

Accretion disc time lag distributions: applying CREAM to simulated AGN light curves

D. A. Starkey,¹* Keith Horne¹ and C. Villforth^{1,2}

¹*SUPA School of Physics and Astronomy, St Andrews, KY16 9SS, Scotland, UK*

²*Department of Physics, University of Bath, Claverton Down, Bath BA2 7AY, UK*

Accepted 2015 November 18. Received 2015 November 12; in original form 2015 September 17

ABSTRACT

Active galactic nuclei (AGN) vary in their brightness across all wavelengths. Moreover, longer wavelength ultraviolet–optical continuum light curves appear to be delayed with respect to shorter wavelength light curves. A simple way to model these delays is by assuming thermal reprocessing of a variable point source (a lamp post) by a blackbody accretion disc. We introduce a new method, CREAM (Continuum REprocessed AGN Markov Chain Monte Carlo), that models continuum variations using this lamp post model. The disc light curves lag the lamp post emission with a time delay distribution sensitive to the disc temperature–radius profile and inclination. We test CREAM’s ability to recover both inclination and product of black hole mass and accretion rate $M\dot{M}$, and show that the code is also able to infer the shape of the driving light curve. CREAM is applied to synthetic light curves expected from 1000 s exposures of a 17th magnitude AGN with a 2-m telescope in Sloan g and i bands with Signal-to-Noise Ratio (SNR) of 500–900 depending on the filter and lunar phase. We also test CREAM on poorer quality g and i light curves with SNR = 100. We find in the high-SNR case that CREAM can recover the accretion disc inclination to within an uncertainty of 5° and an $M\dot{M}$ to within 0.04 dex.

Key words: accretion, accretion discs – methods: statistical – quasars: supermassive black holes.

1 INTRODUCTION

Active galactic nuclei (AGN) at the centres of distant galaxies are thought to contain an accretion disc orbiting a supermassive black hole (SMBH), clouds of both broad and narrow-line-emitting gas (the Broad Line Region and Narrow Line Region) that surround the AGN with uncertain geometry and kinematics, and a dusty torus that obscures the disc component when viewed close to edge on (Urry & Padovani 1995; Elvis 2000). The accretion disc emission is thought to give rise to the ‘Big Blue Bump’ feature observed in the spectra of AGN and quasars (Malkan 1983; Shang et al. 2005). The Broad Line Region (BLR) extents over 10–200 light days (Bentz et al. 2013) and we can infer from observed continuum–continuum light-curve time delays (Cackett, Horne & Winkler 2007; Edelson et al. 2015) and light-travel-time arguments that the accretion disc occupies a region several light days in size. AGN distances make resolving their structure impractical at this time. A number of studies (Kishimoto et al. 2011; Hönig et al. 2014) use infrared interferometry to resolve the inner edge of the dusty torus, but all AGN are too remote to resolve the BLR much less the accretion disc. To obtain information on the accretion disc, one can therefore proceed along one of two routes. The half light radii of the accretion disc at

some wavelength can be inferred from microlensing observations of gravitational lensed quasars (Morgan et al. 2012; MacLeod et al. 2015). Lensing produces multiple images of the quasar whose flux ratios are sensitive to the size of the emitting region at the specified wavelength. One can alternatively obtain information on the structure and kinematics of the BLR and accretion disc in the time rather than space domain via Reverberation Mapping (RM; Blandford & McKee 1982). RM assumes that changes in the accretion disc emission drive the BLR line variability and, by measuring time delays between the disc and broad line emission, one can infer the black hole mass and size of the BLR.

As well as continuum–line delays, one can perform RM studies between continuum light curves at different wavelengths and this is the subject of this paper. AGN and quasar continuum emission exhibits time variability that can be modelled by a damped random walk (DRW; Kelly, Bechtold & Siemiginowska 2009). A DRW is characterized by a stochastic differential equation whose power spectrum is a broken power law with a slope that steepens from 0 to -2 above a break time-scale. In AGN observed in the Sloan digital sky survey (SDSS), the break time-scale is upwards of 100 d, correlates positively with black hole mass and wavelength and is a weak function of luminosity (MacLeod et al. 2010).

Light curves at longer wavelengths appear to lag behind those at shorter wavelengths (Cackett et al. 2007; Chelouche 2013; Edelson et al. 2015) with evidence of a mean delay scaling with wavelength

* E-mail: ds207@st-andrews.ac.uk

like $\langle \tau \rangle \propto \lambda^{4/3}$. The inhomogeneous disc model suggests that variability arises from temperature fluctuations that occur stochastically throughout the accretion disc (Dexter & Agol 2011). While this can explain the magnitude and time-scales of variability observed by e.g. MacLeod et al. (2010), it does not easily explain the short-to-long-wavelength time lag. This feature emerges naturally if the disc is irradiated by hard photons close to the black hole, with time delays increasing towards longer wavelengths due to light-travel-time effects. It is not known whether the photons driving the variability are emitted at far-UV or harder X-ray wavelengths. *SWIFT* observations of NGC 5548 (McHardy et al. 2014) show evidence of X-rays leading UV but accurate models of the delay distribution function are crucial for understanding the origin of disc variability.

Cross-correlating short and long wavelength light curves is the traditional way to measure both continuum–line and continuum–continuum time delays. Two cross-correlation techniques to measure lags despite the uneven sampling (for example incurred from bad weather, blockage by the sun or observing time constraints) are the interpolated cross-correlation function (ICCF; Gaskell & Peterson 1987) and discrete cross-correlation function (DCCF; Edelson & Krolik 1988). ICCF works by linearly interpolating uneven data into evenly spaced data, whereas the DCCF discrete correlation function (DCF) bins the uneven data into equal time widths or centroid of the cross correlation function (CCF) gives an estimate of the mean time delay.

A more recent development, SPEAR (Zu, Kochanek & Peterson 2011) (and the PYTHON implementation JAVELIN) fits the DRW model to an input light curve that is assumed to drive the AGN variability. The driver is usually assumed to be the continuum light curve observed at the smallest wavelength. SPEAR then smooths and delays this light curve in an attempt to best model the continuum and line light curves. This method has been successfully applied to model both continuum–line delays (Zu et al. 2013) and continuum–continuum delays (Edelson et al. 2015). Time delays measured from CCF and SPEAR are in general agreement with one another but each of these methods returns just a single number for the time delay (e.g. the mean, CCF peak, or centroid of an assumed top hat delay function).

In this work, we present a Markov Chain Monte Carlo (MCMC) approach to fit a reverberating accretion disc model and test this using synthetic light curves. Our approach differs from SPEAR and CCF routines as it returns a distribution of continuum time delays that is a function of accretion disc parameters such as the inclination and product $M\dot{M}$ of black hole mass M and accretion rate \dot{M} . Furthermore, the MCMC code (CREAM) does not require an input light curve to act as the driver of variability, it instead infers the shape of the driving light curve.

The layout of the paper is as follows. Section 2 describes the assumed accretion disc model. In Section 3, we detail the statistical framework of the MCMC code. We describe how the synthetic data sets used to test the code are constructed in Section 4. Section 5 presents results of applying Continuum Reprocessing Echo AGN Markov Chain Monte Carlo (CREAM) to the synthetic light curves and shows the posterior probability distributions for $M\dot{M}$ and disc inclination. We present discussion and conclusions in Section 6.

2 ACCRETION DISC THERMAL REPROCESSING MODEL

We model continuum light-curve variability as temperature fluctuations due to variable irradiation of a flat, blackbody accretion disc.

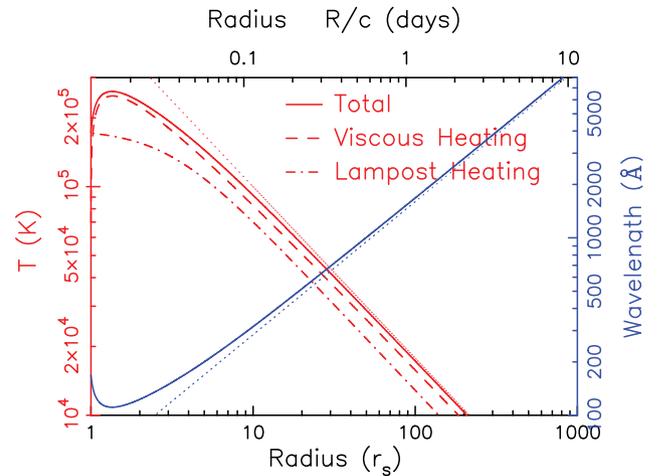


Figure 1. Temperature radius profile due to the viscous heating (dashed red line) and lamp post heating (dot-dashed red line). The total effect is shown in solid red and the dotted red line shows a power-law slope of $-3/4$. The red and blue lines are described by the y-axis on the left- and right-hand sides of the plot, respectively. The solid blue line shows the peak wavelength in the blackbody spectrum $B_v(\lambda_{\text{peak}})$ as a function of radius assuming a temperature radius law given by equation (2). For a flat disc $\lambda_{\text{peak}} \propto r^{4/3}$. We plot a $4/3$ line in dotted blue for comparison. The plot is appropriate for a SMBH of $10^8 M_{\odot}$ accreting at $1 M_{\odot} \text{ yr}^{-1}$.

Variability arises due to accretion disc reprocessing of photons from a variable point source located just above the black hole. In this section, we describe the underlying physics of the model and introduce the time delay distribution appropriate to our accretion disc model.

2.1 Steady state temperature radius structure

The thermal reprocessing model assumes that each point in a flat, optically thick accretion disc has an associated blackbody temperature T and emits at wavelength λ according to the Planck function,

$$B_v(\lambda, T) = \frac{2hc}{\lambda^3} \frac{1}{e^{hc/\lambda kT} - 1}, \quad (1)$$

where h and k are the Planck and Boltzmann constants, and c is the speed of light. We model the source of irradiation as a ‘lamp post’ with bolometric luminosity $L_b(t)$, located a height h_x above the black hole. The lamp post approximates the region driving the variability as a point source from which disc variability at radius r and azimuth ϕ is delayed by a single light-travel-time induced lag. To compute the disc flux at a given wavelength, we require a temperature–radius relation. The $T(r)$ profile (Fig. 1) combines the effects of viscous heating due to differential rotation and irradiation by the lamp post and can be expressed as (Frank, King & Raine 2002; Cackett et al. 2007)

$$T^4 = \frac{3GM\dot{M}}{8\pi\sigma r^3} \left(1 - \sqrt{\frac{r_{\text{in}}}{r}}\right) + \frac{L_b(1-a)h_x}{4\pi\sigma x^3}, \quad (2)$$

where M is the black hole mass, \dot{M} the accretion rate, a the disc albedo, r the orbital radius of the surface element, h_x the lamp post height, $x = \sqrt{r^2 + h_x^2}$ the distance from the lamp post point to surface element and r_{in} is the inner most stable circular orbit ($3 r_s$ for a Schwarzschild black hole). In the limit $r \gg r_{\text{in}}$, h_x , equation (2) simplifies to

$$T = T_0 \left(\frac{r_0}{r}\right)^{3/4}, \quad (3)$$

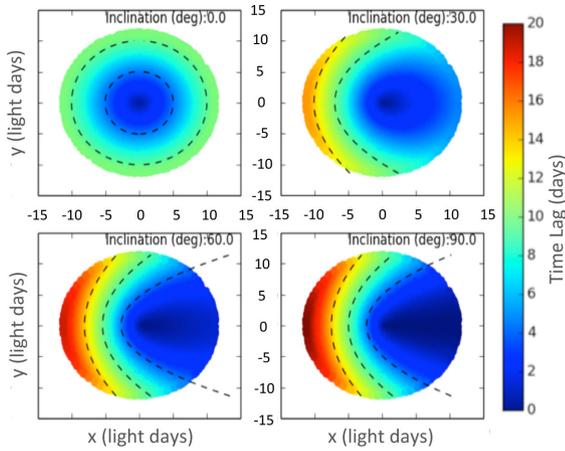


Figure 2. Accretion disc time delay profiles $\tau(r, \phi, i)$ evaluated at 0° (top left), 30° (top right), 60° (bottom left) and 90° (bottom right). Dashed lines represent 5, 10 and 15 d isodelay surfaces.

where

$$T_0^4 = \frac{3GM\dot{M}}{8\pi\sigma r_0^3} + \frac{h_x(1-a)L_b}{4\pi\sigma r_0^3}. \quad (4)$$

The lamp post’s bolometric luminosity is $L_b = \eta\dot{M}c^2$ with an efficiency parameter η typically around 0.1 (Shankar, Weinberg & Miralda-Escudé 2009).

The X-ray emitting region has been constrained by microlensing techniques to within just $6r_s$ of the SMBH (Morgan et al. 2012) and McHardy et al. (2015) find evidence of X-ray variability leading that at UV wavelengths. The point source lamp post approximation can therefore be justified with reference to Fig. 1. The shortest wavelength light curve we use in this work is in the Sloan u filter with central wavelength 3545 Å. For a fiducial case black hole with $M = 10^8 M_\odot$ and $\dot{M} = 1 M_\odot \text{ yr}^{-1}$ the radius of peak emission at this wavelength is $\sim 300r_s$ (Fig. 1). The $6r_s$ X-ray emitting region effectively appears as a point source at this radius.

2.2 Disc time delays

Photons emitted isotropically by the lamp post intercept the disc surface. Once absorbed, the lamp post photons are assumed to heat the disc instantaneously. This is known as an instantaneous thermal reprocessing model (Cackett et al. 2007) and causes the disc at the location of impact to radiate as a blackbody with a higher temperature. The reprocessed photons are seen by a distant observer with a light-travel-time-induced lag, $\tau(r, \phi, i)$, dependent on the radius, r , and azimuth, ϕ , of reprocessing site and inclination, i , of the disc with respect to the observer ($i = 0$ corresponds to a face on disc). For a flat disc,

$$\tau(r, \phi) = \frac{r}{c} (1 + \cos\phi \sin i). \quad (5)$$

Surfaces of constant time delay form circles for a face on disc ($i = 0^\circ$), parabolas for an edge on disc and ellipses for an intermediately inclined disc (Fig. 2).

2.3 Disc response function

Fig. 2 shows that the continuum reprocessing occurs at reprocessing sites over a range of radii r and azimuths ϕ corresponding to a distribution of time lags. To obtain the total disc flux $F(\lambda, t)$ we convolve

the variable component of the driving light-curve flux $\Delta F_x(t)$ with a time delay distribution $\psi(\tau|\lambda)$ and add on a background spectrum $\bar{F}_v(\lambda)$,

$$F_v(\lambda, t) = \bar{F}_v(\lambda) + \Delta F_v(\lambda) \sum_{i=0}^{\tau_{\max}} \psi(\tau_i|\lambda) \Delta F_x(t - \tau_i) \Delta\tau. \quad (6)$$

Here, $\psi(\tau|\lambda)$ is dimensionless and normalized such that $\sum \psi(\tau|\lambda) \delta\tau = 1$ and $\Delta F_v(\lambda)$ is the variable component of the spectrum. The response function $\psi(\tau|\lambda)$ for a flat disc is derived in Cackett et al. (2007) and is a function of $M\dot{M}$ and inclination. We plot several example response functions showing the $M\dot{M}$, inclination and wavelength dependences in Fig. 3. The mean delay, as shown by the vertical lines in Fig. 3, is independent of inclination. Tilting the disc makes the response function more skewed; it peaks at shorter lags and develops a tail towards large lags. This effect arises because higher inclinations increase the lag on the far side and decrease the lag of the near side of the disc relative to a face on inclination (see Fig. 2). If we increase $M\dot{M}$, reprocessing sites emitting at a fixed temperature move to larger radii. This scales the response function without altering its shape with mean delays scaling like $\langle\tau\rangle \propto (M\dot{M})^{1/3} \lambda^{4/3}$. Both the mean delay and response function width increase with wavelength. These effects occur due to the reduction in temperature with radius (equation 3). We see from Fig. 1 that the assumed blackbody-emitting disc allows the longer wavelength emission to extend to larger radii with higher time delays.

3 CREAM: CONTINUUM REPROCESSING ECHO AGN MARKOV CHAIN MONTE CARLO CODE

With the effects of disc inclination and black hole $M\dot{M}$ in mind, we fit a multiparameter model to a set of light curves. Our Continuum Reprocessing Echo AGN Markov Chain Monte Carlo code (CREAM) is detailed in these subsections. CREAM ingests light-curve data at N_λ wavelengths and uses a Metropolis-within-Gibbs parameter sampling approach to sample the posterior parameter distributions and find the best-fitting parameter values and uncertainties (for both $M\dot{M}$, inclination and the parameters in the driving light curve). CREAM fits equation (6) to an input set of continuum light curves. The mean spectrum $\bar{F}_v(\lambda)$ models all constant contributions to these echo light curves. This primarily addresses host galaxy contamination but considers also any other slowly varying or invariant emission sources. The $\Delta F_v(\lambda)$ parameters scale the echo light-curve variations to match the echo light curve data.

3.1 Driving Light curve: $\Delta F_x(t)$

Many CCF studies of continuum light curves assume the shortest wavelength continuum light curve acts as the driver of variability. Lags of optical to X-ray light curves have been observed in several (Arévalo et al. 2008; McHardy et al. 2015) objects but in other instances X-rays appear only weakly to relate to UV and optical variability (Edelson et al. 2015). It is therefore not fully understood what wavelength regime acts as the variability-driving lamp post. CREAM has the ability to infer the shape of the true driving light curve from the shape of the echo light curves alone. In order to achieve such a realization of the driving light curve, we need parameters that fully specify the shape of the driving light curve and an ability to constrain these parameters so that the observed power spectrum

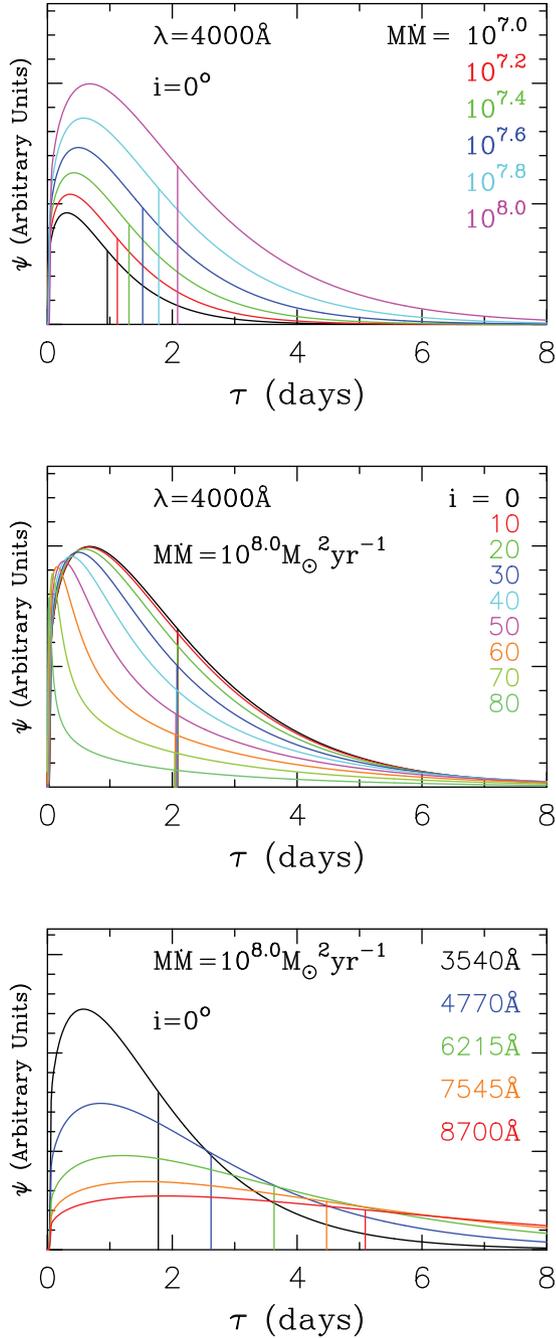


Figure 3. The numerically evaluated response functions $\psi(\tau|\lambda)$ for a face-on, flat disc around a $10^8 M_\odot$ black hole accreting at $1 M_\odot \text{ yr}^{-1}$. The lamp post is set at $h_x = 3r_s$ above the disc plane. We then show the effect of varying $M\dot{M}$ (top), inclination (middle) and wavelength λ (bottom). Vertical lines show the mean delay $\langle\tau\rangle$.

is consistent with observations. We represent the driving light curve by a Fourier time series,

$$\Delta F_x \equiv F_x(t) - \bar{F}_x = \sum_{k=1}^{N_k} C_k \cos(\omega_k t) + S_k \sin(\omega_k t), \quad (7)$$

where $\omega_k = k\Delta\omega$ is the k th Fourier frequency, S_k and C_k are the Fourier amplitude coefficients. \bar{F}_x is the mean of the driving light curve. The reference level \bar{F}_x is therefore somewhat arbitrary and we use $\bar{F}_x = 0$ for our tests. If however we choose to include driving

light-curve data, we set the Fourier sum in equation (7) to correspond to the log of the driving light-curve variable component. This enforces positivity in $\Delta F_x(t)$ that then oscillates about 1. The reference level \bar{F}_x is then included as a multiplicative parameter. The low and high frequencies (ω_{low} and ω_{hi}) in the Fourier sum (equation 7) can be specified as the code is initialised. For these tests we use,

$$\omega_{\text{low}} = \Delta\omega = \frac{1}{2} \frac{2\pi}{T_{\text{rec}}}, \quad (8)$$

$$\omega_{\text{hi}} = \frac{2\pi}{\Delta t} = N_k \Delta\omega. \quad (9)$$

Here, T_{rec} , the recurrence time of the Fourier series, should be longer than the timespan of the data plus the width of the response function (Fig. 3). Δt is the mean time separation between adjacent points. For the specific tests in this work, our light curves span 100 d and we choose $T_{\text{rec}} = 200$ d. We discuss the choice of upper frequency in Section 6 but the default value specified in equation (9) satisfies the Nyquist sampling requirement of fewer than two data points within 1 period of the high-frequency component.

3.2 The Badness of Fit

Adopting a Bayesian framework, we begin this discussion with Bayes' theorem, relating the posterior probability distribution, $P(\Theta|\mathbf{D})$, of N_{par} parameters, $\Theta = (\theta_1, \dots, \theta_{N_{\text{par}}})$, with priors $P(\Theta)$, fitted to data, $\mathbf{D} = (D_1, \dots, D_N)$,

$$P(\Theta|\mathbf{D}) = P(\mathbf{D}|\Theta) \frac{P(\Theta)}{P(\mathbf{D})}. \quad (10)$$

Here, $P(\mathbf{D}|\Theta)$ is the 'likelihood function', $P(\mathbf{D})$ normalizes the posterior probability

$$\int P(\Theta|\mathbf{D}) d\Theta = 1 \quad (11)$$

The 'Badness of Fit' (BOF) is defined as $-2 \ln(P(\Theta|\mathbf{D}))$ and takes the form

$$\text{BOF} = \chi^2 + \sum_{i=1}^N \ln(\sigma_i^2) - 2 \ln(P(\Theta)) + \text{const}, \quad (12)$$

where

$$\chi^2 = \sum_{i=1}^N \left(\frac{D_i - M_i}{\sigma_i} \right)^2. \quad (13)$$

The model $\mathbf{M} = (M_1 \dots M_N)$ is a function of the parameters Θ , and we assume Gaussian errors on the N light-curve data points with error bars σ_i .

3.3 Priors $P(\Theta)$

The model parameters and their priors are summarized in Table 1. Counting the model parameters, we have $2N_k$ sine and cosine amplitudes specifying the shape of the driving light curve about a mean level \bar{F}_x . We also have a response function parametrized by i and $M\dot{M}$ that smooths and delays the driving light curve. The mean flux $\bar{F}_v(\lambda)$ and scaling terms $\Delta F_v(\lambda)$ map the smoothed-delayed driving light curve on to the reprocessed-echo light-curve data. In total, $N_{\text{par}} = 2(N_k + N_\lambda + 1) + 1$.

The inclination prior is uniform in $\cos i$ reflecting the assumption that the disc is randomly orientated in the sky. The Fourier amplitudes, S_k and C_k , are constrained with Gaussian priors as described

Table 1. Summary of priors on the CREAM parameters (σ_k is defined in equation (16)). The table shows (left column) the parameter, (middle column) the number of relevant parameters and (right column) the prior on the parameter(s).

Parameter	N_{par}	Prior
S_k and C_k	$2N_k$	Gaussian ($\langle S_k \rangle = \langle C_k \rangle = 0$, $\langle S_k^2 \rangle = \langle C_k^2 \rangle = \sigma_k^2$)
$\cos i$	1	Uniform
$\log M\dot{M}$	1	Uniform
$\log \Delta F_\nu$	N_λ	Uniform
$\log \bar{F}_\nu(\lambda)$	N_λ	Uniform

in equations (14) to (17). For necessarily positive parameters $M\dot{M}$, $\Delta F_\nu(\lambda)$ and $\bar{F}_\nu(\lambda)$ our priors are uniform in log.

We now discuss our choice of the prior used to constrain the driving light curve. With uniform priors on the Fourier amplitudes the driving light curve over-fits the data by over utilizing the high-frequency Fourier amplitudes. This effect is demonstrated in Fig. 4 and motivates the choice of a more suitable prior to constrain the shape of the driving light curve. McHardy et al. (2014) suggest that X-ray light curves may drive AGN variability and we choose to constrain the S_k and C_k parameters using a ‘random walk prior’. This prior drives the driving light curve’s S_k and C_k Fourier parameters

(equation 7) towards values giving rise to a random walk power spectrum,

$$\langle S_k^2 \rangle + \langle C_k^2 \rangle = P(\omega_k)\Delta\omega = P_0\Delta\omega \left(\frac{\omega_0}{\omega_k}\right)^\alpha, \quad (14)$$

with $\alpha = 2$ motivated by studies of X-ray variability (Uttley, McHardy & Papadakis 2002; McHardy et al. 2006). We incorporate the power spectrum prior in equation (14) as a Gaussian prior on the S_k and C_k parameters with mean $\langle S_k \rangle = \langle C_k \rangle = 0$, and variance $\langle S_k^2 \rangle = \langle C_k^2 \rangle = \sigma_k^2$,

$$P(\Theta) = \prod_{k=1}^{N_k} \frac{e^{-\frac{1}{2} \frac{C_k^2 + S_k^2}{\sigma_k^2}}}{2\pi\sigma_k^2}, \quad (15)$$

where,

$$\sigma_k^2 = \frac{P_0\Delta\omega}{2} \left(\frac{\omega_0}{\omega_k}\right)^2. \quad (16)$$

The BOF, substituting equations 15 in (12), is then

$$\begin{aligned} \text{BOF} = \chi^2 + \sum_{i=1}^N \ln(\sigma_i^2) \\ + \sum_{k=1}^{N_k} \left(2 \ln(\sigma_k^2) + \frac{C_k^2 + S_k^2}{\sigma_k^2} \right) + \text{const.} \end{aligned} \quad (17)$$

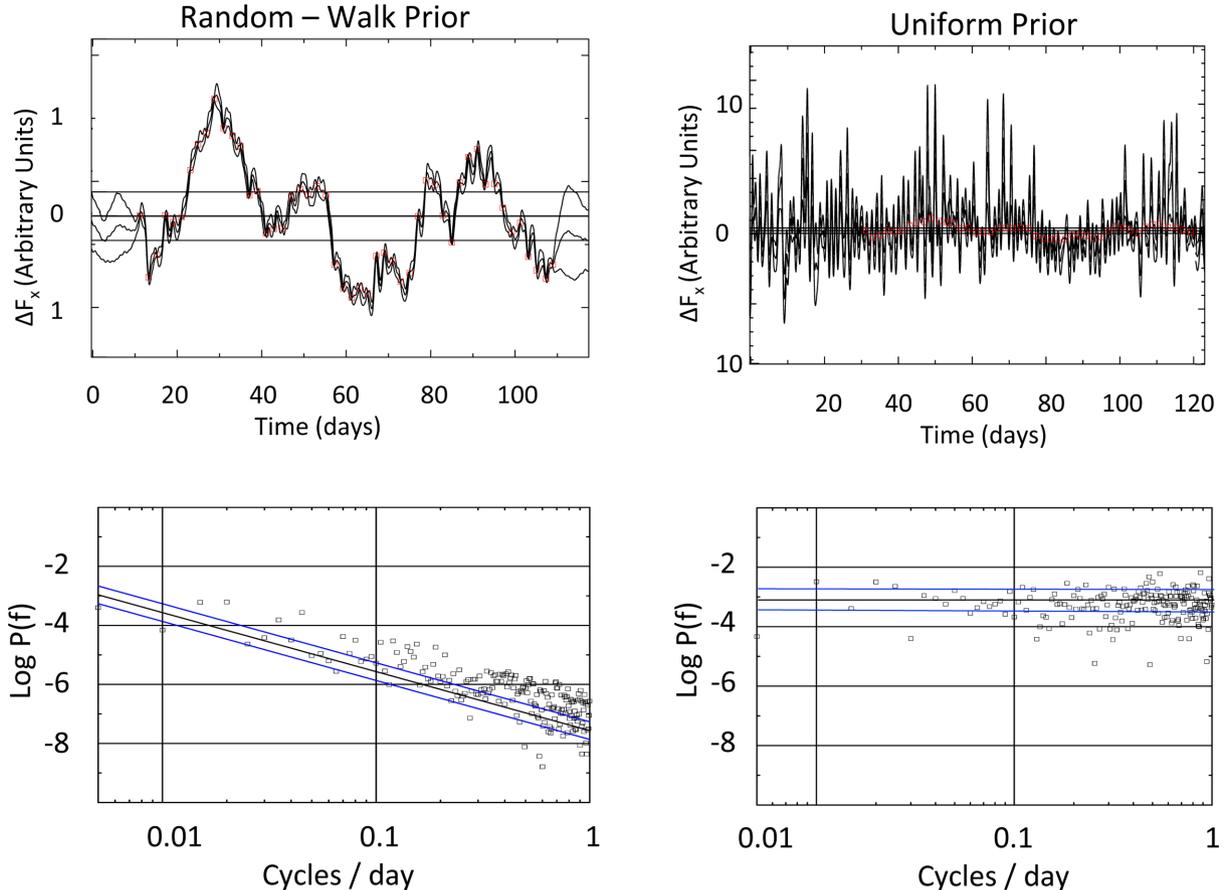


Figure 4. Left: CREAM fit to synthetic light curve with data points in red and mean $\pm 1\sigma$ envelopes in black. Horizontal lines indicate the reference position with 1σ envelopes. The bottom panel shows the power spectrum of the resulting fit with the black squares showing $S_k^2 + C_k^2$ for each Fourier frequency. The black and blue lines indicate the random walk prior (equation 14). On the right we show the effect of fitting the same data with a uniform prior.

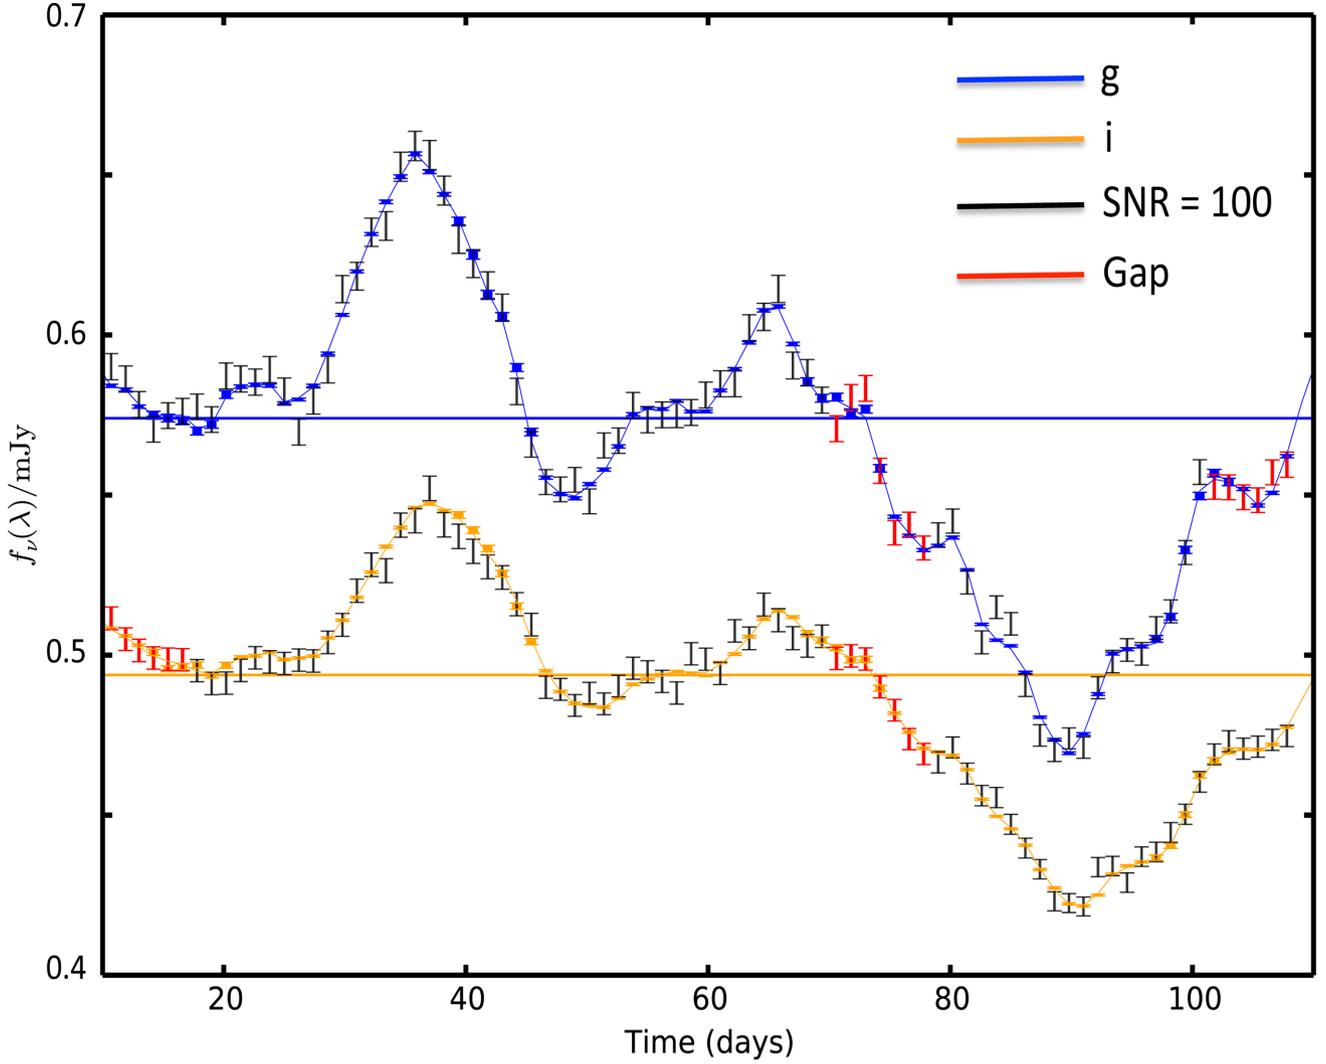


Figure 5. The synthetic light curves used to test CREAM. Noise-free model light curves at g and i wavelengths are shown by the wavy blue and orange lines. The stars indicate a high-quality light curve with mean SNR = 713 and 690, respectively but these vary depending on the lunar phase. We also test CREAM with SNR = 100 light curves in g and i shown by the black error bars. Red light bars are omitted from the SNR = 100 light curves to simulate week long data gaps as discussed in Section 5.2. The analysis of the high-quality light curves, the SNR=100 light curve and the SNR = 100 light curves omitting the red points is discussed in Section 5.

This ‘random walk prior’ gives a more satisfactory fit, as shown in Fig. 4.

4 SYNTHETIC LIGHT CURVES

RM observing campaigns require high-cadence, long-duration light curves. We choose to test CREAM on synthetic light curves (Fig. 5) generated assuming the observing capabilities of the Las Cumbres Observatory Global Telescope (LCOGT) network in Sloan g and i filters with 1000 s exposures. The global longitude coverage of the LCOGT network allows targets to be observed without daily obstruction by the Earth and the g and i filters offer both high signal to noise and broad wavelength coverage (4775–7540 Å).

We first generate a uniform time grid $t_i = i\Delta t$ on which to evaluate the model light curves. This is convenient for evaluating the convolution integral (equation 6). We use $\Delta t = 0.1$ d and gen-

erate the driving light curve using a random walk about a mean $\bar{F}_x = 0$,

$$\Delta F_x(t_i) \equiv F_x(t_i) - \bar{F}_x = \Delta F_x(t_{i-1}) + \sum_{k=1}^i G(0, \sigma\sqrt{\Delta t}). \quad (18)$$

Here, $G(a, b)$ represents a random sample from a Gaussian distribution with mean a and standard deviation b . We adopt $\sigma = 1$ for the standard deviation of the random walk after 1 d.

The echo light curves are then obtained from the driver by convolution with the response function (Fig. 3) evaluated at inclination 30° and $M\dot{M} = 10^8 M_\odot^2 \text{yr}^{-1}$. These are the true parameters for our simulated light curves. The mean level of the g light curve is set at AB magnitude 17 and the mean level of the i light curve scales relative to this reference level like $F_v(\lambda) \propto \lambda^{-1/3}$, as expected for a flat blackbody-emitting accretion disc. This brightness is appropriate for a target observed at $z = 0.27$, $\Omega_\Lambda = 0.7$, $\Omega_m = 0.3$, $H_0 = 70 \text{ km s}^{-1} \text{Mpc}^{-1}$, with an accretion rate, $\dot{M} = 1 M_\odot \text{yr}^{-1}$ and a black hole mass $M = 10^8 M_\odot$, fairly typical for an AGN.

Table 2. Mean and rms of the synthetic *ugriz* light curves. The rms is a function of the DRW time-scale τ_B and SF_∞ parameters measured by MacLeod et al. (2010). These are a function of z , M , \dot{M} and λ and we quote the relevant values below.

Filter	τ_B (days)	SF_∞ (mag)	(AB) (mag)	rms (mag)
<i>u</i> 3540 Å	143	0.138	16.89	0.098
<i>g</i> 4770 Å	151	0.121	17.00	0.085
<i>r</i> 6215 Å	159	0.107	17.10	0.076
<i>i</i> 7545 Å	165	0.098	17.17	0.069
<i>z</i> 8700 Å	170	0.092	17.22	0.065

Table 3. The SNR's obtained from the LCOGT exposure time calculator for 1000 s observations on the *ugriz* filters of the 2-m spectral cameras at airmass 1.3. The SNR is evaluated for three lunar phases. We show also the exposure time required to achieve SNR = 100.

Filter	SNR for $t_{\text{exp}} = 10^3$ s			t_{exp} for SNR = 100		
	Dark	Half	Full	Dark	Half	Full
<i>u</i>	471	435	234	160	174	447
<i>g</i>	999	806	333	57	64	189
<i>r</i>	967	889	443	42	46	142
<i>i</i>	930	776	364	59	68	189
<i>z</i>	674	569	257	139	173	719

MacLeod et al. (2010) models AGN continuum light curves as DRWs with a characteristic break time-scale τ_B . We use the empirical relation between light-curve variance, z , \dot{M} , M and λ (equation 7 of MacLeod et al. 2010) to determine the appropriate variability amplitude for our light curves. This equation returns a parameter SF_∞ that is related to the light-curve variance by

$$\sigma^2(F_v) = \frac{SF_\infty^2}{2} (1 - e^{-t_{\text{len}}/\tau_B}), \quad (19)$$

where t_{len} is the length of the light curve. τ_B is again empirically constrained by MacLeod et al. (2010) in terms of the parameters λ , z , \dot{M} and M and we quote in Table 2 the values of τ_B and SF_∞ used to construct our synthetic light curves.

4.1 Observational uncertainties

We mimic irregular sampling over the 100 d observing period by linearly interpolating the 0.1 d cadence echo light curves on to a new light curve sampled at times $t_i = 1, \dots, N$ according to

$$t_i = t_{i-1} + G(0, \overline{\Delta t}), \quad (20)$$

where $\overline{\Delta t}$ is the mean time separation between adjacent points (for our light curves, we use $\overline{\Delta t} = 1$). We also require $t_i > t_{i-1}$.

The error bars, generated using the LCOGT exposure time calculator,¹ are appropriate for observations of a 17th magnitude target observed in *ugriz* at airmass 1.3 with 1000 s exposures. We compare this with a first order signal to noise calculation assuming just Poisson noise and a circular sky aperture with inner and outer radius of 10 and 20 arcsec. Both methods produced similar Signal-to-Noise Ratio (SNR) values, and we chose to adopt those from the LCOGT calculator. These SNR values shown in Table 3 are strong functions of the phase of the moon. We incorporate this effect into

the synthetic light curves by modulating the error bars with a lunar cycle. Gaussian noise is added to the data in the usual way,

$$F_v(\lambda, t) = F_{0v}(\lambda, t) G\left(1, \frac{1}{\text{SNR}}\right), \quad (21)$$

where the SNR value chosen is appropriate for the lunar phase in each week long segment, F_{0v} are the perfect data.

The final light curves in Fig. 5 show significant variability as expected from the observations of MacLeod et al. (2010). Delays of the *i* with respect to the *g* light curve are clearly present and best seen by eye when comparing positions of individual peaks and troughs in the blue and orange lines.

5 CREAM FITTING RESULTS

CREAM is applied to the *g* and *i* light curves plotted in Fig. 5. We launch three independent chains with random initial values of i between 0 and 90° and $M\dot{M}$ between 10^7 and $10^9 M_\odot \text{yr}^{-1}$. The Fourier S_k and C_k amplitudes start at 0. The offset and scaling terms for the echo light curve ($F_v(\lambda)$ and $\Delta F_v(\lambda)$) are started at the mean and rms of each light curve. CREAM then cycles through the $2(N_k + N_\lambda + 2) + 1$ parameters and explores the parameter space subject to the priors discussed in Section 3.3, stopping after 10^5 iterations.

CREAM fits to the various synthetic light curves are shown in Figs 6 to 9. The inferred driving light curve $\Delta F_x(t)$ is shown in Panel (a) and is convolved with the response functions $\psi(\tau|\lambda)$ in Panels (b) and (d) to yield the model *g* and *i* light curves, respectively. The response functions are parametrized by $M\dot{M}$ and inclination. Using inclination and $M\dot{M}$, CREAM infers a delay distribution function from which we calculate the mean delays for the *g* and *i* light curves. Mean delays are proportional to $(M\dot{M})^{1/3} \lambda^{4/3}$ and are annotated with uncertainties in Figs 6 to 9. Mean delays are 2.62 ± 0.05 d and 4.78 ± 0.08 d, respectively, for the *g* and *i* light curves. It can be seen that CREAM achieves a superb fit to the data with tightly constrained response functions. The posterior probability distributions for the fits in Figs 6 to 9 are shown in Figs 10 to 13. We also see from the red line in Panel (a) of Figs 6 to 9 that CREAM is able to correctly infer the shape of the driving light curve with no data present. We see that the inferred model begins to deviate from the true driving light curve at times earlier than around 10–15 d before the first data point in *g* and *i*. The response functions for *g* and *i* (Panels (b) and (d) in Fig. 6 are approximately 0 after a ‘look back’ time of 15 d. The model at the time of the first data point is therefore insensitive to the driving light-curve behaviour earlier than 15 d before the start of the observations.

Fig. 10 shows the posterior distribution of the parameters $M\dot{M}$ and i , corresponding to the fit in Fig. 6. Colours represent individual chains. We note from the figure and the results in Table 4 that CREAM is able to estimate inclination to $\pm 5^\circ$ and $M\dot{M}$ to ± 0.04 dex. We regard this as a benchmark best case scenario and now explore the extent to which these results degrade with noisier observations.

5.1 Less favourable observations

Because systematic errors often limit CCD photometry, we consider CREAM's ability to recover $M\dot{M}$ and inclination from light curves with a poorer SNR. We use the same driving light curve as with the previous test and create the *g* and *i* response light curves using the prescription described in Section 4. For this new set of *g* and *i* light curves, we generate error bars and perturb the data points assuming an SNR of 100. We show in Table 3 the exposure time needed to

¹ http://lcogt.net/files/etc/exposure_time_calculator.html

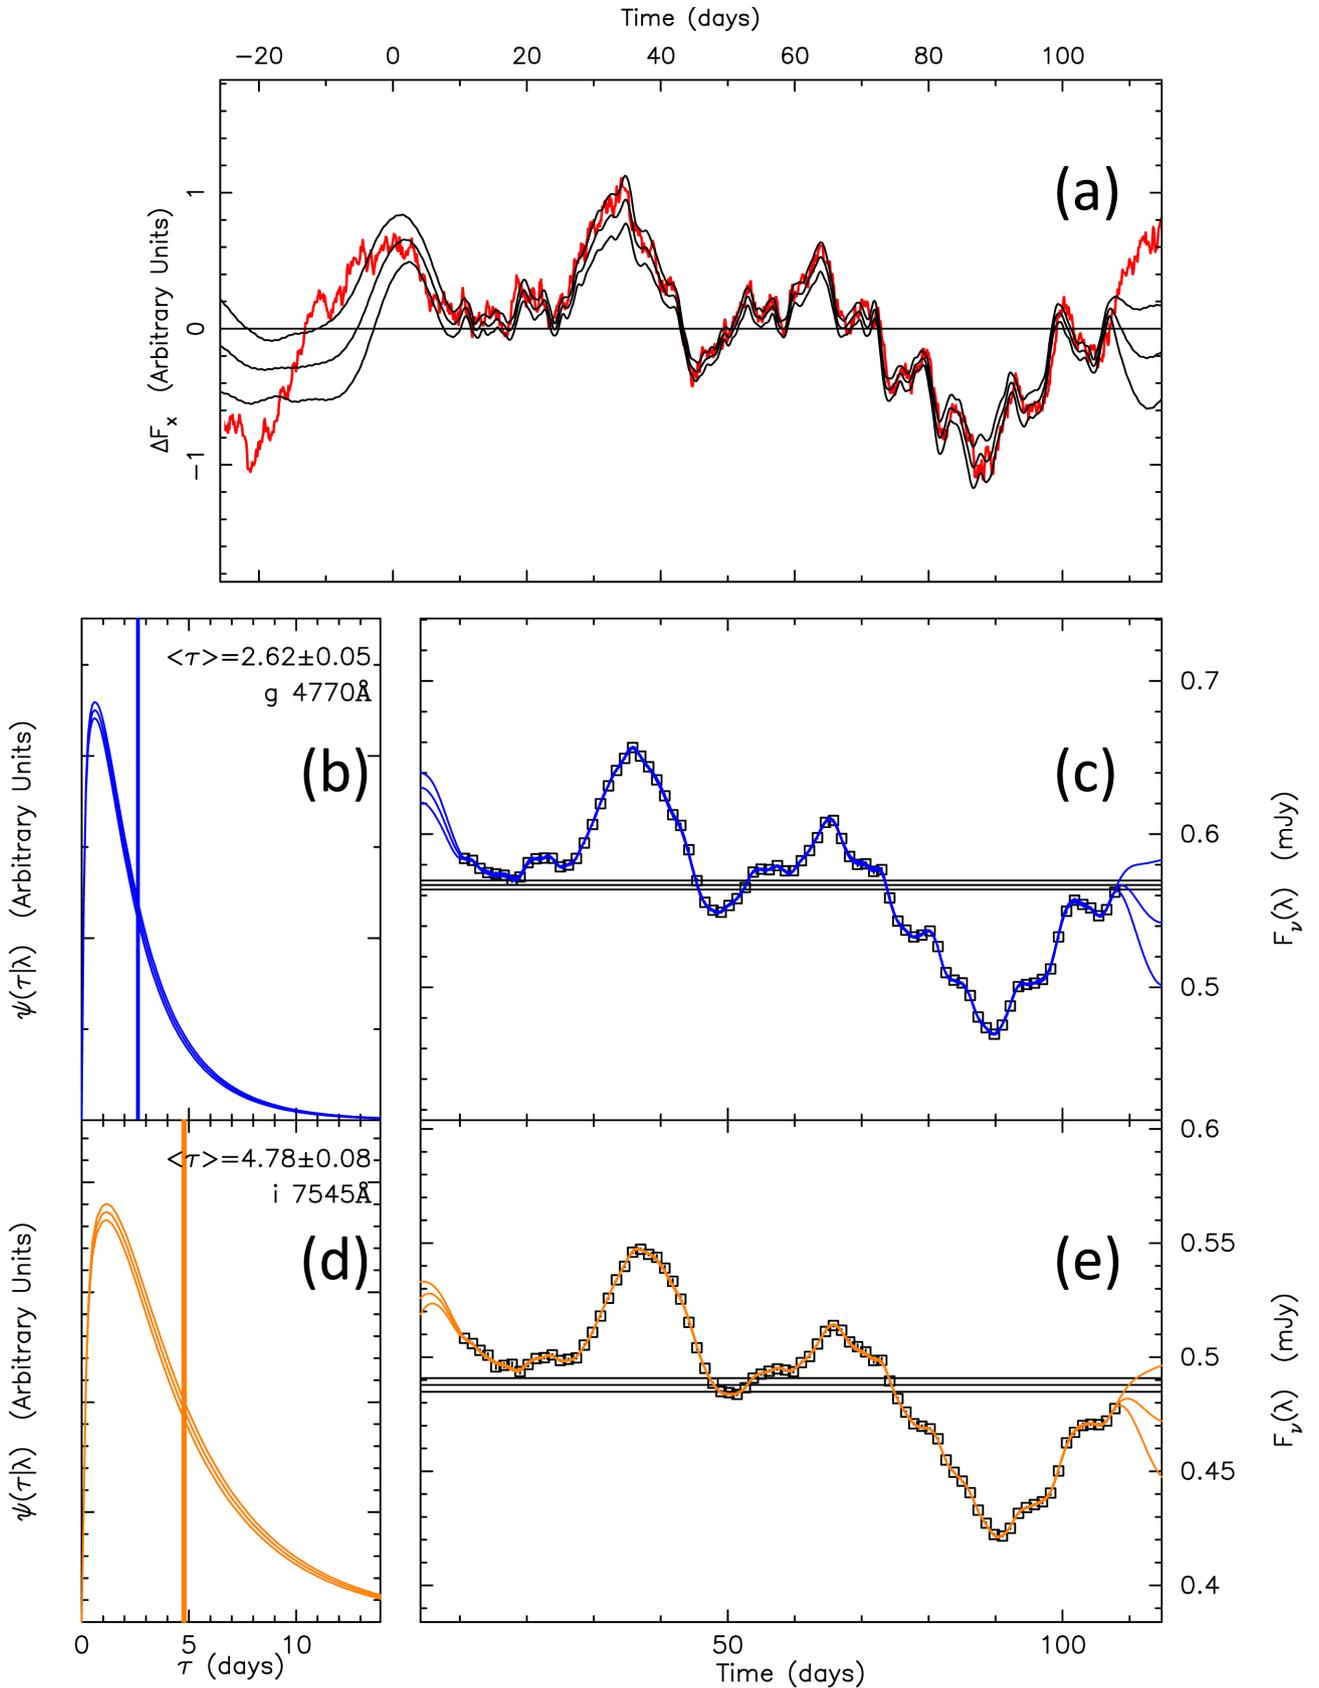


Figure 6. The top plot shows the reconstructed driving light curve (equation 7). The plots below this show (left) the inferred response function (Fig. 3) with vertical lines showing the mean time delay (τ). The right plots show the response light curves in the *g* and *i* filters including 1σ uncertainty envelopes. Horizontal lines indicate the offset parameters and uncertainty envelopes for the response light curves. The red line in Panel (a) shows the driving light curve used to generate the synthetic *g* and *i* light curves, normalized to the rms of the inferred driving light curve.

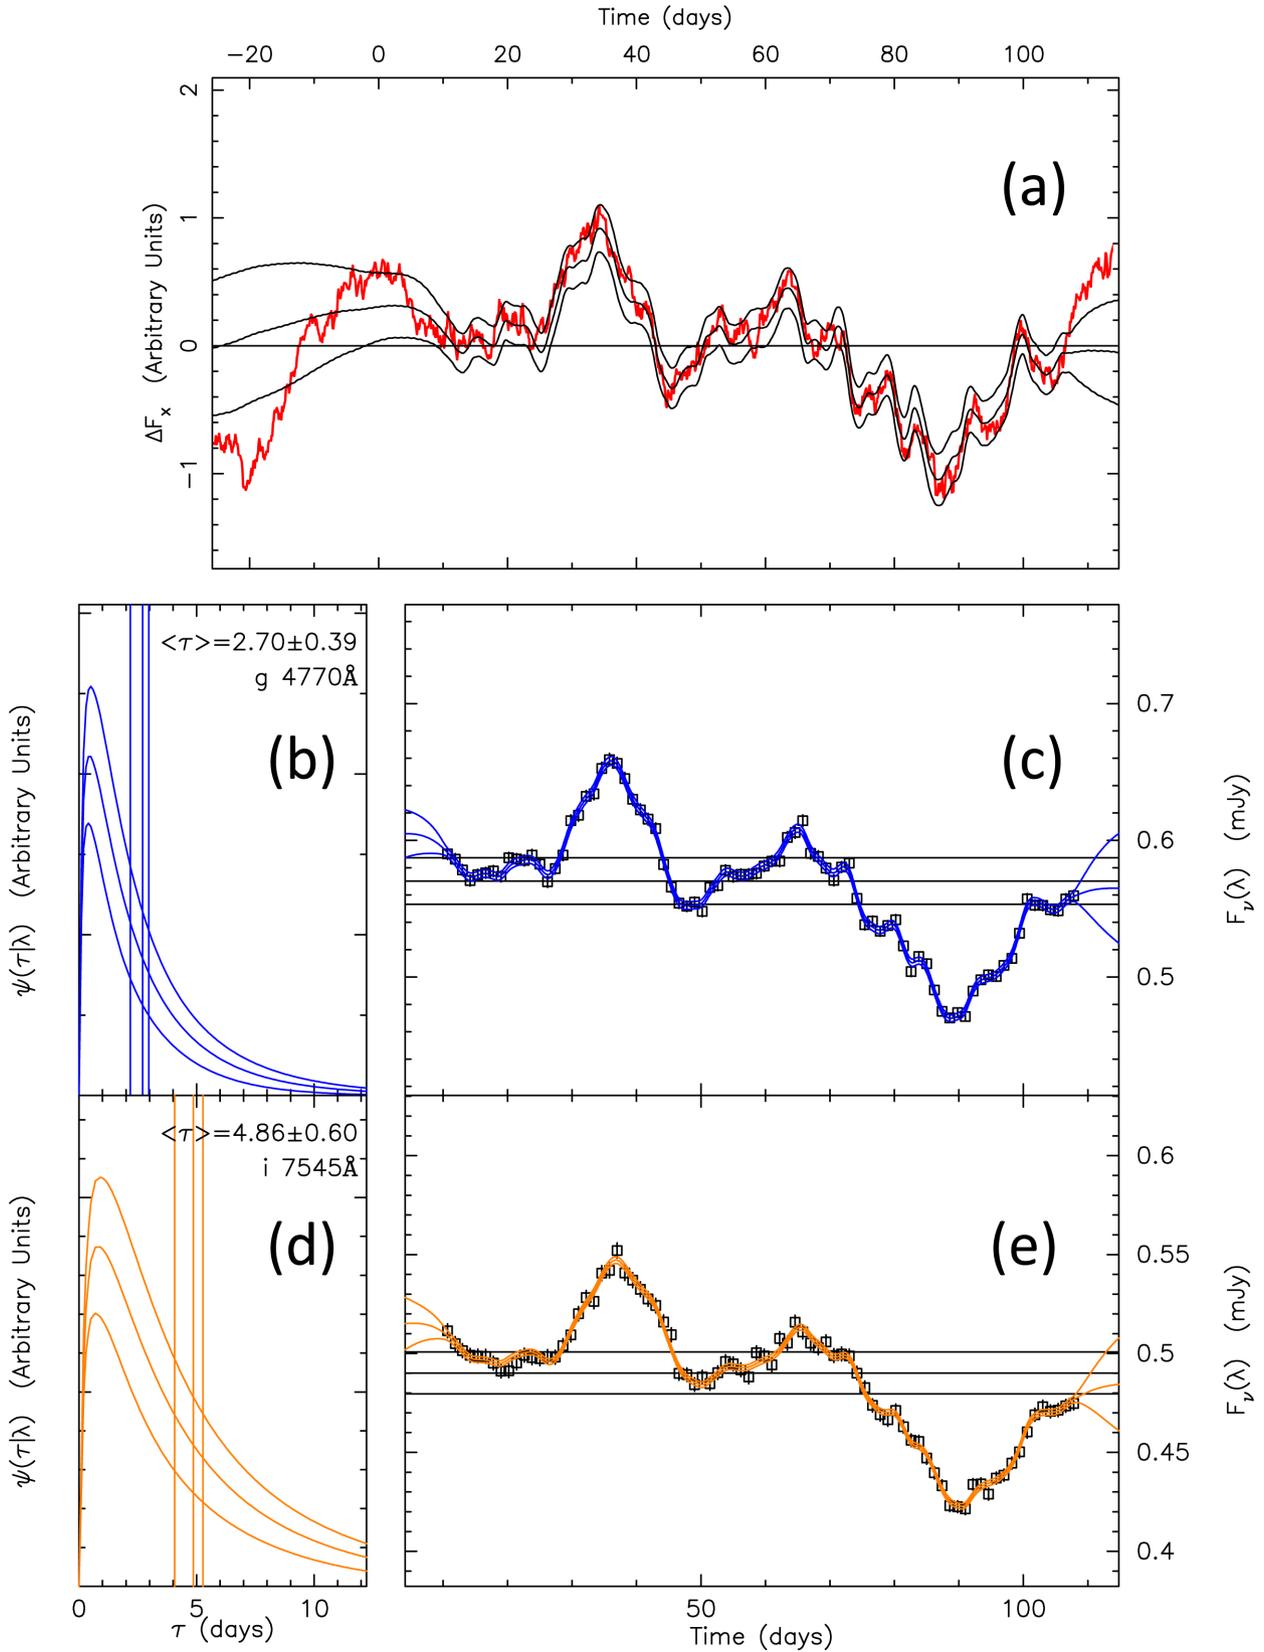


Figure 7. As with Fig. 6 but showing CREAM applied to the noiser SNR = 100 light-curve data.

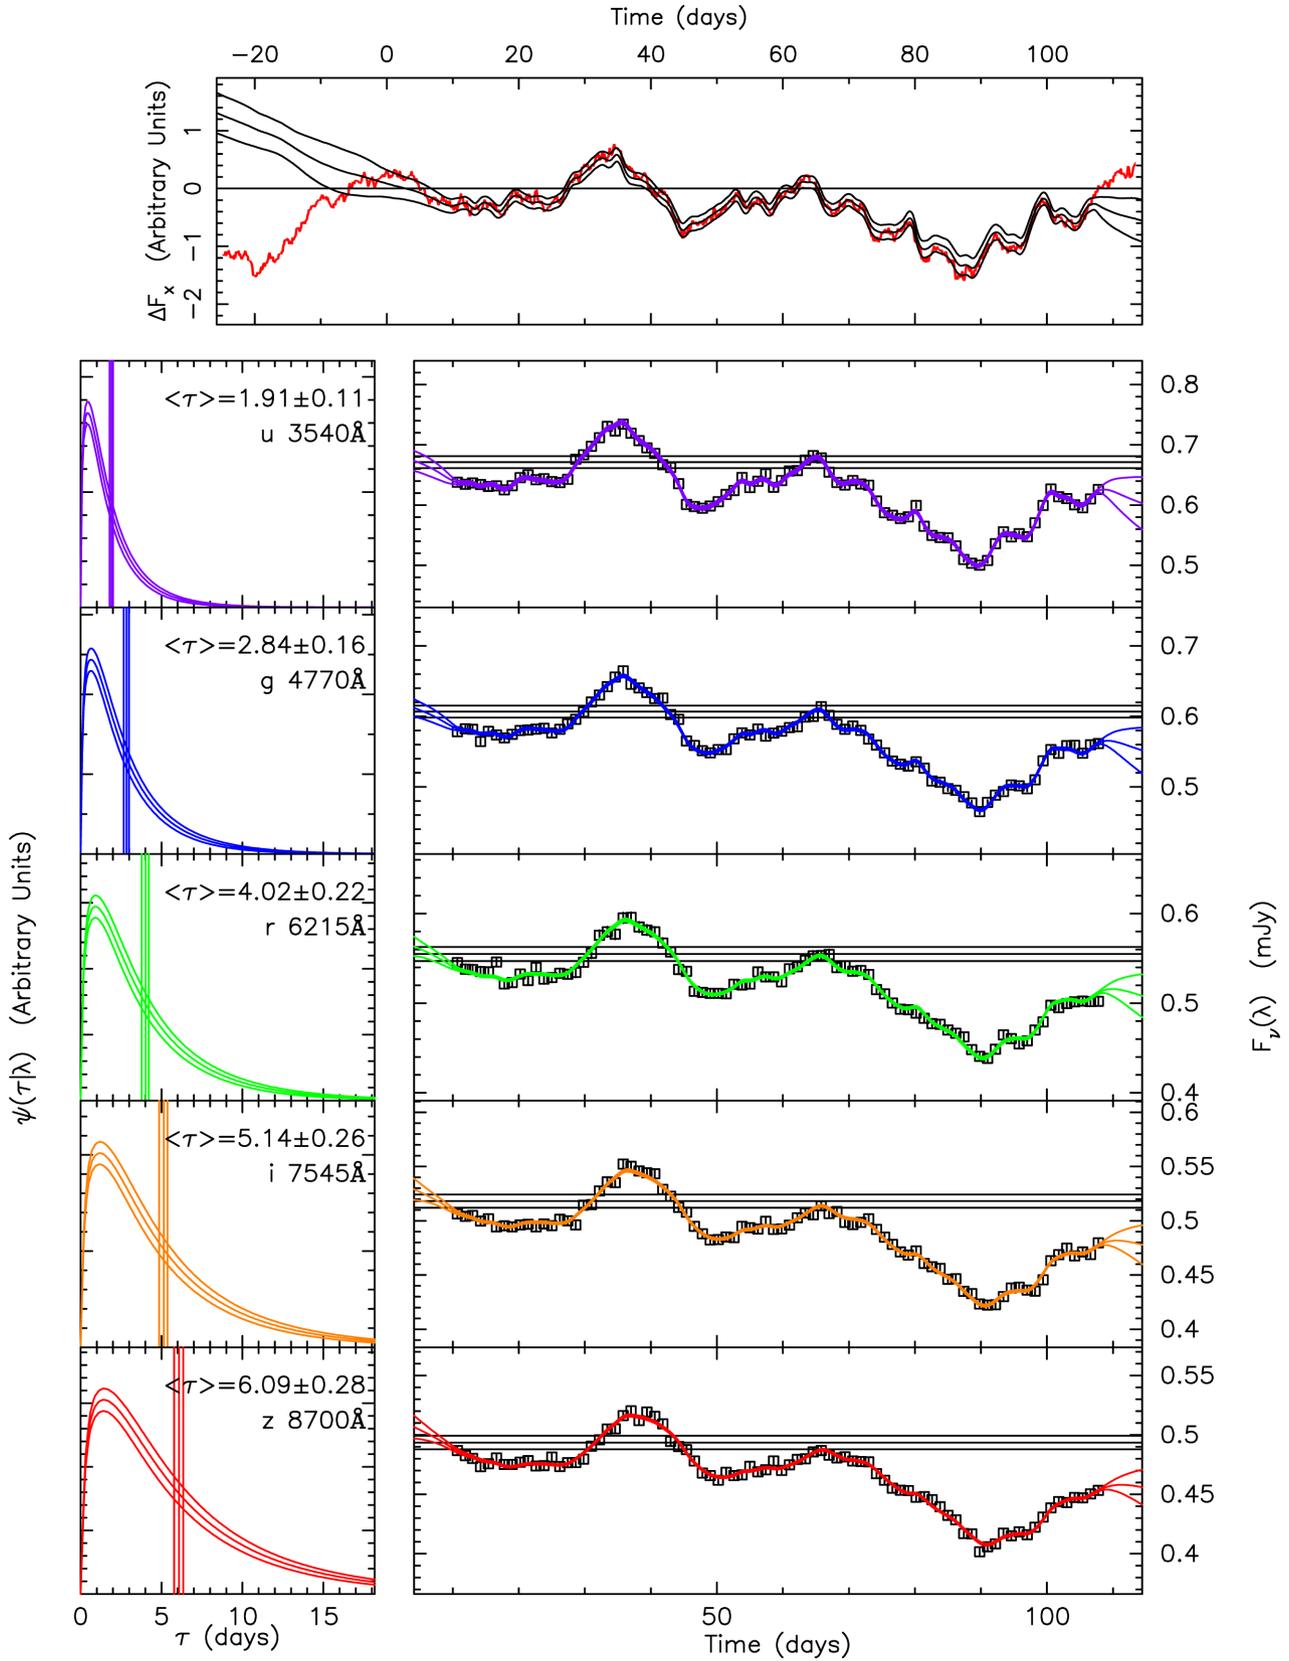


Figure 8. As with Fig. 6 but showing CREAM applied to the SNR = 100 light curves in *ugriz*.

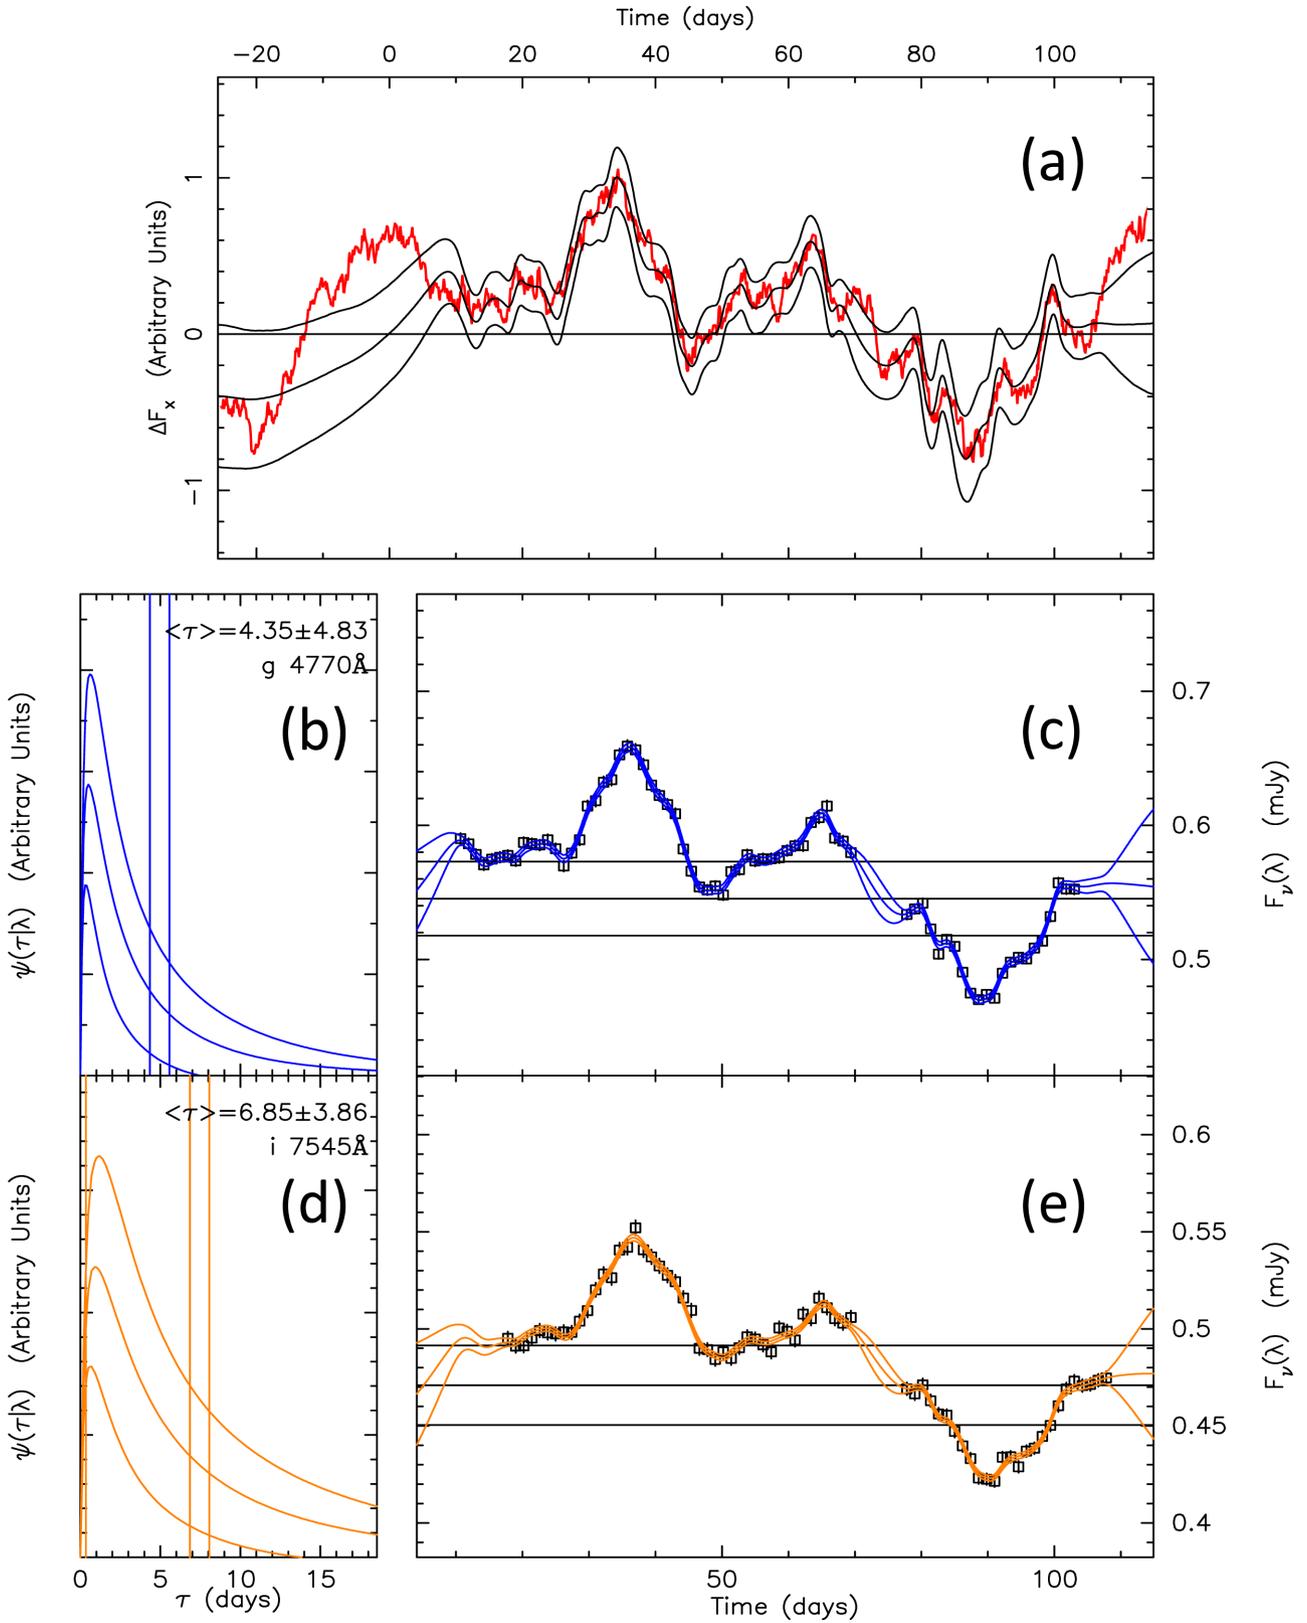


Figure 9. As with Fig. 6 but including the four 1-week gaps in the light curves. Gaps occur at the start of the *g* light curve, in *g* and *i* at 50–57 d and at the end of the *g* light curve.

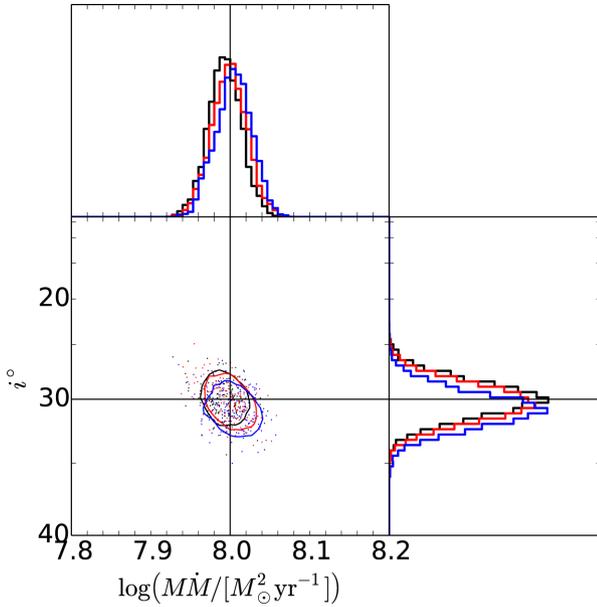


Figure 10. Posterior probability distributions of $M\dot{M}$ and i corresponding to the high-SNR g and i light curves. The three colours indicate separate chains. Contours show the 1σ error regions. Black vertical and horizontal lines show the true parameters ($i = 30^\circ$ and $M\dot{M} = 10^8 M_\odot^2 \text{yr}^{-1}$).

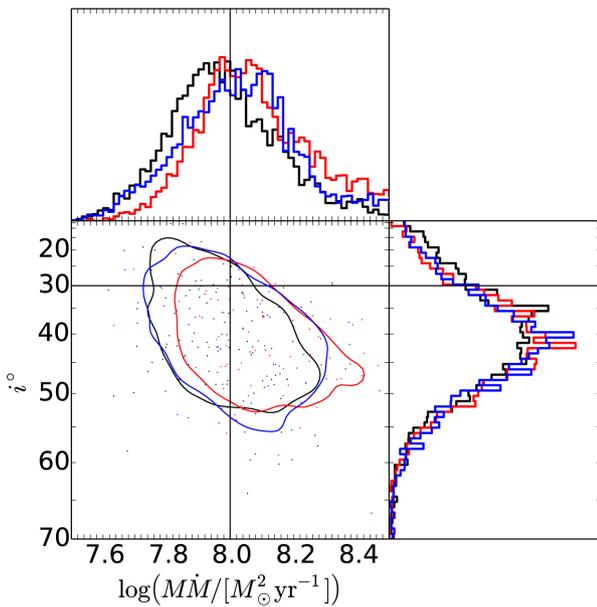


Figure 11. As with Fig. 10 but for the SNR = 100 g and i light curves.

obtain this SNR for our synthetic target. The CREAM fit to the noisier light curve is shown in Fig. 7 and posterior probability distributions in Fig. 11. The uncertainty in inclination increases to $\pm 15^\circ$ and the uncertainty in $M\dot{M}$ rises to ± 0.19 dex. Mean delays now exhibit uncertainties of 0.39 and 0.60 d for g and i , respectively.

These results, while still accurate enough to be of scientific interest, are less attractive than those from the analysis of the higher quality light curves presented in Fig. 6. We therefore investigate how the results improve with the addition of SNR = 100 light curves in u , r and z . These are generated using the same procedure described in Section 4 and the relevant exposure times, and DRW parameters are given in Tables 3 and 2. The same driving light curve is used to

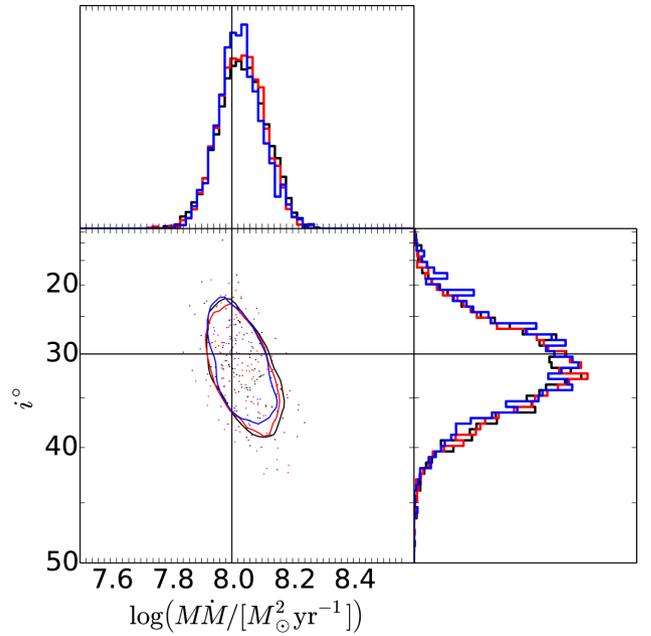


Figure 12. As with Fig. 10 but for CREAM fits to u, g, r, i and z light curves with SNR = 100 (Fig. 8). These light curves are driven by the same driving light curve as the g and i light curves in Fig. 5.

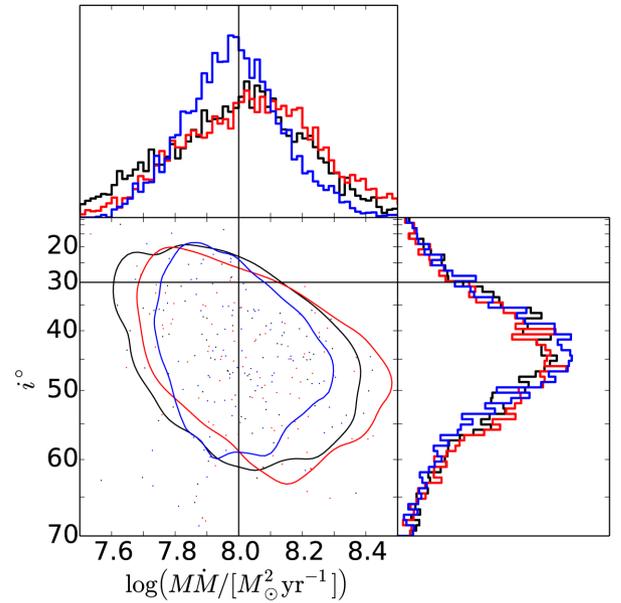


Figure 13. As with Fig. 10 but for the SNR = 100 g and i light curves including the four week-long data gaps (Fig. 9).

generate the additional u , r and z light curves as the earlier g and i data in Fig. 5. A fit to these light curves is shown in Fig. 8. Posterior probability distributions in Fig. 12 show that adding the r , i and z light curves reduces the error by a factor of 2. We see that, for $ugriz$ observations with SNR = 100, CREAM obtains an inclination accurate to 8.5° and $M\dot{M}$ accurate to 0.1 dex.

We note a slight bias to high inclinations in Fig. 11. We test this further by generating an additional 3 driving light curves and obtaining g and i light curves for each using the prescription in Section 4. The posterior probabilities for $M\dot{M}$ and inclination are shown for each of these drivers in Fig. 14. We see that the inferred inclination is sometimes above and sometimes below the true value

Table 4. The $M\dot{M}$ and inclination estimates for each of the three CREAM runs to the light curves shown in the figures stated. A fit to the light curves with high, lunar-phase-dependant SNR's is shown in Fig. 6. The g and i SNR = 100 light curves are fitted in Fig. 7. The $ugriz$ light curves are fitted in Fig. 8. CREAM is fitted to g and i light curves from three different drivers and the posterior probabilities for $M\dot{M}$ and inclination are shown in Fig. 14. CREAM runs on the SNR = 100 light curves with the four week-long gaps are displayed in Figs 9 and 13.

Light curve set	$\log(M\dot{M})$	i (deg)
Figs 6 and 10	8.00 ± 0.03	29.28 ± 5.04
	8.00 ± 0.04	30.84 ± 4.90
	8.01 ± 0.03	31.45 ± 4.52
Figs 7 and 11	7.97 ± 0.17	41.58 ± 13.32
	8.06 ± 0.19	42.40 ± 14.12
	8.03 ± 0.18	43.07 ± 14.62
Figs 8 and 12	8.03 ± 0.08	31.9 ± 8.48
	8.02 ± 0.08	31.97 ± 7.84
	8.03 ± 0.10	31.31 ± 7.96
Fig. 14	8.15 ± 0.18	38.30 ± 8.52
	8.26 ± 0.18	45.10 ± 9.49
	7.93 ± 0.13	18.77 ± 15.62
Figs 9 and 13	8.01 ± 0.24	45.99 ± 18.02
	7.98 ± 0.15	47.78 ± 16.33
	8.05 ± 0.23	46.61 ± 17.23

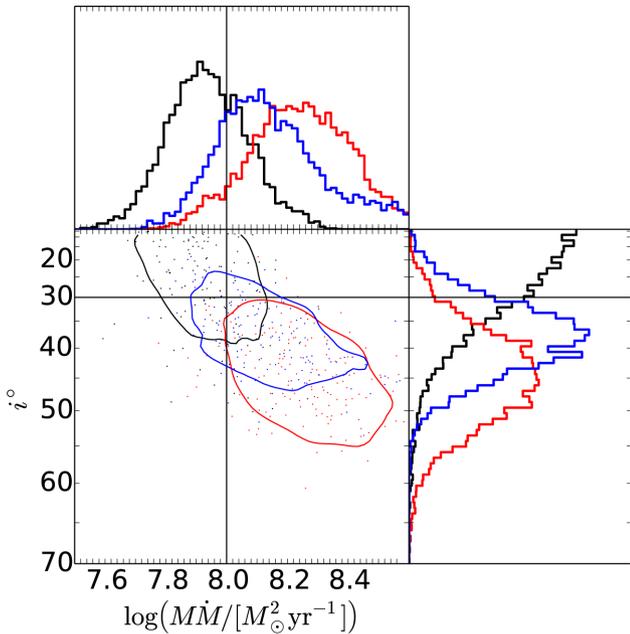


Figure 14. As with Fig. 10 but for CREAM fits to g and i light curves with colours now showing three different driving light curves.

of 30° . This scatter is expected as the 2 parameter 1σ contours should include the true model in 68 per cent of the driving light curves tested.

5.2 Data gaps

The light curves in Fig. 5, while not equally sampled, are fairly regular with mean 1 d cadence. To investigate the possible impact of data gaps (e.g. due to equipment failure or episodes of poor weather), we remove four one-week-long segments from the light curves, the first week from the start of the i light curve, the second

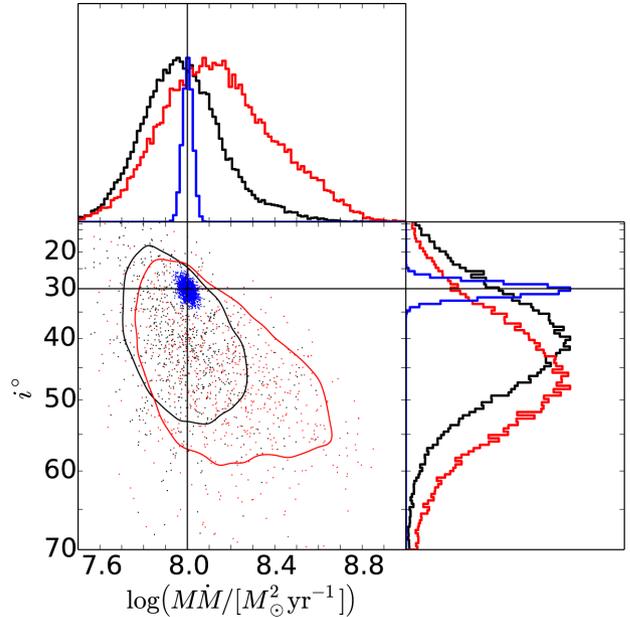


Figure 15. Here, we add together the results from all three chains in Figs 10, 11, and 13. Blue represents the high-SNR light curves, black is the SNR = 100 sample with no gaps and red includes the four week-long gaps in the light curve.

and third weeks from both g and i light curves simultaneously at 50–57 d, and the final week from the end of the i light curve. We perform this test on the noisier g and i light curves with SNR = 100. This is the worst case observing scenario we consider. The CREAM fit to light curves with the data gaps are shown in Fig. 9. The posterior probability distributions for the inclination and $M\dot{M}$ are plotted in Fig. 13. The error envelope for the inferred driving and echo light curves expands inside the gaps as appropriate. It is apparent that CREAM's ability to recover inclination is adversely affected when the gaps are included in the light curves. The uncertainty in the inclination now rises to $\pm 18^\circ$ and the uncertainty in $M\dot{M}$ rises to 0.24 dex. Light curve mean delay uncertainties for g and i again increase to 0.74 and 0.95 d, respectively.

A summary plot showing the posterior probabilities for inclination and $M\dot{M}$ is displayed for ease of comparison between the various tested observing strategies in Fig. 15. The accuracy achieved in all these tests would be of scientific interest if achieved in practice.

5.3 High Frequency Resolution Limits: ω_{hi}

CREAM fits consume computer time due to the large number of parameters in the Fourier series (equation 7). We ran tests lowering ω_{hi} to speed up the fitting. We find no effect on the $\cos i$ or $\log M\dot{M}$ parameter estimates so long as the driving light curve is evaluated with upper frequencies of 0.5 cycle / day or higher. This approximately matches the inferred $\langle \tau \rangle$ for the i light curve in Fig. 11 suggesting future observations at shorter wavelengths than those presented here may need to be run with larger values of ω_{hi} higher resolution.

6 DISCUSSION AND CONCLUSIONS

We have introduced a new approach (CREAM) to measure the continuum lag distribution of reverberating AGN accretion discs. The model assumes a flat optically-thick accretion disc that undergoes

temperature perturbations due to irradiation by a variable driving point source (lamp post) located above the disc centre. Time lags increase with λ and arise due to light-travel-time effects that are longer for the cooler, longer-wavelength-emitting regions located at larger radii. Observations of the power spectrum of X-ray light curves suggest that X-rays may lead the variability (McHardy et al. 2014) making them a strong candidate for the lamp post. This variability exhibits a random walk power spectrum $P(f) \propto f^{-2}$ that we adopt to constrain our model parameters.

This approach is similar to SPEAR and JAVELIN (Zu et al. 2011) which use DRW light curves in an attempt to model an input light curve (usually the shortest wavelength light curve is a proxy for the driving light curve) and model each echo light curve individually assuming a top hat delay distribution. CREAM can infer both MM and disc inclination parameters of the delay distribution, and the driving light curve with no input driver required.

To test CREAM we generated synthetic optical light curves using the thermal reprocessing model assuming a delay distribution appropriate for an $MM = 10^8 M_{\odot}^2 \text{ yr}^{-1}$ accretion disc inclined at 30° . The light curves are randomly sampled with mean 1 d cadence over 100 d and mimic the SNR expected from 1000 s observations in g and i with a 2-m telescope.

We emphasize that CREAM does not require data to act as a proxy for the variable irradiation (the driving light curve). Since we do not fully understand the origin of continuum variability (Lohfink et al. 2014) this is an exciting feature of the code. We anticipate comparing CREAM's inference of the driving light curve with various UV and X-ray light curves to investigate which of these bands CREAM best agrees with.

In this work, we use tests with simulated data to demonstrate CREAM's ability to recover MM to ± 0.04 dex and disc inclinations to $\pm 5^\circ$ with observations at 1 d cadence in g and i filters. We note also that with noisier light curves (SNR = 100) our inferred inclination uncertainty increases to $\pm 15^\circ$. The uncertainty in MM increases to 0.2 dex. If we include overlapping SNR = 100 light curves in u , r and z we are able to infer MM to ± 0.1 dex and inclination to $\pm 8.5^\circ$. Any of these would be of scientific interest.

While CREAM recovers delay distributions for continuum-continuum lags, it does so by applying a number of simplifications. CREAM assumes that variability in continuum light curves arises solely due to lamp post thermal reprocessing. There may in fact be contaminating contributions to light curve variability from the BLR. Korista & Goad (2001) note the presence of diffuse continuum emission from the BLR. The diffuse-cloud-continuum emission is a combination of reflected accretion disc continuum and thermal-non-line emission. If present, diffuse continuum emission might increase the time delays, causing CREAM to overestimate the MM and mean delay. This would also affect the mean delays as obtained by SPEAR and CCF.

In an upcoming paper, we will apply CREAM to 19 overlapping, multiwavelength continuum light curves taken by the STORM collaboration (De Rosa et al. 2015; Edelson et al. 2015). We aim to both measure the disc inclination, accretion rate and infer the shape of the driving light curve; testing the agreement between the X-ray light curve and CREAM's inferred driving light curve to determine how the X-rays are related to continuum variability.

In summary, the ability to constrain accretion disc delay distribution functions affords CREAM many potential applications. Not least

of these is the information we can obtain on the accretion rate and disc inclination. Multiwavelength observations of NGC 5548 (De Rosa et al. 2015; Edelson et al. 2015) offer real data sets on which to apply CREAM to measure the accretion disc parameters. Finally, since no data are needed to act as the driver of variability, the driving light curve can be recovered from CREAM fits to observed echo light curves and compared to observations at different wavelengths to help identify the true driver of AGN variability.

ACKNOWLEDGEMENTS

We thank the referee for the helpful comments and suggestion made during the submission of this manuscript. DAS acknowledges the support of the Science and Technologies Funding Council studentship. KH acknowledges support from the UK Science and Technology Facilities Council (STFC) consolidated grant to St Andrews (ST/M001296/1).

REFERENCES

- Arévalo P., Uttley P., Kaspi S., Breedt E., Lira P., McHardy I. M., 2008, MNRAS, 389, 1479
- Bentz M. C. et al., 2013, ApJ, 767, 149
- Blandford R. D., McKee C. F., 1982, ApJ, 255, 419
- Cackett E. M., Horne K., Winkler H., 2007, MNRAS, 380, 669
- Chelouche D., 2013, ApJ, 772, 9
- De Rosa G. et al., 2015, ApJ, 806, 128
- Dexter J., Agol E., 2011, ApJ, 727, L24
- Edelson R. A., Krolik J. H., 1988, ApJ, 333, 646
- Edelson R. et al., 2015, ApJ, 806, 129
- Elvis M., 2000, ApJ, 545, 63
- Frank J., King A., Raine D., 2002, *Accretion Power in Astrophysics*. Cambridge Univ. Press, Cambridge
- Gaskell C. M., Peterson B. M., 1987, ApJS, 65, 1
- Hönig S. F., Watson D., Kishimoto M., Hjorth J., 2014, Nature, 515, 528
- Kelly B. C., Bechtold J., Siemiginowska A., 2009, ApJ, 698, 895
- Kishimoto M., Hönig S. F., Antonucci R., Barvainis R., Kotani T., Tristram K. R. W., Weigelt G., Levin K., 2011, A&A, 527, A121
- Korista K. T., Goad M. R., 2001, ApJ, 553, 695
- Lohfink A. M., Reynolds C. S., Vasudevan R., Mushotzky R. F., Miller N. A., 2014, ApJ, 788, 10
- McHardy I. M., Koerding E., Knigge C., Uttley P., Fender R. P., 2006, Nature, 444, 730
- McHardy I. M. et al., 2014, MNRAS, 444, 1469
- McHardy I. et al., 2015, in *The Extremes of Black Hole Accretion*, 49 available at: http://xmm.esac.esa.int/external/xmm_science/workshops/2015_science/
- MacLeod C. L. et al., 2010, ApJ, 721, 1014
- MacLeod C. L. et al., 2015, ApJ, 806, 258
- Malkan M. A., 1983, ApJ, 268, 582
- Morgan C. W. et al., 2012, ApJ, 756, 52
- Shang Z. et al., 2005, ApJ, 619, 41
- Shankar F., Weinberg D. H., Miralda-Escudé J., 2009, ApJ, 690, 20
- Urry C. M., Padovani P., 1995, PASP, 107, 803
- Uttley P., McHardy I. M., Papadakis I. E., 2002, MNRAS, 332, 231
- Zu Y., Kochanek C. S., Peterson B. M., 2011, ApJ, 735, 80
- Zu Y., Kochanek C. S., Kozłowski S., Udalski A., 2013, ApJ, 765, 106

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