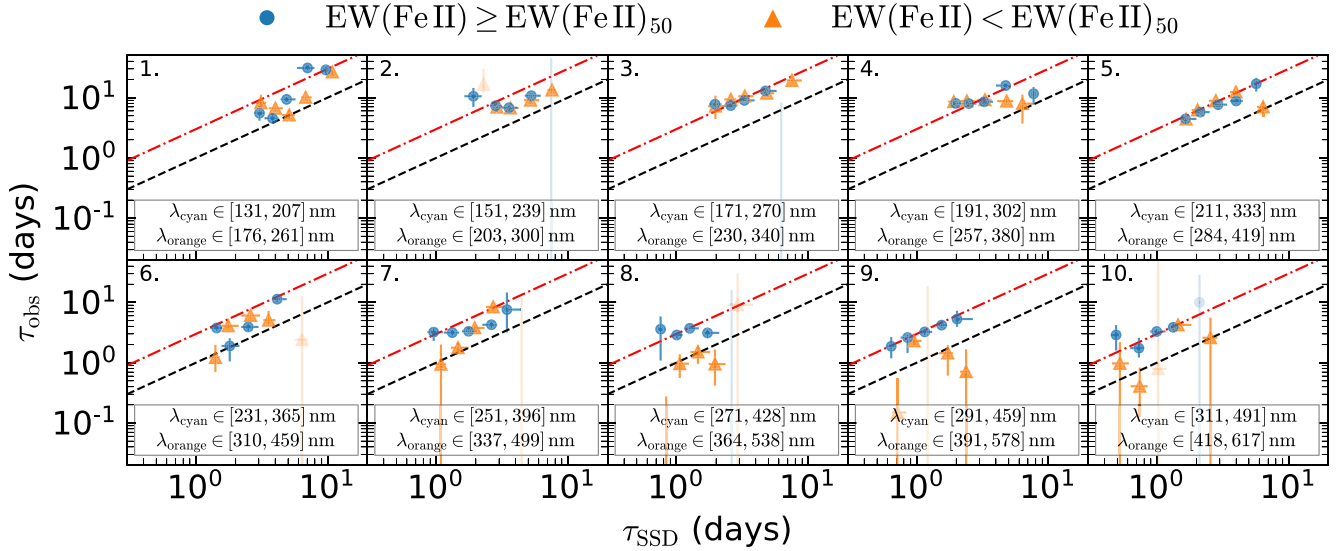


Figure 15. Similar to Fig. 14, but now split by Fe II EW. Higher optical Fe II EW sources display longer lags, with no notable difference when splitting by UV Fe II EW.

UV/optical emission ratio (J. A. Baldwin et al. 2004). The two regimes are also expected to behave differently at the high column densities required to explain the systematic redshifts of Fe II (C. Hu et al. 2008; H. A. N. Le & J.-H. Woo 2019). At these high densities, optical Fe II emission (relative to H β) is expected to increase with increasing column density (which is the theorized connector between L/L_{Edd} and BLR spectral properties X.-B. Dong et al. 2011), whereas UV Fe II emission (relative to H β) saturates (G. J. Ferland et al. 2009). Reflecting previous literature (X.-B. Dong et al. 2011), we also see a moderate correlation between optical Fe II EW and L/L_{Edd} and weak anticorrelation between UV Fe II and L/L_{Edd} . As the EWs between the UV and optical complexes are of similar magnitude and show no redshift evolution (Fig. 3), a delay dependence on optical Fe II and not UV is likely due to confounding factors from the systemic differences in their physics rather than a differing degree of continuum band contamination to our delays. More theoretical modelling is needed to rule out whether the effect seen here may be due to optical Fe II reprocessing.

We aim to examine a delay dependence on BLR structure more directly by using R_{Fe} , a principal variable in classifying AGN spectroscopic diversity (J. W. Sulentic et al. 2000; P. Marziani et al. 2010). We measure delays for the R_{Fe} divided subsets established in Section 2.1.4. As expected (due to the inherent correlation with L/L_{Edd}), we find longer lags in quasars with higher Fe II ratios (in both the optical and UV regimes; Fig. 16). Fig. 17 shows the distribution of Eddington ratios in our quasars for the optical plane of the 4DE1 space and its UV analogue. Broadly following theoretical predictions for the optical diagnostics (R. Zamanov & P. Marziani 2002), we see our quasars are in a regime where Eddington ratio is more strongly separated by H β FWHM than R_{Fe} . We still do see appreciable Eddington ratio separation by subdividing at the median optical R_{Fe} value, confirming that the distributions and mean values of $\log(L/L_{\text{Edd}})$ are significantly different with a Kolmogorov–Smirnov test and two-sample t -test, respectively. We note that this separation in Eddington ratio is also present with UV R_{Fe} , but to a much lesser degree, perhaps owing to the weaker relationship between the UV Fe II emission and Eddington ratio.

Despite the much weaker relationship with Eddington ratio, the UV R_{Fe} -divided subset achieves a similar magnitude effect on inter-band delays to dividing directly by Eddington ratio (comparing panels 1–6 in Figs 14 and 16). As the ratio of UV Fe II EW to Mg II EW may serve as a proxy for chemical enrichment (G. De Rosa et al. 2011), this may signal a metallicity dependence in our delays (although this ratio is dependent on excitation conditions in the BLR and not solely on Fe abundance; J. A. Baldwin et al. 2004). Discs with high iron content may have inherently different disc structure (Y.-F. Jiang, S. W. Davis & J. M. Stone 2016) with more powerful intrinsic disc variability (Y.-F. Jiang & O. Blaes 2020). If such a metallicity dependence is present, it runs counter to the previously seen anticorrelation of continuum delays with Fe II/Mg II EW ratio (Y.-F. Jiang et al. 2017). We, however, have a much larger sample size to simultaneously control for luminosity and wavelength effects. Unfortunately, we lack the sensitivity to further divide our sample to separate out the potential UV R_{Fe} effect from the L/L_{Edd} effect and any other BLR structural dependencies.

5.4 High outflow lags

The physical conditions in AGNs can give rise to outflowing winds as material is ejected from the accretion disc. Such a model is generally used to explain the blueshifted C IV absorption (N. Murray et al. 1995) and emission (K. M. Leighly 2004; G. T. Richards et al. 2011) features seen in the ultraviolet spectra of AGNs. Evidence for C IV outflows is often more strongly present in high mass, high Eddington ratio AGNs (A. L. Rankine et al. 2020; M. J. Temple et al. 2023) where the ionizing continuum is softer (as indicated by decreased He II EW) and radiation line-driving mechanisms are likely dominant (M. Giustini & D. Proga 2019).

The presence of disc winds may help explain the inter-band delay behaviour seen here and in the literature. The ejection of material from the disc creates a radially dependent mass-accretion rate, flattening the temperature profile and requiring hotter overall normalization in the outer disc to recover observed luminosities. In this picture, a greater fraction of the variable and

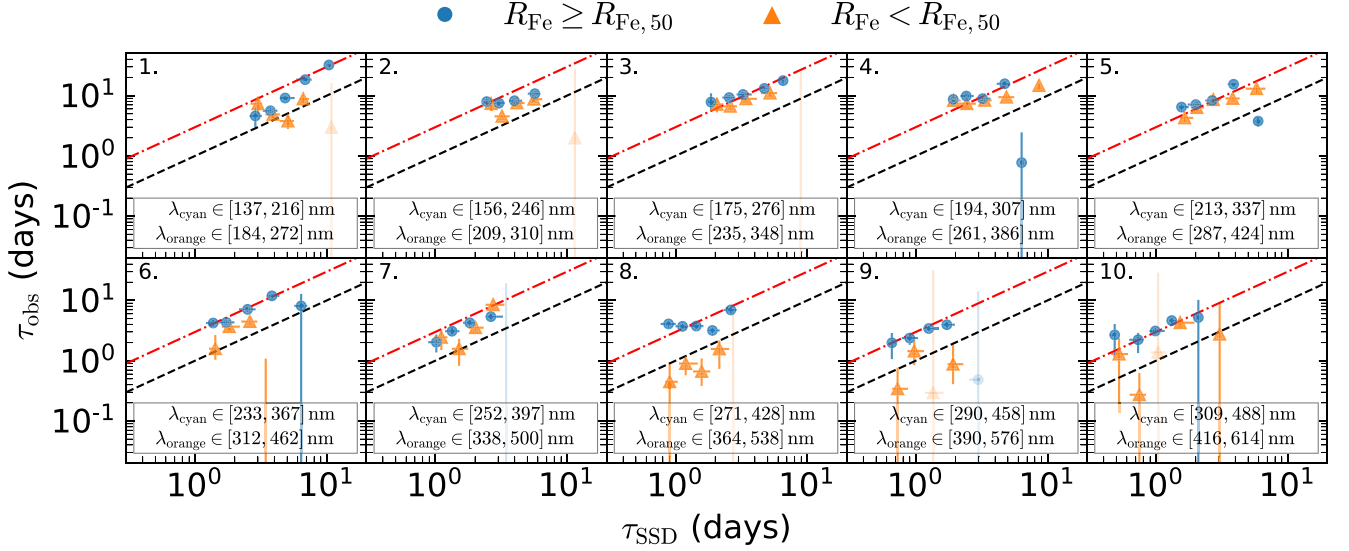


Figure 16. Similar to Fig. 14, but now split by the ratio of Fe II EW to Mg II (panels 1–6)/H β (panels 7–10) EW. Longer lags are seen in quasars with larger iron ratios.




Figure 17. Distribution of quasars in the plane of low-ionization line FWHM and iron ratio. The distribution is coloured by the median $\log(L/L_{\text{Edd}})$, illustrating the correlation with R_{Fe} . This correlation is less strong with the UV R_{Fe} diagnostic.

static emission comes from more distant radii than expected in the SSD model, explaining the inflated size measurements (Y.-P. Li et al. 2019; M. Sun et al. 2019). Alternatively, this spectral reddening could shorten delays, with a weaker ionizing continuum (for fixed UV–optical luminosity) lessening the contribution from variable BLR emission to continuum band delays (H. Netzer 2025).

The effects of winds on continuum-band delays are generally discussed in the context of quasars with broad absorption lines (BAL quasars) in the UV or high column density obscuration in the X-ray. Both properties are variable on the time-scales of days to years, with the two often varying in tandem (J. Ebrero et al. 2016; M. Mehdipour et al. 2022; E. R. Partington et al. 2023). These obscuring winds have been linked to longer continuum (C. Lewin et al. 2024, 2025) and UV BLR delays (Y. Homayouni et al. 2024). This correlation is interpreted as the ejected material re-processing the incident X-ray/disc emission (C. Lewin et al. 2025) which has the potential to explain the increased delay signal seen in AGNs (S. Hagen & C. Done 2023; S. Hagen, C. Done & R. Edelson 2024).

Given the lack of available time-coincident observations of any X-ray obscuring or UV absorbing winds in our large sample, we test for disc wind effects in our lags via a proxy: strong C IV emission line outflows. Reconstructed BAL quasar C IV emission line profiles are co-located with their non-BAL counterparts in C IV EW-blueshift space, with a corresponding non-BAL quasar available with equivalent properties (luminosity, black hole mass, Eddington ratio, He II EW) for each BAL quasar (A. L. Rankine et al. 2020). This suggests that these two classes of quasars are not distinct but come from the same underlying population. Observing a broad absorption trough may thus be a phase or time-dependent property, with several AGNs seen transitioning between BAL and non-BAL states (N. Filiz Ak et al. 2012; J. A. Rogerson et al. 2018; Sameer et al. 2019). By splitting our sample by C IV emission line blueshift, we hope to preferentially select a population of quasars with the largest potential effect should they enter a BAL state, as the strongest absorption troughs are seen in quasars with more blueshifted C IV emission (A. L. Rankine et al. 2020; P. Rodríguez Hidalgo & A. L. Rankine 2022).

For our subsample of quasars with good quality C IV blueshift measurements, we split our sample according to Section 2.1.5. This maintains reasonable sample sizes in panels 1–3 of our binning scheme. We redefine our five uniform $\log(\lambda L_{3000})$ bins to reflect the changes in luminosity range. Without redistributing objects, we do not have sufficient sample counts in our lowest and highest luminosity bins. Having only three bins per rest-frame

panel inhibits our ability to assess if trends are real or spurious under the uncertainties. High C IV blueshift bins do have notably lower counts across the parameter space.

Fig. 18 (left) compares delays measured in quasars with and without strong C IV emission line outflows. Overall, delays are consistent with the more well-constrained main sample, although there is larger scatter in the more blueshifted C IV bins likely owing to the smaller number statistics. A slight preference for longer lags in quasars with more blueshifted C IV emission is seen in panels 1 and 2, with an inversion of this behaviour in panel 3. Such rapid changes in behaviour at adjacent redshifts imply a localized wavelength effect, perhaps stemming from the direct contamination of the cyan filter by the C IV line. As C IV blueshift is anticorrelated with equivalent width, it is difficult to disentangle the effects of the two line properties for the contaminated redshift bins. If the default behaviour implies longer lags in quasars with less blueshifted (i.e. larger EW) C IV emission (i.e. panel 3), then the shorter delays observed in panels 1 and 2 may stem from the greater contribution from the more distant BLR to the leading cyan band. This conclusion is difficult to verify as there is only one uncontaminated redshift bin with corresponding C IV information available.

Physical explanations for longer lags in our low blueshift populations may lie in the link between C IV morphology and the SED of the EUV continuum (M. J. Temple et al. 2023). C IV blueshift is anticorrelated with the He II EW (A. Baskin, A. Laor & F. Hamann 2015; A. L. Rankine et al. 2020; M. J. Temple et al. 2023) which, as a recombination line, serves as a proxy for the number of ionizing photons above 54 eV. Thus, higher C IV may signal a softer/less luminous ionizing continuum, which would result in a smaller $L_{\text{diffuse}}/L_{\text{disk}}$ and measured delay (H. Netzer 2025). However, a delay dependence on other ionizing continuum proxies ($L[\text{O III}]/L(\text{H}\beta)$) has not been observed (H. W. Sharp et al. 2024). Alternatively, He II EW may probe the BLR covering fraction seen by the EUV continuum (although the former interpretation is favoured; see discussion in A. Baskin, A. Laor & F. Hamann 2013). Quasars with less blueshifted C IV emission (weaker He II EW) would then have larger BLR covering fractions and thus longer delays (H. Netzer 2022). Interestingly, we confirm with a non-equal variance two-sample t -test that the high C IV blueshift objects have, on average, larger Eddington ratios. Again, with only one band not directly probing C IV emission it is difficult to ascertain whether this is a universal trend acting in opposition to the previously seen Eddington ratio dependence (Section 5.2) and the C IV connection to the 4DE1 space (J. W. Sulentic et al. 2000, 2007).

We discuss an alternate hypothesis that longer lags are seen in our low C IV blueshift sample due to contamination from BAL quasars. While the fraction of quasars with broad absorption troughs increases with C IV emission line blueshift for fixed EW, this fraction peaks below our chosen cut-off blueshift (see fig. 10 of A. L. Rankine et al. 2020). While quasars displaying broad absorption troughs in C IV were removed from this subsample, as this phenomenon is time-varying, our C IV separated sample may achieve a greater current fraction of BAL quasars in weakly blueshifted C IV bins. This may explain our previous interpretation of longer delays in these bins with previous evidence that disc wind obscuration increases delay times (C. Lewin et al. 2024, 2025).

We attempt to probe delays in BAL quasars directly from their historical SDSS flags. HiBALS, which show absorption troughs for high-ionization lines such as C IV, make up only 7 per cent

of our sample in the three highest redshift bins (with no information available at lower redshifts). We compare lags derived from only these historic BALs to a non-BAL sample (removing LoBALS, HiBALS, and those whose status was unmeasurable) in Fig. 18 (right). Results are broadly consistent with HiBAL quasars exhibiting longer delays at the available wavelengths, although these delays are less well constrained with such low sample counts (some equivalent HiBAL bins are empty). Further dispersion is expected as we use potentially outdated information about the time-variant BAL states. As our C IV information is obtained from observed-frame optical spectra, we are unable to investigate whether this behaviour extends to longer wavelengths as has been seen in other studies (C. Lewin et al. 2024, 2025), highlighting the challenges of survey-based investigations into disc wind properties. Clearly, future continuum RM campaigns with coincident disc wind measurements and more precise/wider wavelength coverage photometry are needed to investigate the role of disc winds in inter-band delay measurements.

5.5 Colouring lags by residual colour

While the standard model of N. I. Shakura & R. A. Sunyaev (1973) accretion discs broadly explains the shape and strength of AGN UV–optical continua, it fails to capture spectral behaviour in detail. The NUV and FUV continuum slopes of AGNs are often redder than predicted (E. W. Bonning et al. 2007; S. W. Davis, J.-H. Woo & O. M. Blaes 2007) and appear bounded below the expected $\alpha = 1/3$ ($f_{\nu} \propto \nu^{\alpha}$) slope (X. Xie et al. 2016). Potential reasons for the unexpectedly red continua include advection cooling in the disc (H. Netzer & B. Trakhtenbrot 2014; A. Kubota & C. Done 2019), disc-wind outflows (O. Slone & H. Netzer 2012; A. Laor & S. W. Davis 2014), and dust reddening (G. T. Richards et al. 2003; D. M. Capellupo et al. 2015; D. Baron et al. 2016; X. Xie et al. 2016), all of which may induce longer lags. As such, we search for a colour dependence in continuum delays here.

In Section 2.1.3, we established a broad distribution of colour in our sample relative to a redshift-dependent median. After removing the red-tail, this distribution is approximately Gaussian. We divide this residual colour distribution into three uniform bins with the resulting lags shown in Fig. 19. We see delays are often ordered by their residual colour, with redder quasars possessing longer lags. The magnitude and coherence of this observed effect is not uniform across our parameter space. The largest scatter is seen in the lowest redshift bins where we have the smaller number statistics (binning in $1/(1+z)$) and residual host-galaxy contamination is most likely present. It is difficult to disentangle whether this behaviour is separate from that seen in Section 5.2 as several redshift bins have significant Spearman rank correlations between Eddington ratio and residual colour (higher Eddington ratio quasars are redder). With the larger scatter observed here in the colour-separated delays, further subdividing this sample by multicollinear explanatory variables is beyond the limits of our data.

Internal AGN reddening presents a potential explanation for these colour-dependent lags. Redder residual colours would then indicate AGNs experience a larger degree of dust extinction and have their luminosity more greatly underestimated (and hence underestimating their delay times; C. M. Gaskell 2017; C. M. Gaskell et al. 2023). As we control for the observed luminosity, this presents as redder quasars having longer delays in their respective panels. The hypothesis of dust reddening driving the ob-

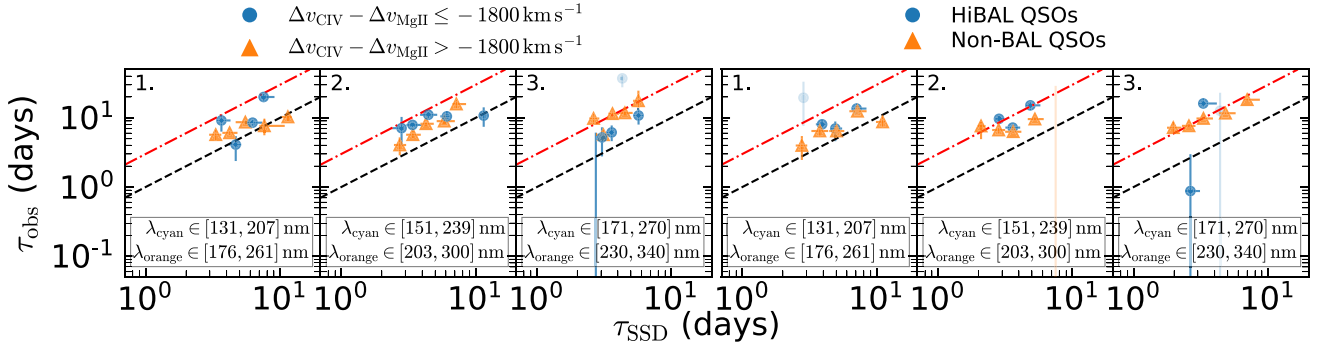


Figure 18. Similar to Fig. 14, now splitting by C IV emission line blueshift (left) and BAL status (right). Longer delays are seen in quasars with higher (lower) C IV blueshift in panels 1–2 (panel 3) and in quasars with C IV BAL troughs (all panels).

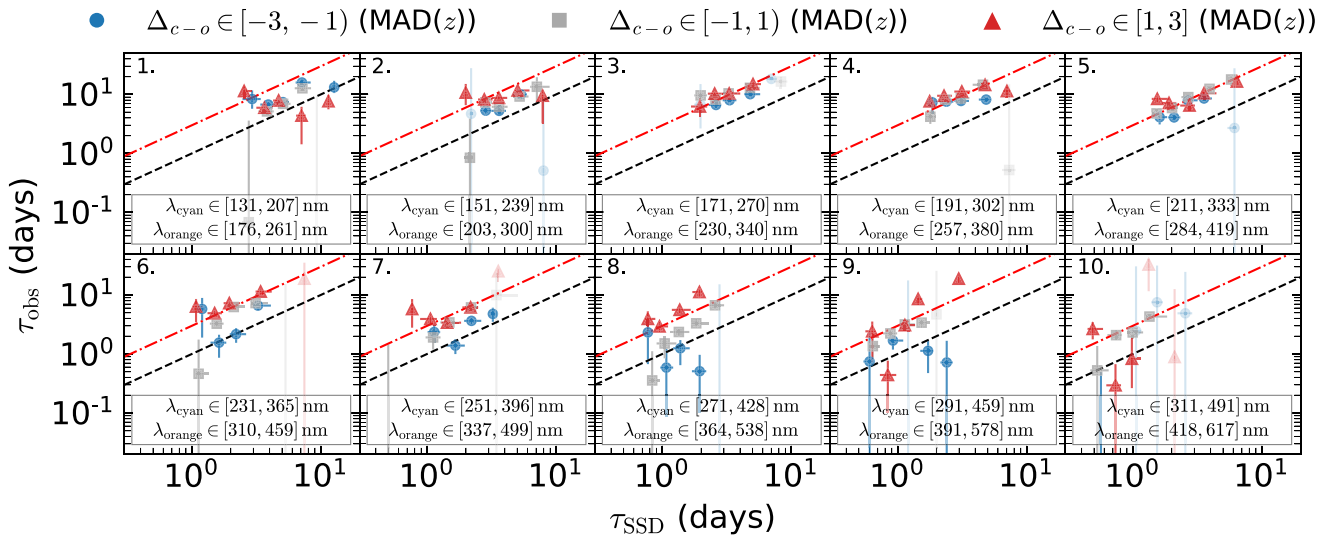


Figure 19. Similar to Fig. 14, now split by quasar residual colour into bluer, average, and redder quasars, respectively. Longer delays are seen in quasars with increasingly red colours.

served delayed behaviour would be strongly supported if the red-tail quasars excluded in Section 2.1.3 exhibited the longest lags. Quasars belonging to the red-tail of the residual colour distribution are thought to be strong evidence of dust extinction in AGNs (G. T. Richards et al. 2003). As a consequence of this dust reddening, these very red quasars are among the faintest in our sample with many hovering just above the minimum *Gaia* R_p magnitude ($R_p \leq 18$). As *Gaia* R_p occupies similar wavelengths as the ATLAS orange band, these very red quasars will be even fainter in the cyan band, which is the limiting passband in this analysis. As a result, there is insufficient signal to constrain delays for quasars in the dust-reddened tail using ATLAS data, even with our stacked methodology. The future Legacy Survey of Space and Time (LSST; Ž. Ivezić et al. 2019) will offer a fantastic opportunity to study these very red quasars with its superior photometric sensitivity. If dust reddening was responsible for colour-dependent delays, then this may open up the exciting prospect of using continuum delays to measure reddening laws in quasars.

With evidence of diffuse nebular emission contaminating our delays, the observed colour-dependent lags may alternatively be explained by a varying BLR covering factor. Quasars with larger BLR covering factors would possess more strongly emitting BLR

regions relative to the accretion disc, lengthening delays (H. Netzer 2022). Such spectral signatures would also then be expected in the time-averaged SEDs and photometric colours. In this picture, our colour dependence has a more nuanced interpretation (unlike with the dust reddening case which is largely monotonic with wavelength), as our colours are derived at a fixed observed-frame wavelength. As the Balmer and Paschen continua are redder than the underlying disc continuum, when both filters probe the same nebular emission, quasars with larger BLR covering factors would appear redder and have longer delays. When filters straddle the Balmer jump, quasars with larger covering factors would have larger nebular emission contributions to the cyan filter, making the time-averaged photometric colours appear more blue and shortening the observed inter-band delay. As we only have two filters, the differential delay operator that provides evidence for this non-monotonic wavelength behaviour (Fig. 6) does not rule out larger delay signatures in an absolute sense. For example, redder quasars in panel 8 of Fig. 19 are seen to be at least a factor three larger than anticipated in the SSD. This same (spanning the Balmer jump) wavelength range is where the observed-to-predicted delay ratio experiences its minima in our main sample delays, perhaps signalling that the underlying disc

lags (with less BLR contamination, i.e. redder residual colours at this redshift) are still larger than previously anticipated. This difference is more easily reconciled without revisions to accretion disc theory in the dust extinction picture which behaves similarly at all wavelengths.

6 SUMMARY AND OUTLOOK

We have analysed high-cadence ATLAS light curves to measure some of the most precise population-level inter-band continuum delays for ~ 9500 of the most UV–optical luminous quasars in the sky. With our large sample size, we were able to correlate delays over a range of quasar properties and help break degeneracies between several population literature models. We summarize the main findings of this work as follows:

(i) We test our ability to recover continuum band delays on simulated data and demonstrate that selecting delays based on significance introduces a bias towards longer lags (Fig. 9). When we combine posterior estimates without significance cuts, no bias is introduced, but with our data, the delay uncertainties are very large. Only by stacking delay information at the inference level, rather than combining individual posteriors, do we greatly improve our ability to measure the underlying signal. Some benefit is seen on the simulated data set when analysing only the more well-correlated light curves, but we show that the inter-band correlation is dependent on quasar colour (relative to a redshift-dependent median; Fig. 10). As such, cutting by correlation biases our sample towards bluer quasars.

(ii) Using the stacked methodology, we compile lags across a range of redshifts ($z \sim 0.3\text{--}2.5$) and luminosities ($\lambda L_{3000} \sim 10^{45}\text{--}10^{47}\text{erg s}^{-1}$). When we compare quasars within a given redshift bin, we do not see an anticorrelation between the disc size discrepancy and luminosity (Fig. 11). These findings disfavour the CHAR model to explain our delays which predicts the aforementioned anticorrelation when controlling for light curve duration and wavelength as done here. We further find a non-monotonic trend of delay amplitude with rest-frame wavelength, disfavoured the K21 model as the primary source of elevated disc sizes as it predicts a monotonic change. We find our results are naturally explained by diffuse contamination from the BLR, suggesting it is commonplace in the wider AGN population.

(iii) After separating our sample by Eddington ratio, we observed that quasars with higher Eddington ratios often prefer longer lags (Fig. 14). This behaviour is observed across our redshift range, suggesting it is not solely driven by the advection cooling of slim discs, which is expected to contribute significantly only at higher redshifts in our sample. The increasing delay magnitude is also unlikely to be from changes in disc scale height with increasing Eddington ratio as the scale height profiles in slim discs are predicted to be concave (M. A. Abramowicz et al. 1988; A. Laor & H. Netzer 1989) and thus produce shorter delays (D. A. Starkey et al. 2023). Instead, given the strong evidence for diffuse contamination in our delays, the observed Eddington ratio dependence likely stems from its connection to BLR structure and the 4DE1 space (J. W. Sulentic et al. 2007; P. Marziani et al. 2010).

(iv) We find evidence that delay magnitude increases with optical Fe II EW, but not UV Fe II EW. Given the similar equivalent width of the two regimes (and the lack of redshift dependence; Fig. 3), we believe this is less likely to be evidence of a direct response by the iron BLR, but instead related to BLR structure as UV Fe II is less dependent on L/L_{Edd} than the optical complex (X.-

B. Dong et al. 2011). We further find longer delays in quasars with larger $R_{\text{Fe}} = \text{Fe II}/\text{H}\beta$ (and the UV analogue Fe II/Mg II) which is a principal variable in explaining AGN spectroscopic diversity (J. W. Sulentic et al. 2000, 2007). As we see a similar magnitude effect with Fe II/Mg II as L/L_{Edd} , while the two are more weakly related (Fig. 17), this may signal a possible metallicity dependence in our lags (G. De Rosa et al. 2011; J. Shin et al. 2019).

(v) Motivated by recent observations linking disc winds to excess continuum delays (C. Lewin et al. 2024, 2025), we test to see if delays change with C IV emission line blueshift. We see evidence that quasars with smaller C IV blueshifts have longer lags, with this relationship inverting when the C IV line enters the cyan band (Fig. 18, left). However, this relationship is difficult to verify with only one panel not directly probing C IV emission and considering its inconsistency with the Eddington ratio and 4DE1 interpretation. We additionally examine the effect of broad absorption lines by comparing delays between historically confirmed BAL and non-BAL quasars. We find some evidence of longer delays in BAL quasars, although this relation is more tenuous given the low sample counts and transient nature of BALs. This transient nature may help explain our C IV results as BAL frequency is highest in the low C IV blueshift subset of our division of the C IV EW-blueshift plane (A. L. Rankine et al. 2020).

(vi) We find a tendency for quasars with redder continua to have longer lags (Fig. 19). There are several plausible mechanisms that can create redder continua, including: slim discs, disc winds, and dust reddening. We explored the first two mechanisms independently (see points iii and v, respectively). We are unable to quantify whether residual colour or Eddington ratio is the driver behind this trend, with both potentially being explained through either dust reddening or a varying BLR covering factor. The hypothesis of dust reddening would be more strongly supported if quasars in the dust-reddened tail displayed the longest lags. However, we are unable to further test this hypothesis as these highly reddened quasars have the lowest SNR.

We have shown that lags may depend on several physical parameters, and disentangling them will require more precise measurements over a larger parameter space than given here. While we have shown how to extract a greater degree of signal from light curves that individually struggle to constrain lags, testing for all the different dependencies is expensive and non-feasible (only so much independent information can be extracted). The upcoming LSST survey will provide an excellent opportunity to further study inter-band continuum delays and correlate them with several physical parameters simultaneously as its superior photometric precision will allow for measuring delays in thousands of AGN individually (A. B. Kovačević et al. 2022; F. Pozo Nuñez et al. 2023). The deeper photometry of LSST will also allow us to examine delays in the dust-reddened tail of the AGN colour distribution, further shedding light onto the potential role of dust extinction in continuum RM.

ACKNOWLEDGEMENTS

We kindly thank the anonymous referee for their insightful comments and assistance in improving this manuscript. ZS was supported by the Australian Government Research Training Program (RTP). KWS acknowledges funding from the Royal Society. AC acknowledges the support from ANID, Chile, through grants AIM23-0001 (Millennium Science Initiative) and FONDECYT No. 1251692. This work used data from the University of Hawaii’s

ATLAS project, funded through NASA grants NN12AR55G, 80NSSC18K0284, and 80NSSC18K1575, with contributions from the Queen's University Belfast, STScI, the South African Astronomical Observatory, and the Millennium Institute of Astrophysics, Chile.

This work made use of SDSS spectroscopic data. Funding for the Sloan Digital Sky Survey IV was provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High Performance Computing at the University of Utah. The SDSS website is <http://www.sdss.org>. SDSS-IV is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, Center for Astrophysics| Harvard & Smithsonian, the Chilean Participation Group, the French Participation Group, Instituto de Astrofísica de Canarias, the Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU)/University of Tokyo, the Korean Participation Group, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatório Nacional/MCTI, The Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University.

This work made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC; <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC was provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

DATA AVAILABILITY

The data underlying the analysis in this paper may be shared upon reasonable request.

REFERENCES

- Abramowicz M. A., Fragile P. C., 2013, *Living Rev. Relativ.*, 16, 1
 Abramowicz M. A., Czerny B., Lasota J. P., Szuszkiewicz E., 1988, *ApJ*, 332, 646
 Alonso-García J., Mateo M., Sen B., Banerjee M., Catelan M., Minniti D., von Braun K., 2012, *AJ*, 143, 70
 Alston W. N. et al., 2020, *Nat. Astron.*, 4, 597
 Baldwin J. A., 1977, *ApJ*, 214, 679
 Baldwin J. A., Ferland G. J., Korista K. T., Hamann F., LaCluzé A., 2004, *ApJ*, 615, 610
 Baron D., Stern J., Poznanski D., Netzer H., 2016, *ApJ*, 832, 8
 Barth A. J. et al., 2013, *ApJ*, 769, 128
 Barthel P. D., 1989, *ApJ*, 336, 606
 Barvainis R., 1987, *ApJ*, 320, 537
 Baskin A., Laor A., Hamann F., 2013, *MNRAS*, 432, 1525
 Baskin A., Laor A., Hamann F., 2015, *MNRAS*, 449, 1593
 Beard M. W. J. et al., 2025, *MNRAS*, 537, 293
 Becker R. H., White R. L., Helfand D. J., 1995, *ApJ*, 450, 559
 Bentz M. C. et al., 2013, *ApJ*, 767, 149
 Blackburne J. A., Kochanek C. S., Chen B., Dai X., Chartas G., 2014, *ApJ*, 789, 125
 Blandford R. D., McKee C. F., 1982, *ApJ*, 255, 419
 Bonning E. W., Cheng L., Shields G. A., Salvander S., Gebhardt K., 2007, *ApJ*, 659, 211
 Boroson T. A., Green R. F., 1992, *ApJS*, 80, 109
 Breedt E. et al., 2009, *MNRAS*, 394, 427
 Buisson D. J. K. et al., 2018, *MNRAS*, 475, 2306
 Caballero-García M. D., Papadakis I. E., Dovčiak M., Bursa M., Svoboda J., Karas V., 2020, *MNRAS*, 498, 3184
 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., Bentz M. C., Kara E., 2021, *iScience*, 24, 102557
 Cackett E. M., Zoghbi A., Ulrich O., 2022, *ApJ*, 925, 29
 Cackett E. M. et al., 2023, *ApJ*, 958, 195
 Capellupo D. M., Netzer H., Lira P., Trakhtenbrot B., Mejía-Restrepo J., 2015, *MNRAS*, 446, 3427
 Chanchaiworawit K., Sarajedini V., 2024, *ApJ*, 969, 131
 Chelouche D., Rafter S. E., Cotlier G. I., Kaspi S., Barth A. J., 2014, *ApJ*, 783, L34
 Chen J., Sun M., Zhou S., 2024a, *Res. Notes Am. Astron. Soc.*, 8, 253
 Chen J., Sun M., Zhang Z.-X., 2024b, *ApJ*, 962, 134
 Clavel J. et al., 1991, *ApJ*, 366, 64
 Coatman L., Hewett P. C., Banerji M., Richards G. T., 2016, *MNRAS*, 461, 647
 Collier S. J., 1998, Phd thesis, Univ. St Andrews
 Collier S. J. et al., 1998, *ApJ*, 500, 162
 Condon J. J., Cotton W. D., Greisen E. W., Yin Q. F., Perley R. A., Taylor G. B., Broderick J. J., 1998, *AJ*, 115, 1693
 Davis S. W., Woo J.-H., Blaes O. M., 2007, *ApJ*, 668, 682
 De Rosa G., Decarli R., Walter F., Fan X., Jiang L., Kurk J., Pasquali A., Rix H. W., 2011, *ApJ*, 739, 56
 Dean R. T., Dunsmuir W. T. M., 2016, *Behav. Res. Methods*, 48, 783
 Dexter J. et al., 2019, *ApJ*, 885, 44
 Dong X.-B., Wang J.-G., Ho L. C., Wang T.-G., Fan X., Wang H., Zhou H., Yuan W., 2011, *ApJ*, 736, 86
 Ebrero J., Kriss G. A., Kaastra J. S., Ely J. C., 2016, *A&A*, 586, A72
 Edelson R. et al., 2019, *ApJ*, 870, 123
 Edelson R., Peterson B. M., Gelbord J., Horne K., Goad M., McHardy I., Vaughan S., Vestergaard M., 2024, *ApJ*, 973, 152
 Event Horizon Telescope Collaboration, 2019, *ApJ*, 875, L1
 Fausnaugh M. M. et al., 2016, *ApJ*, 821, 56
 Fausnaugh M. M. et al., 2017, *ApJ*, 840, 97
 Ferland G. J., Hu C., Wang J.-M., Baldwin J. A., Porter R. L., van Hoof P. A. M., Williams R. J. R., 2009, *ApJ*, 707, L82
 Filiz Ak N. et al., 2012, *ApJ*, 757, 114
 Fine S. et al., 2013, *MNRAS*, 434, L16
 Fitzpatrick E. L., 1999, *PASP*, 111, 63
 Gardner E., Done C., 2017, *MNRAS*, 470, 3591
 Gaskell C. M., 2017, *MNRAS*, 467, 226
 Gaskell C. M., Anderson F. C., Birmingham S. Á., Ghosh S., 2023, *MNRAS*, 519, 4082
 Georgakakis A. et al., 2024, *MNRAS*, 531, 4524
 Giustini M., Proga D., 2019, *A&A*, 630, A94
 Gravity Collaboration, 2018, *Nature*, 563, 657
 Grier C. J. et al., 2017, *ApJ*, 851, 21
 Guo H. et al., 2022a, *ApJ*, 927, 60
 Guo H., Barth A. J., Wang S., 2022b, *ApJ*, 940, 20
 Hagen S., Done C., 2023, *MNRAS*, 521, 251
 Hagen S., Done C., Edelson R., 2024, *MNRAS*, 530, 4850
 Hall P. B., Sarrrouh G. T., Horne K., 2018, *ApJ*, 854, 93
 Hanawa T., 1989, *ApJ*, 341, 948
 Heinze A. N. et al., 2018, *AJ*, 156, 241
 Hernández Santisteban J. V. et al., 2020, *MNRAS*, 498, 5399

- Homayouni Y. et al., 2019, *ApJ*, 880, 126
Homayouni Y. et al., 2022, *ApJ*, 926, 225
Homayouni Y. et al., 2024, *ApJ*, 963, 123
Hu C., Wang J.-M., Ho L. C., Chen Y.-M., Zhang H.-T., Bian W.-H., Xue S.-J., 2008, *ApJ*, 687, 78
Ivezić Ž. et al., 2019, *ApJ*, 873, 111
Jiang Y.-F., Blaes O., 2020, *ApJ*, 900, 25
Jiang Y.-F., Davis S. W., Stone J. M., 2016, *ApJ*, 827, 10
Jiang Y.-F. et al., 2017, *ApJ*, 836, 186
Jin C., Ward M., Done C., Gelbord J., 2012, *MNRAS*, 420, 1825
Kammoun E., Papadakis I. E., Dovčiak M., Panagiotou C., 2024, *A&A*, 686, A69
Kammoun E. S., Papadakis I. E., Dovčiak M., 2021a, *MNRAS*, 503, 4163
Kammoun E. S., Dovčiak M., Papadakis I. E., Caballero-García M. D., Karas V., 2021b, *ApJ*, 907, 20
Kammoun E. S., Robin L., Papadakis I. E., Dovčiak M., Panagiotou C., 2023, *MNRAS*, 526, 138
Kara E. et al., 2021, *ApJ*, 922, 151
Kara E. et al., 2023, *ApJ*, 947, 62
Kaspi S., Brandt W. N., Maoz D., Netzer H., Schneider D. P., Shemmer O., Grier C. J., 2021, *ApJ*, 915, 129
Kellermann K. I., Verschuur G. L., 1988, *Galactic and Extragalactic Radio Astronomy*, Springer New York, NY
Kellermann K. I., Sramek R., Schmidt M., Shaffer D. B., Green R., 1989, *AJ*, 98, 1195
Kelly B. C., Bechtold J., Siemiginowska A., 2009, *ApJ*, 698, 895
Kelly B. C., Treu T., Malkan M., Pancoast A., Woo J.-H., 2013, *ApJ*, 779, 187
Kessler R. et al., 2015, *AJ*, 150, 172
Koratkar A. P., Gaskell C. M., 1991, *ApJS*, 75, 719
Korista K. T., Goad M. R., 2001, *ApJ*, 553, 695
Korista K. T., Goad M. R., 2019, *MNRAS*, 489, 5284
Kovačević-Dojčinović J., Popović L. Č., 2015, *ApJS*, 221, 35
Kovačević A. B. et al., 2022, *ApJS*, 262, 49
Kozłowski S., 2016, *ApJ*, 826, 118
Krolik J. H., Horne K., Kallman T. R., Malkan M. A., Edelson R. A., Kriss G. A., 1991, *ApJ*, 371, 541
Kubota A., Done C., 2019, *MNRAS*, 489, 524
Laor A., Davis S. W., 2014, *MNRAS*, 438, 3024
Laor A., Netzer H., 1989, *MNRAS*, 238, 897
Lawther D., Goad M. R., Korista K. T., Ulrich O., Vestergaard M., 2018, *MNRAS*, 481, 533
Le H. A. N., Woo J.-H., 2019, *ApJ*, 887, 236
Leighly K. M., 2004, *ApJ*, 611, 125
Lewin C., Kara E., Cackett E. M., Wilkins D., Panagiotou C., García J. A., Gelbord J., 2023, *ApJ*, 954, 33
Lewin C. et al., 2024, *ApJ*, 974, 271
Lewin C., Kara E., Panagiotou C., Cackett E. M., Gelbord J., Hernández Santisteban J. V., Horne K., Kriss G. A., 2025, *ApJ*, 993, 245
Li J. et al., 2017, *ApJ*, 846, 79
Li T. et al., 2021, *ApJ*, 912, L29
Li Y.-P., Yuan F., Dai X., 2019, *MNRAS*, 483, 2275
Liu H., Luo B., Brandt W. N., Brotherton M. S., Gallagher S. C., Ni Q., Shemmer O., Timlin III J. D., 2021, *ApJ*, 910, 103
MacLeod C. L. et al., 2010, *ApJ*, 721, 1014
Malik U. et al., 2024, *MNRAS*, 531, 163
Marculewicz M., Sun M., Wu J., Zhang Z., 2023, *ApJ*, 956, 126
Marinello M., Rodríguez-Ardila A., García-Rissmann A., Sigut T. A. A., Pradhan A. K., 2016, *ApJ*, 820, 116
Marziani P., Sulentic J. W., Negrete C. A., Dultzin D., Zamfir S., Bachev R., 2010, *MNRAS*, 409, 1033
Masci F. J. et al., 2019, *PASP*, 131, 018003
McHardy I. M. et al., 2014, *MNRAS*, 444, 1469
McHardy I. M. et al., 2018, *MNRAS*, 480, 2881
Mehdipour M., Kriss G. A., Costantini E., Gu L., Kaastra J. S., Landt H., Mao J., 2022, *ApJ*, 934, L24
Mineshige S., Kawaguchi T., Takeuchi M., Hayashida K., 2000, *PASJ*, 52, 499
Morales A. M., Miller J. M., Cackett E. M., Reynolds M. T., Zoghbi A., 2019, *ApJ*, 870, 54
Morgan C. W., Kochanek C. S., Morgan N. D., Falco E. E., 2010, *ApJ*, 712, 1129
Morgan C. W. et al., 2012, *ApJ*, 756, 52
Morganson E. et al., 2014, *ApJ*, 784, 92
Mosquera A. M., Kochanek C. S., 2011, *ApJ*, 738, 96
Mudd D. et al., 2018, *ApJ*, 862, 123
Murray N., Chiang J., Grossman S. A., Voit G. M., 1995, *ApJ*, 451, 498
Mushotzky R. F., Edelson R., Baumgartner W., Gandhi P., 2011, *ApJ*, 743, L12
Netzer H., 2022, *MNRAS*, 509, 2637
Netzer H., 2025, *MNRAS*, 539, 3242
Netzer H., Trakhtenbrot B., 2014, *MNRAS*, 438, 672
Netzer H. et al., 2024, *ApJ*, 976, 59
Neustadt J. M. M., Kochanek C. S., 2022, *MNRAS*, 513, 1046
Neustadt J. M. M. et al., 2024, *ApJ*, 961, 219
Panagiotou C., Papadakis I., Kara E., Kammoun E., Dovčiak M., 2022a, *ApJ*, 935, 93
Panagiotou C., Kara E., Dovčiak M., 2022b, *ApJ*, 941, 57
Panagiotou C. et al., 2025, *ApJ*, 983, 132
Paolillo M. et al., 2023, *A&A*, 673, A68
Papoutsis M., Papadakis I. E., Panagiotou C., Dovčiak M., Kammoun E., 2024, *A&A*, 691, A60
Partington E. R. et al., 2023, *ApJ*, 947, 2
Partington E. R. et al., 2024, *ApJ*, 977, 77
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
Pooley D., Blackburne J. A., Rappaport S., Schechter P. L., 2007, *ApJ*, 661, 19
Pozo Nuñez F., Bruckmann C., Deesamutara S., Czerny B., Panda S., Lobban A. P., Pietrzyński G., Polsterer K. L., 2023, *MNRAS*, 522, 2002
Pozo Nuñez F., Bañados E., Panda S., Heidt J., 2025, *A&A*, 700, L8
Prokhorenko S. A. et al., 2024, *MNRAS*, 528, 5972
Rakić N., La Mura G., Ilić D., Shapovalova A. I., Kollatschny W., Rafanelli P., Popović L. Č., 2017, *A&A*, 603, A49
Rakshit S., Stalin C. S., Kotilainen J., 2020, *ApJS*, 249, 17
Rankine A. L., Hewett P. C., Banerji M., Richards G. T., 2020, *MNRAS*, 492, 4553
Richards G. T. et al., 2003, *AJ*, 126, 1131
Richards G. T. et al., 2006, *ApJS*, 166, 470
Richards G. T. et al., 2011, *AJ*, 141, 167
Rodríguez Hidalgo P., Rankine A. L., 2022, *ApJ*, 939, L24
Rogerson J. A., Hall P. B., Ahmed N. S., Rodríguez Hidalgo P., Brandt W. N., Filiz Ak N., 2018, *ApJ*, 862, 22
Rybicki G. B., Press W. H., 1992, *ApJ*, 398, 169
Sameer et al., 2019, *MNRAS*, 482, 1121
Schechter P. L., Mateo M., Saha A., 1993, *PASP*, 105, 1342
Schlafly E. F. et al., 2012, *ApJ*, 756, 158
Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, 500, 525
Secunda A., Jiang Y.-F., Greene J. E., 2025, *ApJ*, 984, 19
Selsing J., Fynbo J. P. U., Christensen L., Krogager J.-K., 2016, *A&A*, 585, A87
Shakura N. I., Sunyaev R. A., 1973, *A&A*, 24, 337
Sharp H. W. et al., 2024, *ApJ*, 961, 93
Shen Y. et al., 2015, *ApJS*, 216, 4
Shen Y. et al., 2016, *ApJ*, 831, 7
Shen Y. et al., 2024, *ApJS*, 272, 26
Shin J., Nagao T., Woo J.-H., Le H. A. N., 2019, *ApJ*, 874, 22
Shingles L. et al., 2021, *Transient Name Server AstroNote*, 7, 1.
Simm T., Salvato M., Saglia R., Ponti G., Lanzuisi G., Trakhtenbrot B., Nandra K., Bender R., 2016, *A&A*, 585, A129
Slone O., Netzer H., 2012, *MNRAS*, 426, 656
Smith K. L., Mushotzky R. F., Boyd P. T., Malkan M., Howell S. B., Gelino D. M., 2018, *ApJ*, 857, 141
Starkey D. A., Horne K., Villforth C., 2016, *MNRAS*, 456, 1960
Starkey D. A., Huang J., Horne K., Lin D. N. C., 2023, *MNRAS*, 519, 2754

- Stone Z., Shen Y., 2023, *MNRAS*, 524, 4521
- Stone Z. et al., 2022, *MNRAS*, 514, 164
- Sulentic J. W., Zwitter T., Marziani P., Dultzin-Hacyan D., 2000, *ApJ*, 536, L5
- Sulentic J. W., Bachev R., Marziani P., Negrete C. A., Dultzin D., 2007, *ApJ*, 666, 757
- Sun M., Grier C. J., Peterson B. M., 2018a, PyCCF: Python Cross Correlation Function for Reverberation Mapping Studies, Astrophysics Source Code Library, record ascl:1805.032
- Sun M., Xue Y., Richards G. T., Trump J. R., Shen Y., Brandt W. N., Schneider D. P., 2018b, *ApJ*, 854, 128
- Sun M., Xue Y., Trump J. R., Gu W.-M., 2019, *MNRAS*, 482, 2788
- Sun M. et al., 2020, *ApJ*, 902, 7
- Tang J.-J., Wolf C., Tonry J., Lai S., Yong S. Y., Steyn Z., 2023, *MNRAS*, 523, 4441
- Tang J.-J., Wolf C., Tonry J., 2024, *MNRAS*, 535, 2260
- Temple M. J. et al., 2023, *MNRAS*, 523, 646
- Tie S. S., Kochanek C. S., 2018, *MNRAS*, 473, 80
- Timmer J., König M., 1995, *A&A*, 300, 707
- Tonry J. L. et al., 2018, *PASP*, 130, 064505
- Uttley P., McHardy I. M., Papadakis I. E., 2002, *MNRAS*, 332, 231
- Uttley P., Edelson R., McHardy I. M., Peterson B. M., Markowitz A., 2003, *ApJ*, 584, L53
- Vanden Berk D. E. et al., 2004, *ApJ*, 601, 692
- 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
- Vio R., Wamsteker W., 2001, *PASP*, 113, 86
- Wanders I. et al., 1997, *ApJS*, 113, 69
- Wang J.-M., Zhou Y.-Y., 1999, *ApJ*, 516, 420
- Wang S., Guo H., Woo J.-H., 2023, *ApJ*, 948, L23
- Watarai K.-y., 2006, *ApJ*, 648, 523
- Welsh W. F., 1999, *PASP*, 111, 1347
- Xie X., Shao Z., Shen S., Liu H., Li L., 2016, *ApJ*, 824, 38
- Yang Q. et al., 2020, *ApJ*, 900, 58
- Yao P. Z., Secunda A., Jiang Y.-F., Greene J. E., Villar A., 2023, *ApJ*, 953, 43
- Yu Z. et al., 2020a, *ApJS*, 246, 16
- Yu Z., Kochanek C. S., Peterson B. M., Zu Y., Brandt W. N., Cackett E. M., Fausnaugh M. M., McHardy I. M., 2020b, *MNRAS*, 491, 6045
- Zamanov R., Marziani P., 2002, *ApJ*, 571, L77
- Zoghbi A., Miller J. M., Cackett E., 2021, *ApJ*, 912, 42
- Zu Y., Kochanek C. S., Peterson B. M., 2011, *ApJ*, 735, 80

APPENDIX A: DELAY PIPELINE DEPENDENCY

We investigate what influence our analysis pipeline choices have over the measured delays to ensure that we are not introducing any systematic biases through these decisions. We choose the main sample as our representative validation set for the real data. This sample is where we have the greatest number of statistics and thus where any potential biases may be best separated from noise. Two possible concerns include whether homogenous detrending of light curves removes genuine physical variability and whether binning without concern for cadence suppresses short-term features. Both, if present, would distort the correlation structure and hence our inter-band delays. As such, we display our fiducial choices alongside two other comparative applications of ICCF and JAVELIN with/without binning and without stacking (Fig. A1).

Results from ICCF seem to qualitatively prefer the detrended approach in more panels (1, 6–10) than not (2–5) based on the a priori expectation that delays follow a power-law scaling with luminosity at fixed redshift. We additionally compare ICCF performance with and without detrending on the simulated data set

and find detrending is necessary for lag recovery, reflecting earlier recommendations in the literature (W. F. Welsh 1999). This detrending helps remove the low-frequency components of light curves, which contribute more significantly than high-frequency components to the construction of the CCF (W. F. Welsh 1999; R. Vio & W. Wamsteker 2001), thus improving our ability to constrain delays on short time-scales. We find detrending important here as a form of local normalization, which helps reduce spurious correlations in time series with large autocorrelation (R. T. Dean & W. T. M. Dunsmuir 2016), as is common for systems like AGNs where physical limitations of variability impose dependence on previous epochs (B. C. Kelly et al. 2009).

In Fig. A1, we see JAVELIN performs better on light curves without detrending. Detrending here appears to introduce a systematic bias towards smaller lags (most notably in panels 5–10). Again, this behaviour is reflected in the simulated data. JAVELIN is biased when additional sources of variability are present that are not described by the assumed linear response (one light curve is a shifted, smoothed, and scaled version of the other; equation 5). As we detrend each filter independently with a quadratic fit, we may be artificially introducing additional long-term variability outside of the model assumption, particularly in noisier light curves where the common long-term trends may be masked (we see a notable decrease in inter-band correlation when detrending; see Fig. 8).

Our findings suggest that each algorithm does not optimally perform on the same light curve set, with or without detrending. Ideally, we would compare these methods on the same set of curves as detrending removes some longer term variability that may be associated with reprocessing from the BLR (E. M. Cackett et al. 2022; C. Lewin et al. 2023) or intrinsic disc variability (P. Uttley et al. 2003; E. Breedt et al. 2009; J. V. Hernández Santisteban et al. 2020) that propagates oppositely to X-ray reprocessing lags (P. Z. Yao et al. 2023; J. M. M. Neustadt et al. 2024). We find that this mismatched comparison is necessary to compute reliable lags with each method. We also note that these recommendations are also verified on simulated data for which only the disc reprocessing scenario is present. Given this, and the overall consistency between detrending and not detrending with each algorithm, we do not see strong evidence that choosing one or the other significantly biases the presence of other delay phenomena.

Previously, we discussed (see Section 4.3) how selecting individual inter-band delays based on significance cuts, inter-band correlation, and fractional variability can bias the population distribution of lags. Thus, to compare our stacked lags to the individual object measurements, we compute running median (and 16th and 84th percentile) trends based on the complete, individual lag population. Individual delays are measured using the optimal detrending choice for each algorithm in the stacking scenario. Fig. A1 shows that the joint distribution of individual JAVELIN delays is remarkably similar to the stacked approach, but with much larger uncertainty. While the distribution of individual ICCF delays broadly matches those of the stacked approach under the uncertainties, the individual estimates do seem to favour shorter delays. It is difficult to verify whether stacking ICCF is smoothing out short-time-scale structure in the CCF and lengthening delays without also testing a finer time step resolution in the CCF, but given the consistency with the JAVELIN lags, we find it more likely that the individual ICCF estimates are biased.

Given our success with inferring joint delays, we test whether delays can be further improved by renormalizing light curves to

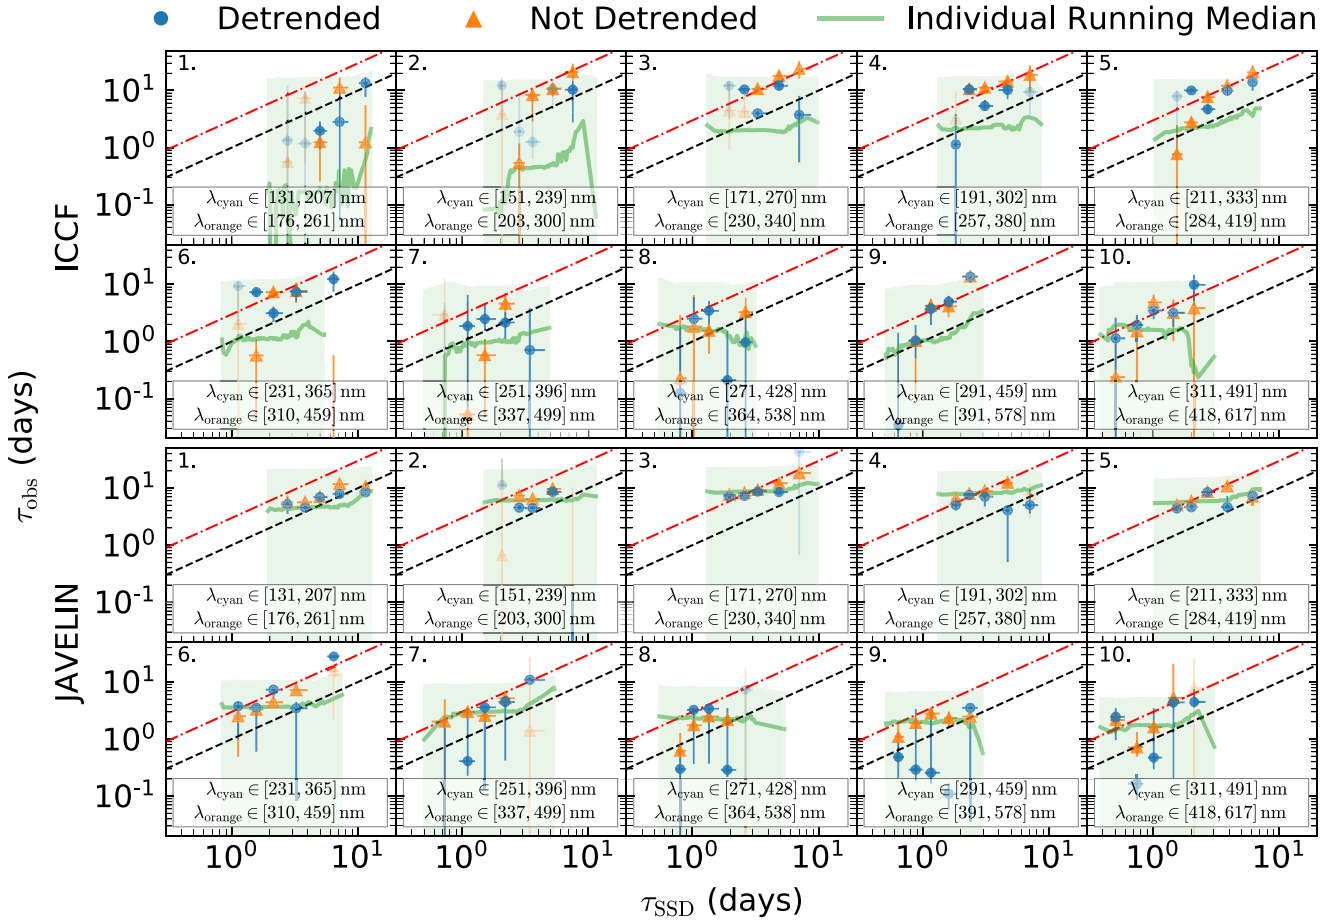


Figure A1. Comparison between the detrended and not detrended stacked delay approaches for ICCF (top) and JAVELIN (bottom) for the main sample. A running median (and 16th and 84th percentiles) derived from applying each algorithm to the individual light curve pairs is also displayed. Delays with multiple peaks in their posterior distributions have reduced opacity.

a common scale (by dividing by their maximum value). Renormalization does not affect ICCF as light curves are renormalized in calculating the Pearson correlation coefficient. For JAVELIN, we observe this renormalization has an adverse effect on the recovered delays in both the real and simulated data. This effect is more notable in light curves without detrending as we may artificially inflate the importance of light curves that do not show strong long-term variability, with the shorter time-scale fluctuations being closer to the order of the photometric noise, and hence diluting the signal.

APPENDIX B: LIGHT CURVE PROPERTIES AND THEIR INFLUENCE ON JOINT DELAYS

During our analysis, we identified that jointly fit delays were poorly constrained in some bins. These bins are identified by either possessing much larger uncertainties than their peers and/or multiple peaks in the posterior delay distributions. We compile all the lags measured during this work on real data finding approximately 27 per cent of ICCF and 11 per cent of JAVELIN delays possess more than one posterior peak (as defined in Section 4.4). The < 9 per cent of bins with delay semi-amplitude uncertainties ($\Delta\tau_{84-16}/2$) larger than 10 d belong almost exclusively to this multi-peaked population.

We investigate which light properties influence whether measured delays are poorly constrained by fitting a random forest model to predict whether a delay distribution will have multiple peaks using properties such as the number of constituent light curve pairs, the level of long and short-time-scale inter-band correlation, and the median relative uncertainty, fractional variability, and (median/minimum) cadence in each passband. The models achieve Area Under the Curve (AUC) scores on a 25 per cent (of all delays) validation set of 0.84 and 0.90 for ICCF and JAVELIN, respectively. We compute permutation feature importances for each model, identifying important predictors in long-time-scale inter-band correlation and median relative photometric error for JAVELIN, and short-time-scale inter-band correlation for ICCF. We see visually that the distributions of these predictors between unimodal and multimodal delay populations differ and confirm with a Kolmogorov–Smirnov test. For JAVELIN, multi-peaked behaviour occurs in regions with high inter-band correlation and low median relative photometric error. While these regions are where we would expect to best constrain lags, the additional important requirement is that the bin has a small number of objects ($\lesssim 20$). We compare the composite distribution of individual JAVELIN runs to the joint fit for several offending bins finding that JAVELIN either begins to resolve the individual delay signals or struggles to converge when the individual signals are confidently disparate and few in number. Whether this dis-

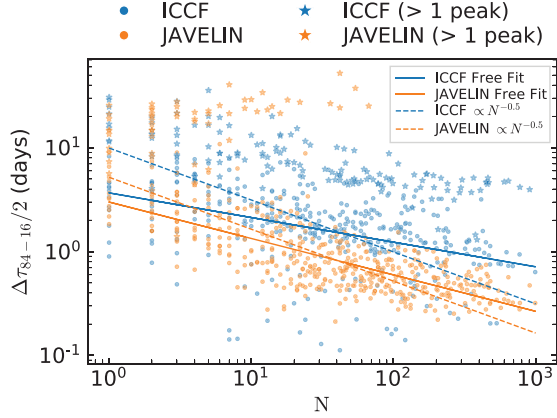


Figure B1. Semi-amplitude (uncertainty) of measured stacked delays against the number of contributing light curve pairs in each bin using both ICCF (blue) and JAVELIN (orange). Power laws with freely varying amplitude and slope are fit to each algorithm (solid lines), along with a fixed $N^{-1/2}$ relation intersecting the median error (dashed lines). Bins where posterior lag distributions showed more than one peak are indicated by stars rather than circles.

parade behaviour spuriously arises from our random sampling of an unfavourable variability episode or some uncontrolled binning parameter is unclear. At this stage, there does not appear to be any systematic issues with JAVELIN that are not explained by poor sample statistics. For ICCF, multimodal posteriors occur regardless of the number of constituent objects, primarily when short-time-scale inter-band correlation is low and spurious peaks in the CCF become more common. As such, ICCF does not perform equally well across our parameter space, disfavouring faint and redder objects and reaffirming our choice in JAVELIN as the preferred lag algorithm.

The remainder of the quoted stacked delay uncertainties outside multimodal bins can be largely explained by the number of constituent light curve pairs contributing to the estimate. Fig. B1 shows both ICCF and JAVELIN posterior delay semi-amplitudes




Figure B2. Comparison of the semi-amplitudes measured by the stacked approach to those derived from the expected SSD delay distribution in each bin. The black dashed line represents the one-to-one relation. Similarly, bins where posterior lag distributions showed more than one peak are indicated by stars rather than circles in blue (orange) for ICCF (JAVELIN).

scale more shallowly than expected for a statistical noise-limited system ($\propto 1/\sqrt{N}$, where N is the number of objects), although a visual comparison maintains $\propto 1/\sqrt{N}$ is still somewhat consistent. The displayed fits do not include the multimodal measurements (shown by stars) which are seen to have systematically larger uncertainties as expected. While the quoted uncertainties do improve with increased number statistics, they appear bounded by the inherent dispersion in the predicted SSD delay distribution (Fig. B2), suggesting delay errors are not significantly underestimated. Although we do note as per our discussion in Section 5.1, we are likely limited by systematic uncertainties regarding our bin compositions.

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