

The evolution of the X-ray luminosity functions of unabsorbed and absorbed AGNs out to $z \sim 5$

J. Aird^{1,2*}, A. L. Coil³, A. Georgakakis^{4,5}, K. Nandra⁴, G. Barro⁶ and P. G. Pérez-González⁷

¹*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge, CB3 0HA*

²*Department of Physics, Durham University, Durham, DH1 3LE*

³*Center for Astrophysics and Space Sciences (CASS), Department of Physics, University of California, San Diego, CA 92093, USA*

⁴*Max Planck Institute für Extraterrestrische Physik, Giessenbachstrasse, 85748 Garching, Germany*

⁵*IAASARS, National Observatory of Athens, GR-15236 Penteli, Greece*

⁶*University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA*

⁷*Departamento de Astrofísica, Facultad de CC. Físicas, Universidad Complutense de Madrid, E-28040 Madrid, Spain*

Submitted to Monthly Notices of the Royal Astronomical Society on 2015 February 24.

ABSTRACT

We present new measurements of the evolution of the X-ray luminosity functions (XLFs) of unabsorbed and absorbed Active Galactic Nuclei (AGNs) out to $z \sim 5$. We construct samples containing 2957 sources detected at hard (2–7 keV) X-ray energies and 4351 sources detected at soft (0.5–2 keV) energies from a compilation of *Chandra* surveys supplemented by wide-area surveys from *ASCA* and *ROSAT*. We consider the hard and soft X-ray samples separately and find that the XLF based on either (initially neglecting absorption effects) is best described by a new flexible model parametrization where the break luminosity, normalization and faint-end slope all evolve with redshift. We then incorporate absorption effects, separately modelling the evolution of the XLFs of unabsorbed ($20 < \log N_{\text{H}} < 22$) and absorbed ($22 < \log N_{\text{H}} < 24$) AGNs, seeking a model that can reconcile both the hard- and soft-band samples. We find that the absorbed AGN XLF has a lower break luminosity, a higher normalization, and a steeper faint-end slope than the unabsorbed AGN XLF out to $z \sim 2$. Hence, absorbed AGNs dominate at low luminosities, with the absorbed fraction falling rapidly as luminosity increases. Both XLFs undergo strong luminosity evolution which shifts the transition in the absorbed fraction to higher luminosities at higher redshifts. The evolution in the shape of the total XLF is primarily driven by the changing mix of unabsorbed and absorbed populations.

Key words: galaxies: active – galaxies: evolution – galaxies: luminosity function, mass function – X-rays: galaxies

1 INTRODUCTION

The luminosity function of Active Galactic Nuclei (AGNs) represents one of the crucial observational constraints on the growth of supermassive black hole (SMBHs) over the history of the Universe. The shape of the luminosity function reflects a combination of the underlying distribution of the SMBH masses and the distribution of their accretion rates or Eddington ratios (e.g. Aird et al. 2013a; Shankar et al. 2013; Schulze et al. 2015). Thus, accurate measurements of the shape and evolution of the luminosity function provide crucial insights into the physical processes that drive SMBH growth over cosmic time.

Many AGNs are surrounded by gas and dust that can obscure their emission at certain wavelengths. Thus, it is vital to understand AGN obscuration in order to obtain accurate measurements of

the luminosity function. Quantifying AGN obscuration also reveals whether SMBHs undergo significant periods of obscured growth, when this takes place within the lifetimes of AGNs, and how it relates to the triggering and fuelling processes.

Obtaining accurate measurements of the luminosity function and revealing the extent of obscuration requires large, unbiased samples of AGNs selected over the widest possible range of redshifts and luminosities. Optical surveys, combined with follow-up spectroscopy, can efficiently cover wide areas (e.g. SDSS: York et al. 2000) but are biased towards the most luminous, unobscured sources. Alternatively, AGNs can be identified in the mid-infrared (mid-IR), which probes the reprocessed emission from the dusty, obscuring material. Mid-IR selection should not be biased against obscured sources, but the contribution of the host galaxy is often significant at these wavelengths, which limits mid-IR selection to luminous sources where the galaxy light is overwhelmed by the AGN (e.g. Donley et al. 2008; Mendez et al. 2013).

* jaird@ast.cam.ac.uk

X-ray surveys can efficiently identify AGNs over a wide luminosity range, including low-luminosity sources where the host galaxy dominates at optical or infrared wavelengths (e.g. Barger et al. 2003). Nevertheless, soft X-ray emission (at energies $\lesssim 2$ keV) will be absorbed by the same gas and dust that obscures the AGN at optical and UV wavelengths. Thus, soft X-ray samples are generally dominated by unobscured AGNs. Absorption biases are reduced at hard X-ray energies ($\sim 2 - 10$ keV), except in the most heavily-obscured, Compton-thick AGNs (equivalent line-of-sight hydrogen column densities $N_{\text{H}} \gtrsim 10^{24} \text{ cm}^{-2}$). However, even Compton-thick sources may still be identified at soft or hard X-ray energies due to scattered emission, including the Compton-scattered emission (“reflection”) from the obscuring material itself.

A large number of deep and wide X-ray surveys have been carried out, taking advantage of the efficiency and power of X-ray selection (see a recent review by ?). A number of studies have measured the X-ray luminosity function (XLF) of AGNs out to high redshifts using these samples (e.g. Ueda et al. 2003; Barger et al. 2005; ?). These studies find that AGNs are a strongly evolving population, with a sharp decrease in their number density between $z \sim 1 - 2$ and today. Bright X-ray-selected AGNs are found to peak in number density at $z \approx 2$, similar to optically-selected QSOs. Fainter AGNs peak later in the history of the Universe ($z \approx 1$) but with a much milder decline to the present day (e.g. Hasinger et al. 2005). These patterns have led several authors to propose a luminosity-dependent density evolution (LDDE) parametrization to describe the evolution of the XLF of AGNs (e.g. Miyaji et al. 2000; Ueda et al. 2003). In this model the XLF is modified by differing degrees of density evolution that vary with luminosity and redshift. This results in an XLF that changes shape over cosmic time.

In Aird et al. (2010, hereafter A10) we challenged this evolutionary model with a detailed study of the hard-band XLF that carefully accounted for numerous uncertainties and biases that were generally not included in prior measurements. These included flux measurement errors, Eddington bias, incompleteness of optical identifications, and the uncertainty in photometric redshift estimates. At high redshifts ($z \sim 2 - 3$) we adopted a rest-frame UV colour pre-selection technique (Aird et al. 2008). By performing a robust model comparison based on Bayesian statistical techniques, we found that the evolution of the XLF could be described by a simpler model in which the XLF retains the same shape at all redshifts but evolves in both luminosity and density (see also Assef et al. 2011; Ross et al. 2013).

While A10 presented a number of important advances, absorption effects were not explored. Other studies have attempted to measure the distribution of absorption column densities and present absorption-corrected measurements of the XLF. Ueda et al. (2003) found that the fraction of absorbed AGNs (those with $N_{\text{H}} > 10^{22} \text{ cm}^{-2}$) was strongly dependent on luminosity, decreasing at higher luminosities. Later studies found that the fraction of absorbed AGNs depends on both luminosity *and* redshift, dropping at high luminosities but increasing (at a given luminosity) to higher redshifts (e.g. La Franca et al. 2005; Hasinger 2008). The extent of any redshift evolution has been a matter of debate (see Akylas et al. 2006; Dwelly & Page 2006), potentially due to difficulties regarding the selection functions for absorbed and unabsorbed sources. Recent work by Ueda et al. (2014) re-examined the evolution of the XLF and the distribution of N_{H} (the “ N_{H} -function”) using a large compilation of both soft and hard X-ray surveys and found that both a luminosity and redshift dependence of the absorbed fraction were required. They also found that an LDDE parametriza-

tion was needed to describe the evolution of the XLF (with some further modifications to describe the evolution at $z \gtrsim 3$, see also Civano et al. 2011; Hiroi et al. 2012).

Recently, the combination of extremely deep X-ray survey data and new approaches to X-ray spectral analysis have enabled improved measurements of N_{H} at $z \sim 0.5 - 2$ and have been used to identify sizable samples of Compton-thick AGNs (e.g. Comastri et al. 2011; Georgantopoulos et al. 2013; Brightman et al. 2014). Building on this work, Buchner et al. (2015) used a flexible, non-parametric method to estimate the space densities of AGNs as a function of redshift, luminosity and N_{H} , effectively measuring the XLF for different column densities. This work also recovered a luminosity and redshift dependence in the evolution of the fraction of absorbed AGNs (although the Compton-thick fraction was consistent with a constant value of ~ 35 per cent). However, a detailed comparison of parametric models for the evolution of the XLF of AGNs was not undertaken.

In this paper we address some remaining issues in studies of the evolution of the XLF of AGNs: the shape of the XLF and how it evolves with redshift, the extent of any luminosity and redshift dependence of the absorbed fraction, and the connection between the absorption properties and the evolution of the AGN population. We combine samples selected at both hard and soft X-ray energies and determine the underlying XLF and distribution of N_{H} that adequately describes the observed fluxes in both samples (similar to the approach of Ueda et al. 2014, cf. the X-ray spectral analysis used in Buchner et al. 2015).

In Section 2 we describe our datasets that we use to define large samples of X-ray sources selected in the hard (2–7 keV) and soft (0.5–2 keV) energy bands. We also compile deep optical, near-infrared and mid-infrared imaging across our fields that we use to robustly identify counterparts to our X-ray sources and calculate photometric redshifts. In Section 3, we describe our Bayesian statistical technique that allows us to incorporate a range of X-ray spectral shapes and account for the effects of absorption. We also introduce an approach to account for the contribution from normal, X-ray detected galaxies on our measurements. We then present measurements of the XLF based on our hard and soft samples individually (Section 4), introducing a new flexible parametrization of the XLF. We show that significant discrepancies between the measurements at all redshifts warrant the further consideration of absorption effects. In Section 5 we separately model the XLF of unabsorbed and absorbed AGNs (including a contribution from Compton-thick sources) and show how the combination of these populations can simultaneously account for both our hard- and soft-band samples. Our results place constraints on the total XLF of AGNs and the absorbed fraction as a function of luminosity and redshift. In Section 6 we compare our results to prior work and discuss the wider implications of our findings. Section 7 summarizes our paper and overall conclusions.

Given the length of this paper, a casual reader may wish to skip to Section 4.4, Section 5 and the discussion in Section 6 (and focus on Figures 7, 8, and 9). We adopt a flat cosmology with $\Omega_{\Lambda} = 0.7$ and $h = 0.7$ throughout this paper.

2 DATA

To constrain the evolution of the XLF we require large samples of X-ray selected AGNs. By selecting samples in both the hard (> 2 keV) and soft (0.5–2 keV) observed energy bands, we can also constrain the distribution of N_{H} and correct for these effects on the

XLF. In this paper we combine a large number of *Chandra* X-ray surveys along with larger area surveys from *ASCA* and *ROSAT*. We give further details of our datasets below. Table 1 summarizes the different surveys and provides the number of hard and soft X-ray selected sources from each.

2.1 *Chandra* X-ray data

In this paper we use *Chandra* X-ray observations from five distinct parts of the sky: the Chandra Deep Field-South (CDFs), Chandra Deep Field-North (CDFN), Extended Groth Strip (EGS), COSMOS, and Bootes fields.

In the CDFS field we identify two different “surveys”: 1) the series of observations that have targeted the central $\sim 0.07 \text{ deg}^2$ of the field and reach a total combined exposure time of $\sim 4\text{Ms}$ (the CDFS-4Ms survey, Xue et al. 2011); and 2) the series of four 250ks observations that surround the central area (the Extended-CDFS, or E-CDFS survey, Lehmer et al. 2005). The X-ray data reduction and source detection procedures were carried out independently for each of these surveys. To combine the two surveys, we define a central region where the CDFS-4Ms survey takes precedence, roughly corresponding to the area within $\sim 9'$ of the center of the field. Outside this central region we adopt the sources detected in the E-CDFS survey only. This procedure ensures we have a well-defined sample where we can accurately determine the sensitivity.

In the EGS field we also identify two distinct surveys: 1) the series of eight pointings that reach a nominal depth of $\sim 200\text{ks}$ and were presented in Laird et al. (2009), which we refer to as the AEGIS-XW(ide) survey; and 2) the AEGIS-XD(eep) survey (?) which took three of the original 200ks to a depth of $\sim 800\text{ks}$. We adopt the deeper AEGIS-XD observations when available.

In each of the remaining fields we adopt data from a single *Chandra* survey: the 2Ms survey in CDFN (Alexander et al. 2003); the $\sim 160\text{ks}$ C-COSMOS observations (Elvis et al. 2009; Puccetti et al. 2009); and the 5ks XBootes survey (Murray et al. 2005).

The X-ray data from all of our surveys were reduced with our own pipeline procedure, which is described in detail by Laird et al. (2009) and ?. We performed point source detection using the procedure described by Laird et al. (2009) and applied a false Poisson probability threshold of $< 4 \times 10^{-6}$ to generate catalogues of detected sources in the soft (0.5–2 keV), hard (2–7 keV), full (0.5–7 keV) or ultrahard (4–7 keV) observed energy bands. We combined the source lists in each band to create a merged catalogue, which is used in the counterpart identification procedure described in Section 2.2 below.

For the 5ks XBootes survey we applied a stricter false probability cut ($< 10^{-8}$) in addition to a requirement of ≥ 5 total detected counts in each band. This cut reduces the sample size but applies an effective X-ray flux limit that helps raise the completeness of the spectroscopic follow-up in this field. In this field, we also restrict our analysis to the $\sim 7.1 \text{ deg}^2$ of the Bootes field that corresponds to the 15 standard sub-fields of the AGES spectroscopic survey (Kochanek et al. 2012). This cut also ensures we have a high spectroscopic completeness (see Section 2.5 below).

We determined X-ray sensitivity maps and area curves for each band as described in Georgakakis et al. (2008), accounting for the stricter false probability cut and minimum counts requirement for the XBootes survey. We convert the sensitivity maps to area curves as a function of flux by assuming a fixed X-ray spectral slope of $\Gamma = 1.4$. We note that the flux calculated with this fixed conversion factor scales directly with the count rate; in Section 3

below we describe our procedure to convert between the count rate and intrinsic quantities (such as luminosity), which allows for a more complex X-ray spectrum and accounts for uncertainties in this conversion factor. The sensitivity map calculation is limited to the footprint of the multiwavelength photometry for each field. For the CDFS and EGS fields we adopt the sensitivity from the appropriate survey at each position and calculate a combined area curve. Figure 1 shows the corresponding area curves for each of our X-ray fields.

We note that previously published X-ray source catalogues, often including multiwavelength counterpart information, are available for the all of our *Chandra* fields (e.g. Alexander et al. 2003; Goulding et al. 2012; Brand et al. 2006). Adopting our own X-ray reductions and source detection procedures ensures we can accurately determine the sensitivity in a consistent manner, which is essential for our Bayesian analysis of the XLF. Our catalogues contain $\sim 10 - 25$ per cent fewer sources than Xue et al. (2011) and Puccetti et al. (2009) in the CDFS-4Ms and C-COSMOS areas, mainly due to our stricter (i.e. lower) false probability threshold. Thus our catalogues are more conservative and should be more robust.

2.2 Multiwavelength counterparts and photometry

In each of our *Chandra* fields, we identify multiwavelength counterparts to sources in our X-ray catalogues using the likelihood ratio (LR) method (e.g. Ciliegi et al. 2003; Brusa et al. 2007; Civano et al. 2012), matching to multiple optical, near-IR and mid-IR bands to ensure a high completeness and reliability. We also compile multiwavelength photometry from a larger number of bands, which we use to calculate photometric redshifts (see Section 2.6 below).

2.2.1 CDFS, CDFN and EGS

In three of our fields—the CDFS, CDFN and EGS—we identify counterparts and extract the multiwavelength photometry using a custom version of the *Rainbow* Cosmological Surveys Database¹ (??Barro et al. 2011a,b), which provides a compilation of the various photometric datasets. Appendix A lists all the different photometric imaging datasets for each field that are used in this paper. We note that full coverage of the entire field is not always available in each photometric band. All the images are registered to a common astrometric reference frame and photometry is performed in consistent apertures to produce spectral energy distributions that span from the UV to mid-IR. The X-ray source matching procedure is described in detail in ?, but we briefly summarise the method here.

First, we reran the *Rainbow* photometric code, extracting all potential counterparts within $3.5''$ of the X-ray positions in any of the bands covered by *Rainbow* (using initial SExtractor catalogues in each of the bands). The counterparts were then cross-matched using a $2''$ search radius to create a single multiband catalogue. We obtained consistent photometry by applying a single aperture across all optical and near-IR bands. We also extracted IRAC photometry, applying the procedure described in ? and Barro et al. (2011a) to deblend the IRAC photometry when a single IRAC source is associated with multiple optical/near-IR counterparts.

Next, we calculated the LR for all candidate counterparts detected in the IRAC $3.6\mu\text{m}$ band and determined an LR threshold

¹ <https://rainbowx.fis.ucm.es>

Table 1. Details of the X-ray surveys used in this paper.

Field	Survey	RA	Dec	X-ray exposure	Survey area (deg ²)	Soft band			Hard band		
		(J2000)	(J2000)			N_X	N_{ctrprt}	N_{spec-z}	N_X	N_{ctrprt}	N_{spec-z}
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
CDFS	CDFS-4Ms	03:32:27.2	-27:47:55	4Ms	0.075	413	397 (96.1%)	240 (58%)	283	273 (97%)	155 (55%)
CDFS	ECDFS	03:32:27.2	-27:47:55	250ks	0.181	334	328 (98.2%)	115 (34%)	273	268 (98%)	101 (37%)
CDFN	CDFN	12:36:49.3	+62:13:19	2Ms	0.112	384	363 (94.5%)	242 (63%)	286	273 (96%)	176 (62%)
EGS	AEGIS-XD	14:19:20.8	+52:50:03	800ks	0.260	698	673 (96.4%)	295 (42%)	552	539 (98%)	233 (42%)
EGS	AEGIS-XW	14:17:15.0	+52:25:31	200ks	0.204	334	332 (99.4%)	130 (39%)	274	274 (100%)	110 (40%)
COSMOS	C-COSMOS	10:00:20.3	+02:11:20	160ks	0.984	1213	1195 (98.5%)	694 (57%)	889	877 (99%)	530 (60%)
Bootes	XBootes	14:31:28.3	+34:28:07	5ks	7.124	754	744 (98.7%)	566 (75%)	257	255 (99%)	196 (76%)
...	ALSS	13:14:00	+31:30:00	...	5.800	34	34 (100%)	33 (97%)
...	AMSS	81.77	109	109 (100%)	107 (98%)
...	ROSAT	20391	221	221 (100%)	221 (100%)
Total						4351	4253	2503	2957	2902	1641

Column (1) name of the field; (2) name of X-ray survey programme within this field; (3,4) approximate center of the survey; (5) nominal X-ray exposure time; (6) total area covered by both the X-ray data and the optical or infrared imaging, excluding areas close to bright stars; (7) number of X-ray sources detected in the soft (0.5–2 keV) energy band; (8) number (and fraction) of the soft X-ray sources that are associated with a robust multiwavelength counterpart; (9) number (and fraction) of the soft X-ray sources that have a spectroscopic redshift; (10) number of X-ray sources detected in the hard (2–7 keV) energy band; (11) number (and fraction) of the hard X-ray sources that are associated with a robust multiwavelength counterpart; (12) number (and fraction) of the hard X-ray sources that have a spectroscopic redshift.

that maximises the sum of the completeness and reliability (see Luo et al. 2010). A candidate counterpart that exceeded this LR threshold was deemed a secure counterpart (taking the counterpart with the highest LR value in cases of > 1 secure candidate). We then repeated the entire LR matching process for the bands indicated in the table in Appendix A, retaining any additional secure counterparts identified in these bands. This procedure enables us to identify secure counterparts for a high fraction ($> 90\%$) of the X-ray sources in the CDFS, CDFN and EGS fields (see Table 1). The vast majority ($\sim 92\%$) of the secure counterparts were identified in the (deblended) IRAC 3.6 μm catalogue. Matching to the additional bands allows us to identify counterparts when the IRAC candidate is faint, blended, or non-existent. No additional cross-matching is required as the full multiband photometry is provided through matched apertures for all sources in the Rainbow database.

2.2.2 COSMOS

In the COSMOS field, which is not currently included in the Rainbow surveys database, we matched directly between our X-ray source lists and two multiwavelength catalogues: 1) the COSMOS Intermediate and Broad Band Photometry Catalogue 2008², which is based on detection in the deep Subaru i^+ imaging of the entire COSMOS field (Capak et al. 2007); and 2) the S-COSMOS IRAC 3.6 μm based catalogue (Sanders et al. 2007). Unlike the Rainbow catalogues, the S-COSMOS IRAC 3.6 μm catalogue has not been deblended. Thus, we first identified secure matches (using the LR method) from the higher-resolution Subaru i^+ catalogues. We found secure counterparts for 1348 of the 1621 X-ray sources (83 per cent) in our overall C-COSMOS catalogue. Next, we applied the LR method to match between the X-ray and S-COSMOS catalogues. We identified secure counterparts for an additional 178 sources, taking our overall completeness to 94 per cent.

Finally, we cross-matched between our secure positions and the original catalogues, again using the LR method to ensure only secure associations are considered. For 1499 of the 1526 secure counterparts, we end up with both a Subaru i^+ and S-COSMOS counterpart. For these sources we obtain photometry in up to 18 broad-band filters spanning from the UV to near-IR, as well as 15 intermediate- or narrow-band optical filters, from the Subaru i^+ catalogue (see Ilbert et al. 2009; McCracken et al. 2010). The photometry was extracted in 3'' diameter apertures from PSF-matched images. We also adopt IRAC photometry in the 3.6, 4.5, 5.8 and 8.0 μm bands from the S-COSMOS catalogue, applying the aperture corrections given in Ilbert et al. (2009).

In 27 cases we have an i^+ counterpart but do not find an S-COSMOS counterpart, either due to the limited depth of the IRAC imaging or because the i^+ source is blended at the IRAC resolution. For these sources we simply ignore the IRAC bands in our photometric redshift estimates (see Section 2.6 below).

In 23 cases, we identify an S-COSMOS counterpart, but no Subaru i^+ source. For these sources we extracted photometry in 3'' diameter apertures at the S-COSMOS position in the PSF-matched images for all the UV to near-IR bands, generally obtaining only upper limits for these optically faint sources.

2.2.3 Bootes

For the Bootes field, we compiled multiwavelength catalogues from the NOAO Deep-Wide Field Survey (Jannuzi & Dey 1999) DR3, SDSS DR9 (Ahn et al. 2012), GALEX GR7³, FLAMINGOS Extragalactic Survey (Elston et al. 2006), and the Spitzer Deep-Wide Field Survey (Ashby et al. 2009). We use the LR method to match our X-ray catalogues to the appropriate selection band for each of these surveys, assigning secure matches from the surveys in the order of priority indicated in the table in Appendix A. The vast

² <http://irsa.ipac.caltech.edu/data/COSMOS/datasets.html>

³ <http://galex.stsci.edu/GR6/>

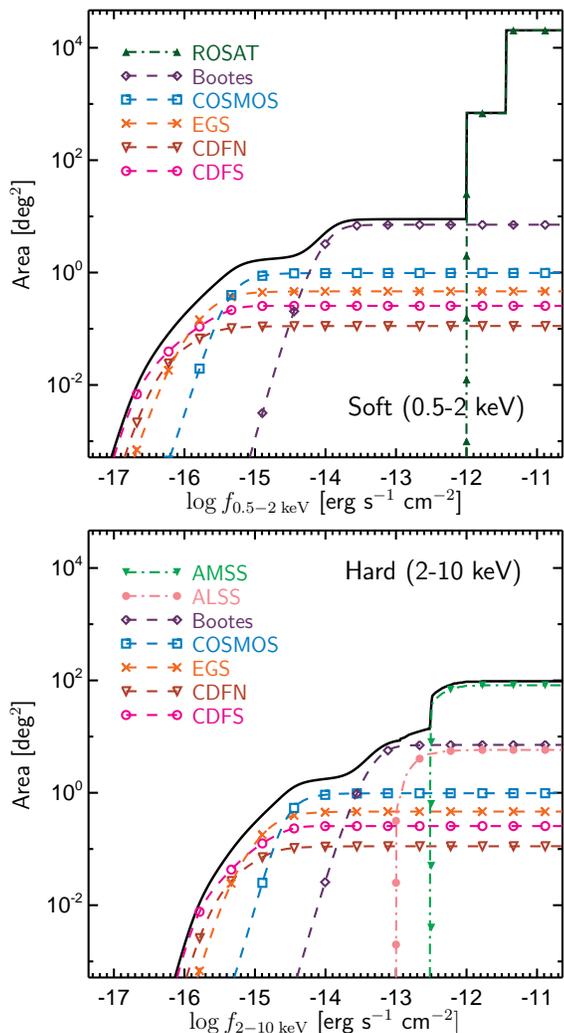


Figure 1. X-ray area curves (sensitive area versus X-ray flux) for each of our surveys in the soft and hard bands. The black line indicates the total area curve for our study. We limit our sensitivity analysis to the area falling within the footprint of our multiwavelength photometry, excluding areas close to bright stars. The ROSAT area curve includes the ROSAT bright survey and SA–N survey (see Section 2.3). The EGS area curve includes both the AEGIS-XD 800ks data and the additional area at 200ks from AEGIS-XW. The CDFS area curve combines the CDFS 4Ms data with the flanking 250ks data from the E-CDFS. These area curves assume a single, fixed conversion factor between the X-ray count rate and flux, corresponding to an X-ray spectrum with a photon index $\Gamma = 1.4$ and Galactic absorption only; the effect of different X-ray spectral shapes—and the resulting uncertainties in the conversion factors—are accounted for in our Bayesian methodology described in Section 3.

majority of our secure matches are identified in the NDWFS *I* band. We identify secure counterparts for 95.8 per cent of our X-ray sources. Finally, we cross-matched between the original catalogues and our secure counterpart positions, again applying the LR method, to identify common sources. The combined surveys provide photometry in up to 17 different bands (see Appendix A).

2.3 Large-area surveys

Combining our five *Chandra* fields provides a sample of over 4000 soft band detections and over 2800 hard band detections from

a total area of $\sim 9 \text{ deg}^2$. However, to accurately constrain the bright end of the XLF requires samples of higher luminosity X-ray sources identified from larger-area surveys. We thus supplement our sample with sources from large-area surveys carried out with *ASCA* (in the hard band) and *ROSAT* (in the soft band).

For our large-area hard-band sample, we include the 34 sources from the *ASCA* Large Sky Survey (ALSS: Ueda et al. 1999), which covers a contiguous area of 5.8 deg^2 near the north Galactic pole. We adopt the optical identifications from Akiyama et al. (2000): 2 sources are optically identified as galaxy clusters, 1 is a star, 1 source remains unidentified and the remaining 30 are associated with AGNs, all of which have spectroscopic redshifts. We also select sources from the *ASCA* Medium Sensitivity Survey (AMSS, Ueda et al. 2001), which combines data from a large number of *ASCA* observations at high Galactic latitudes over an area $\sim 82 \text{ deg}^2$. We include sources from the AMSSn subsample, selected in the hard (2–10 keV) band, with optical identifications presented by Akiyama et al. (2003). The sample includes 87 X-ray sources and has 100 per cent spectroscopic completeness. We include additional sources with from the AMSSs subsample (Ueda & Akiyama, private communication), which includes 20 AGN; 2 sources in this sample remain unidentified. We adopt area curves from the ALSS and AMSS from Akiyama et al. (2003) and Ueda et al. (2003) respectively.

For our soft band sample we combine samples from the *ROSAT* bright survey (Fischer et al. 1998; Schwöpe et al. 2000) and the Selected-Area–North survey (SA–N: Appenzeller et al. 1998), removing duplicate sources that appear in both surveys. We include sources with significant detections in the 0.5–2 keV band and adopt the unabsorbed flux estimates (corrected for Galactic absorption). We cut our sample at flux limits of $f_{0.5-2\text{keV}} > 3.6 \times 10^{-12} \text{ erg s}^{-1} \text{ cm}^{-2}$ for the RBS sample and $f_{0.5-2\text{keV}} > 1 \times 10^{-12} \text{ erg s}^{-1} \text{ cm}^{-2}$ for the SA–N sample. These high flux limits ensure our sources all lie well above the sensitivity limits of the surveys, allowing us to adopt simple sensitivity curves that correspond to the entire area of each survey and cut off sharply at each of the flux limits (see Figure 1). All the sources above these flux limits have spectroscopic classifications and we identify a total of 221 AGNs with spectroscopic redshifts (excluding BLLac type objects).

We note that, in contrast to our *Chandra* fields, our area curves for our large-area surveys do not account for the Poisson nature of the detection. As we restrict the samples from the large-area surveys to highly significant detections, this simplification will have a minimal effect on our XLF measurements. However, differences in the assumed spectral shape can have a significant impact on the estimate of a flux and thus the assumed sensitivity. Uncertainties in the spectral shape and the resulting differences in sensitivity are accounted for by our Bayesian methodology described below.

2.4 Identification and masking of stars

Bright stars can contaminate our photometry, leading to issues with counterpart identification and photometric redshift estimates in these regions. We have therefore masked out areas close to bright stars from all of our *Chandra* fields in a consistent manner. We searched for stars brighter than $V = 15$ in the *HST* Guide Star Catalog 2.3 (Lasker et al. 2008). We masked all areas within a radius, r , given by

$$r = (16 - V) \times 6'' \quad (1)$$

where V is the V -band magnitude from the Guide Star Catalog. We set a maximum masking radius of $40''$. We have removed any X-ray sources within this radius from our samples (which can include the star itself or a nearby source). We also excluded these regions when calculating the X-ray sensitivity and area curves.

We have also identified stars with fainter magnitudes that are detected at X-ray wavelengths and removed them from our samples. In the fields with Rainbow coverage (CDFs, CDFN, EGS) stars were identified by a range of colour and morphology criteria, as described in Barro et al. (2011a). In the COSMOS and Bootes fields we applied a single colour criterion based on the region of colour-colour space occupied by stars in Ilbert et al. (2009),

$$R - [3.6] < 3.0 \times (R - I) - 1.2 \quad (2)$$

where $[3.6]$ is the magnitude in the IRAC $3.6\mu\text{m}$ imaging, R is the magnitude in the Subaru $r+$ or NDWFS R filter, and I is the magnitude in the Subaru $i+$ or NDWFS I filter. For all fields, we also required that the X-ray sources exhibit a low X-ray-to-optical flux ratio, $\log f_X/f_{\text{opt}} < -1$, where the ratio is calculated as

$$\log \frac{f_X}{f_{\text{opt}}} = \log f_{0.5-2\text{keV}} + 5.4 + \frac{I}{2.5}. \quad (3)$$

This cut ensures that we do not exclude bright QSOs from our sample that may satisfy the other stellar criteria. When a spectroscopic classification is available (see Section 2.5) this over-rides our photometric classification. Of the 50 spectroscopically classified stars in our five *Chandra* fields, 40 (80 per cent) were also identified by our photometric procedure.

2.5 Spectroscopic redshifts

All of our *Chandra* surveys have been the subject of intense spectroscopic campaigns. These campaigns include large-scale follow-up of the general galaxy population, in addition to those directly targeting X-ray sources.

In the CDF-S we first searched for spectroscopic redshifts in the catalogue of Xue et al. (2011). The majority of the spectroscopic redshifts for our CDFs-4Ms survey are taken from this catalogue, as well as some of our ECDFS survey (in the area that overlaps with the 4Ms data). We also searched for spectroscopic redshifts from the Arizona CDFs Environment Survey (Cooper et al. 2012b), the spectroscopic sub-sample of sources from the MUSYC sample (Cardamone et al. 2010b), and PRIMUS (Coil et al. 2011).

In the CDF-N we used spectroscopic redshifts from DEEP3 (Cooper et al. 2011) as well as the surveys of Trouille et al. (2008); Barger et al. (2008); Reddy et al. (2006); Wirth et al. (2004); Cowie et al. (2004) and Steidel et al. (2003).

In the EGS we compiled spectroscopic redshifts from a number of surveys including DEEP2, DEEP3, the Canada-France Redshift Survey and *MMT* follow-up of X-ray sources. See ? and references therein for full details.

In COSMOS we initially searched for spectroscopic redshifts of X-ray sources in the C-COSMOS catalogue of Civano et al. (2012). We also search for additional spectroscopic redshifts from the bright zCOSMOS survey catalogue (Lilly et al. 2009) and PRIMUS (Coil et al. 2011).

In Bootes we adopt spectroscopic redshifts from the AGN and Galaxy Evolution Survey (Kochanek et al. 2012). For all fields we also searched for spectroscopic redshifts from the Sloan Digital Sky Survey (SDSS: York et al. 2000).

We matched the spectroscopic catalogues to the secure counterparts of our X-ray sources using a $2''$ search radius, corrected for

any overall astrometric offset, and repeated the matching with a $1''$ radius. We only adopt those spectroscopic redshifts that are flagged as high-quality, reliable redshifts in the original catalogues. Table 1 gives the number of X-ray sources in our hard- and soft-band selected samples with reliable spectroscopic redshifts.

2.6 Photometric redshifts

The levels of spectroscopic completeness vary over our *Chandra* surveys from ~ 35 per cent (in our ECDFS area) to ~ 75 per cent (in the Bootes field). For the remaining X-ray sources we must resort to photometric redshift estimates, which are determined by fitting a set of template spectra to the observed spectral energy distributions (SEDs) of our sources. Such redshifts can be highly uncertain, particularly when considering faint X-ray sources. A key advantage of our Bayesian analysis (see Section 3) is that we are able to account for the uncertainties in photometric redshifts by adopting probability distributions for the redshift, $p(z)$, rather than a single redshift estimate. We thus require that our photo- z approach recovers a $p(z)$ distribution that accurately reflects the uncertainties in our redshift estimates.

We calculate photometric redshifts using the *EaZY* photo- z code (Brammer et al. 2008). We use *EaZY* in two-template mode, allowing for combinations of a galaxy and AGN template. For the galaxy templates we adopt the ‘‘pegase13’’ template set provided with *EaZY*. This template set consists of 259 synthetic galaxy templates drawn from the range of parameters described by Grazian et al. (2006). Additional ‘‘star-forming and dusty’’ templates are included by applying the Calzetti et al. (2000) reddening law for a range of different extinctions to a subset of the galaxy templates. We adopt seven AGN templates from the Salvato et al. (2009) template set⁴, namely the Sey 1.8, Sey 2, Mrk231, pl_TQSO1, pl_QSOH, pl_QSO and the S0-10_QSO2-90 hybrid template. We also include the Type-2 ‘‘Torus’’ template from the Polletta et al. (2007) library⁵, which is not included in the final Salvato et al. (2009) set. We note that some of these templates will include host galaxy contributions (particularly the Sey 1.8, Sey 2, Mrk231 and S0-10_QSO2-90 templates). This is not a major issue as we want to determine the distribution of possible redshifts, rather than perform an accurate host-AGN decomposition. The two-template mode in *EaZY* allows any possible combination of one of our AGN templates and one of our galaxy templates and thus allows for a large amount of flexibility in the fitted template SEDs.

Another advantage of the *EaZY* photo- z code is that it allows for the inclusion of a ‘‘template error function’’. This feature accounts for additional uncertainty in the template SED as a function of (rest-frame) wavelength and thus allows for the fact that our template set may not accurately represent the true diversity of SED shapes. This uncertainty in the true range of template SEDs is particularly useful as our observed SEDs extend into the UV and mid-IR, where the templates are poorly calibrated, especially for AGNs. In addition, as the optical emission from an (unobscured) AGN can vary on timescales of months-to-years, our observed SEDs may not be well-matched by a single underlying template. The template error function can also allow for any overall calibration uncertainties in the diverse sets of photometric observations used to construct

⁴ http://www.mpe.mpg.de/~mara/PHOTOZ_XCOSMOS/

⁵ http://www.iasf-milano.inaf.it/~polletta/templates/swire_templates.html

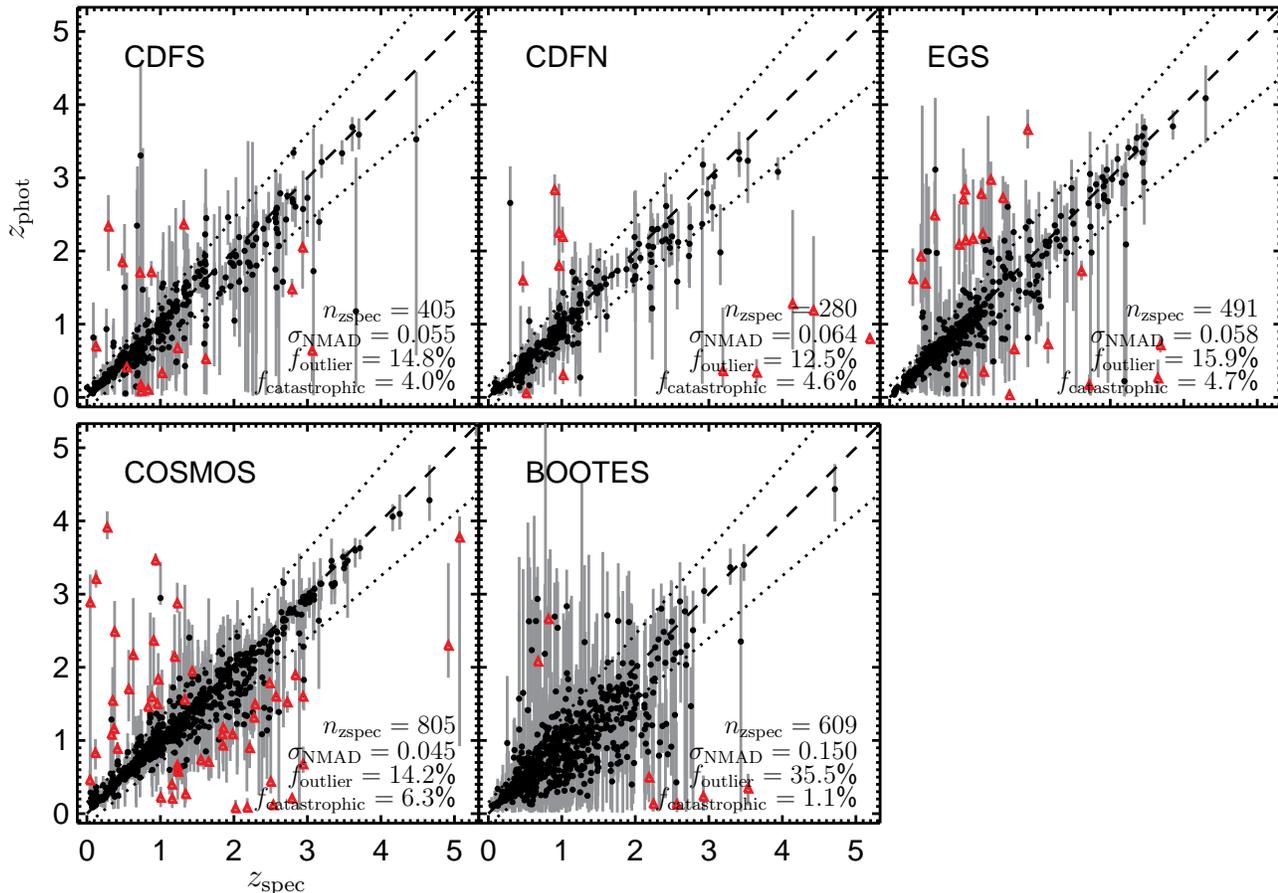


Figure 2. Comparison of photometric redshifts (z_{phot}) and spectroscopic redshifts (z_{spec}) for our five *Chandra* fields. We plot the mean of the $p(z)$ distribution as the best estimate of the photometric redshifts (black circles); error bars indicate the 95 per cent central confidence interval. The dashed line indicates a 1:1 relation, whereas the dotted lines correspond to $\Delta z/(1 + z_{\text{spec}}) = \pm 0.15$ (sources where the best estimate of the photo- z lies outside this range are flagged as outliers). In the legend for each panel we give the number of X-ray sources with reliable spectroscopic redshifts ($n_{z_{\text{spec}}}$), the accuracy based on the normalized median absolute deviation (σ_{NMAD}), the fraction of outliers (f_{outlier}) and the fraction of catastrophic failures ($f_{\text{catastrophic}}$); see text for details. We highlight catastrophic failures with a red triangle. Over four of our fields (CDFS, CDFN, EGS and COSMOS) we obtain a consistent accuracy of $\sigma_{\text{NMAD}} \approx 0.05$, with ~ 15 per cent outliers and ~ 5 per cent catastrophic failures. In the Bootes field—our largest area, shallowest field—we have poorer accuracy and a higher outlier fraction, reflecting the more limited photometric imaging in this field. While our photo- z have a poorer accuracy (σ_{NMAD}) and a higher outlier rate than in some prior works, we choose to adopt our estimates as they have been calculated in a consistent manner across all five of our *Chandra* fields and have representative errors that we can fully track via the $p(z)$ in our Bayesian methodology (see text for further discussion).

our observed SEDs. We derive a template error function on a field-by-field basis—to ensure that it represents the calibration uncertainties in a given data set—using the basic procedure laid out in Brammer et al. (2008). First, we attempt to fit the observed SEDs with our templates, fixing the redshift at the spectroscopic value, where available. We then calculate

$$\Delta f_j = \frac{F_j - T_j}{F_j} \quad (4)$$

where F_j indicates the observed flux in a filter for source j , and T_j is the flux from the best-fitting template. We calculate Δf_j as function of rest-frame wavelength for all sources and filters and take the median of every 400 individual measurements across the rest-frame wavelength range. We subtract the median photometric error, in quadrature, to estimate the contribution from “template error” to the uncertainty as a function of rest-frame wavelength. The value of the template error is typically around 10 per cent (in

flux) but varies between ~ 4 per cent and ~ 20 per cent depending on the wavelength and the data in a given field.

In Figure 2 we compare our photo- z estimates for X-ray sources to secure spectroscopic redshifts across our 5 *Chandra* fields. The panels for the EGS and CDFS fields combine the deep survey areas (AEGIS-XD, CDFS-4Ms), where the best photometry is available, with the larger-area shallow surveys (AEGIS-XW, ECDFS). We include all X-ray sources with a high-quality spectroscopic redshift in these plots and do not apply any cuts based on the estimated quality of the photo- z . Thus, we include sources with extremely broad $p(z)$ distributions, which are often flagged as unreliable and excluded when assessing the success of photo- z techniques.

In each panel of Figure 2 we present a number of summary statistics:

- (i) σ_{NMAD} : the accuracy based on the normalized median absolute deviation between the best photo- z and

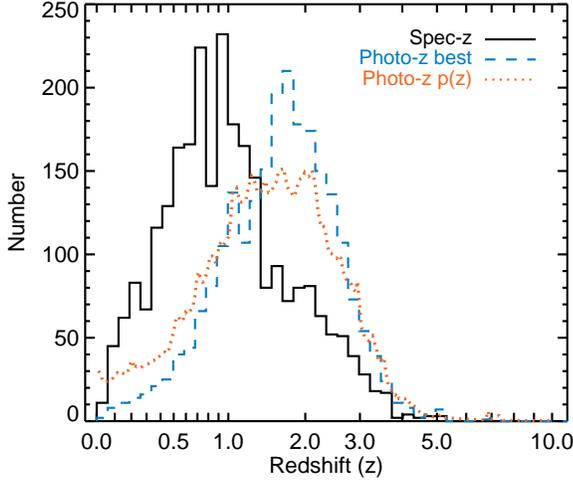


Figure 3. Redshift distribution of X-ray sources in our five *Chandra* fields with high-quality spectroscopic redshifts (black solid line) and those where we adopt photometric redshifts. The blue dashed line indicates the distribution of the “best” photo- z estimates (mean of the $p(z)$ distribution), whereas the orange dotted line shows the distribution obtained by combining the individual $p(z)$ distributions.

the spectroscopic value, defined as $\sigma_{\text{NMAD}} = 1.48 \times \text{median}(|z_{\text{phot}} - z_{\text{spec}}|/(1 + z_{\text{spec}}))$.

- (ii) f_{outlier} : the fraction of outliers, defined as the fraction of sources where $|z_{\text{phot}} - z_{\text{spec}}|/(1 + z_{\text{spec}}) > 0.15$;
- (iii) $f_{\text{catastrophic}}$: the catastrophic outlier rate, calculated as the fraction of sources where less than 5 per cent of the integrated $p(z)$ lies within $-0.15 < (z - z_{\text{spec}})/z_{\text{spec}} < +0.15$.

Over four of our fields (CDFs, CDFN, EGS and COSMOS) we obtain a consistent accuracy of $\sigma_{\text{NMAD}} \approx 0.06$, with ~ 15 per cent outliers. Our approach ensures we assign an appropriate uncertainty, traced by the $p(z)$, to the bulk of our sources and only ~ 5 per cent of sources are thus flagged as catastrophic failures. In the Bootes field—our largest area, shallowest field—we have poorer accuracy and a much higher outlier fraction, reflecting the more limited photometric imaging in this field. However, the catastrophic outlier fraction is only ~ 1 per cent, indicating that this additional uncertainty is represented by our $p(z)$ distributions. We also note that in this field we have the highest spectroscopic completeness (~ 75 per cent) and so only use photometric redshifts for a relatively small fraction of the sources in our eventual analysis of the XLF; when we do resort to a photometric redshift we account for the large uncertainty in the redshift.

In Figure 3 we plot the distribution of redshift estimates for sources in our *Chandra* fields where we have a spectroscopic redshift and those where we adopt the photometric redshift information. Generally the sources with photometric redshifts lie at higher redshifts, a consequence of spectroscopic follow-up programmes being biased towards optically bright sources. We also show the integrated contribution from the $p(z)$ of all the photo- z sources. The integrated $p(z)$ is skewed towards lower redshifts. This skew is due to the possibility that many potential high-redshift sources could actually lie at lower redshifts, which is reflected by their $p(z)$ and must be accounted for in measurements of the XLF.

A small fraction of our X-ray sources (< 2 per cent) lack a multiwavelength counterpart, precluding a photometric redshift es-

timate. We retain these sources in our analysis, ensuring completeness of our sample, but adopt a $p(z)$ that is constant in $\log(1 + z)$ over our allowed redshift range ($0 < z < 10$). This reflects our lack of *a priori* knowledge of the redshift; the X-ray flux information is retained and thus *a posteriori* (after folding the constant $p(z)$ through the final XLF) there may be a preferred redshift solution. The lack of a multiwavelength counterpart could imply that a high redshift solution should be given higher *a priori* preference for such sources, thus our approach is conservative. There is also a possibility that these X-ray sources lack counterparts as they are spurious detections, corresponding to positive fluctuations in the background count rate. Our analysis accounts for the Poisson nature of the X-ray detection and thus allows for this possibility.

Other estimates of photometric redshifts are available for X-ray sources in many of our fields (e.g. Barger et al. 2003; Cardamone et al. 2010a; Hsu et al. 2014). Many of these studies take additional steps to improve the quality of the photo- z estimates. These steps can include optimizing the template set (e.g. Luo et al. 2010), attempting to correct the observed photometry for variability (e.g. Salvato et al. 2009), or applying priors based on the source morphology and X-ray flux (e.g. Salvato et al. 2011). These studies often achieve a higher accuracy and lower outlier rate than our own photo- z analysis. However, these additional steps can lead to underestimates of the true uncertainties in the photo- z . Conversely, we retain a large set of possible templates to ensure we produce $p(z)$ distributions that account for the large uncertainties in the redshift and template degeneracies. While we have higher outlier rates, our fraction of catastrophic failures remains $\lesssim 5$ per cent, indicating that our $p(z)$ distributions are representing the uncertainties. The nominal errors (i.e. the 68 per cent confidence intervals) on our photo- z are also comparable to the residuals between our best photo- z estimate and the available spectroscopic redshifts, in contrast to most previous work where errors may be underestimated by a factor $\sim 2 - 6$ (e.g. Luo et al. 2010; Hsu et al. 2014). Furthermore, we require the full $p(z)$ distribution, which most prior studies do not provide. We thus choose to use our own photometric redshifts: our poorer accuracy and higher outlier rates are accounted for and compensated by our Bayesian analysis that incorporates the full $p(z)$ information.

3 BAYESIAN METHODOLOGY

In this paper we expand on the Bayesian methodology developed by A10, incorporating the distribution of X-ray absorption properties and accounting for the effects on the inferred shape and evolution of the XLF. Our method also accounts for the uncertainty in the measured X-ray flux (due to photon counting statistics), uncertainties in the redshift (for sources with photometric redshifts or no counterparts), uncertainties in the X-ray spectral shape, and the resulting uncertainty in the X-ray luminosity for an individual source. Our methodology is described below.

3.1 Probability distribution function for a single source

For a single source in either our hard- or soft-band sample we can derive the probability distribution function for z , L_X and N_H based on our observed data for that source alone, which is given by $p(z, L_X, N_H | D_i)$ where D_i indicates the observed data from source i in our sample. This function is normalized such that

$$\int dz \int d \log L_X \int d \log N_H p(z, L_X, N_H | D_i) = 1. \quad (5)$$

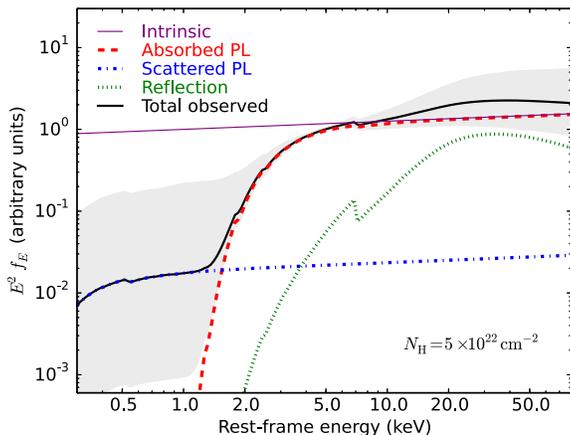


Figure 4. Our assumed X-ray spectral model for an AGN. We assume the intrinsic X-ray spectrum is a power law with photon index $\Gamma = 1.9 \pm 0.2$ (purple line). The observed spectrum is absorbed by the intervening column density, fixed at $N_{\text{H}} = 5 \times 10^{22} \text{ cm}^{-2}$ for this example (red dashed line). We also allow for a fraction ($f_{\text{scatt}} \approx 2$ per cent) of the intrinsic power-law to be scattered, unabsorbed, into the line-of-sight (blue dot-dashed line). In addition, we allow for a component from Compton-reflection from cold, optically thick matter (such as a torus or accretion disc) that leads to the characteristic hump at high energies (green dotted line). The black line indicates the total observed spectrum, while the grey region indicates the 95 per cent confidence interval on the spectrum, allowing for the range of possible spectral parameters (see Section 3.1, Table 2).

We can re-write the probability distribution function as

$$p(z, L_X, N_{\text{H}} | D_i) = p(z | d_i) p(L_X, N_{\text{H}} | z, T_i, b_i) \quad (6)$$

where d_i indicates the multiwavelength data used to estimate the redshift and T_i and b_i correspond to the X-ray data for this source: the total observed X-ray counts in the given band and the estimated background respectively.

Our knowledge of z is based on either a spectroscopic redshift, in which case we assume $p(z | d_i)$ is described by a δ -function at the spectroscopic value, or a photometric redshift, when $p(z | d_i)$ is given by the $p(z)$ from our photometric redshift fitting described in Section 2.6 above. For the small fraction of sources where we were unable to identify a multiwavelength counterpart—and thus have no redshift information—we adopt a $p(z)$ distribution with a constant density in $\log(1+z)$ over $0 < z < 10$.

The observed X-ray data can be described by a Poisson process. Thus, the likelihood of observing T_i counts from a source is given by

$$\mathcal{L}(T_i | c_i, b_i) = \frac{(c_i + b_i)^{T_i}}{T_i!} e^{-(c_i + b_i)} \quad (7)$$

where c_i is the X-ray count rate from source i in the observed energy band and we have assumed that the expected background count rate, b_i , is well determined. The uncertainty in c_i is fully described by the Poisson likelihood given in Equation 7. However, to convert from a count rate, c_i , to an estimate of L_X and N_{H} (given z) we must assume a model for the X-ray spectral shape and fold this model through the appropriate instrumental response.

Figure 4 shows an example of the X-ray spectral model used in this work. We assume the intrinsic X-ray continuum is described by a power law with photon index Γ and a high energy cut-off (we fix the folding energy at 300 keV, although this has a negligible im-

pact on the lower energies we observe). The observed spectrum is attenuated along the line-of-sight by the intervening column density, N_{H} . We include both photoelectric absorption (using the *wabs* model in XSPEC) and Compton-scattering (via *cabs*), which suppresses the continuum further for Compton-thick column densities. A fraction of this power-law, f_{scatt} , is also allowed to emerge as an unabsorbed component that is thought to be scattered into the line-of-sight by ionized gas in the vicinity of the AGN. We also include a contribution due to Compton reflection from cold, optically thick matter, which gives rise to a characteristic “hump” in the X-ray spectrum at ~ 30 keV. Such a reflection component is expected due to reprocessing of the primary X-ray emission by a surrounding, dusty torus, and/or an accretion disc. We adopt the *pexrav* model (Magdziarz & Zdziarski 1995), which is based on Monte-Carlo simulations of Compton reflection from a cold, optically thick slab of material. We assume the incident power-law has the same shape as the directly transmitted component. We fix the inclination angle to 30° as a representative value and allow the intensity to be set by the normalization, R , relative to that expected from a slab subtending a solid angle of 2π (which is allowed to vary between 0 and 2, spanning the extreme cases of no reflection up to an effective 4π solid angle coverage). We absorb the reflection component by the same column density seen by the primary emission. This is a good assumption when the reflection arises from the accretion disc and also provides reasonable agreement with the shape and intensity of the reflection component based on sophisticated models of toroidal obscurers (e.g Brightman & Nandra 2011a). All components are subjected to Galactic absorption, with column densities determined from Dickey & Lockman (1990) via the HEASOFT N_{H} tool. In XSPEC terminology, our model is described by

$$\text{wabs} * [(1 - \text{constant}) * \text{zwabs} * \text{cabs} * \text{zpowerlw} * \text{zhigcut} + \text{constant} * \text{zpowerlw} + \text{zwabs} * \text{pexrav}] \quad (8)$$

where the constant corresponds to the scattered fraction, f_{scatt} . We note that more physically motivated models could be adopted to describe the X-ray spectrum, self-consistently modeling the emission, reflection and absorption due to the accretion disc or torus (e.g. Ross & Fabian 2005; Brightman & Nandra 2011a). However, Buchner et al. (2014) found that more simplistic models such as ours are generally sufficient to reproduce the observed spectral shape of individual, distant AGNs, especially considering our analysis uses broad-band fluxes rather than performing a detailed X-ray spectral analysis.

To describe our X-ray spectral model requires three additional parameters—the photon index, Γ , the scattered fraction, f_{scatt} , and the relative normalization of the reflection component, R —which we introduce as “nuisance” parameters (collectively designated by ξ) in our Bayesian analysis. For a given set of spectral parameters (ξ), z , L_X and N_{H} , we can determine the expected count rate, c_i , and thus link the probability distribution function for L_X , N_{H} and ξ to the Poisson likelihood given in Equation 7 above. Thus,

$$p(L_X, N_{\text{H}}, \xi | z, T_i, b_i) \propto \mathcal{L}(T_i | c_i(z, L_X, N_{\text{H}}, \xi), b_i) \pi(\xi) \quad (9)$$

where $c_i(z, L_X, N_{\text{H}}, \xi)$ is the expected count rate for a source with redshift z , luminosity L_X , absorption column N_{H} , and additional spectral parameters ξ based on our X-ray spectral model, folded through the appropriate instrumental response for source i .

A priori, the spectral parameters ξ for a given source are not well known; $\pi(\xi)$ denotes the prior distribution that we adopt for

Table 2. Priors on the spectral parameters for an individual X-ray AGN, ξ .

Parameter	Prior type	Prior specification
Γ	Gaussian	1.9 ± 0.2^a
$\log f_{\text{scatt}}$	lognormal	-1.73 ± 0.8^b
R	constant	0 – 2

^aBased on observed distribution from Nandra et al. (2007).

^bBased on observed distribution from Winter et al. (2009).

our spectral parameters, which describes the range of possible values and thus encapsulates the uncertainty in X-ray spectral shape. We assume the photon index, Γ , is drawn from a Gaussian distribution with a mean of 1.9 and standard deviation of 0.2 (corresponding to the observed distribution of intrinsic photon indices in X-ray spectral studies of nearby AGNs, e.g. Nandra et al. 2007). We assume the scattered fraction, f_{scatt} , is drawn from a lognormal distribution with mean of 2 per cent and scatter of 0.8 dex based on the observed distribution of partial covering factors of sources in the *Swift*/BAT sample from Winter et al. (2009). For the reflection strength, R , we assume a uniform distribution in the range $0 < R < 2$. We thus allow for a large uncertainty in the strength of the reflection component, which is reasonable to encapsulate uncertainties in the geometry of the accretion disc and/or torus with our simplified modeling of the reflection. Table 2 summarizes this prior information. The coloured lines in Figure 4 show each component evaluated at the prior mean of the spectral parameters and the black line corresponds to the total observed spectrum for these values. The grey region indicates the 95 per cent confidence interval on the total observed spectrum, adopting our prior distributions for the spectral parameters. We note that the large uncertainties in the scattered fraction, f_{scatt} , lead to large uncertainties in the spectrum at softer energies for moderately and heavily absorbed sources and thus leads to large uncertainties in the intrinsic L_X .

To obtain the probability distribution function for L_X and N_H only (for a given z and our observed X-ray data), we must marginalize over the nuisance spectral parameters. Thus,

$$p(L_X, N_H | z, T_i, b_i) = \int d\xi p(L_X, N_H, \xi | z, T_i, b_i). \quad (10)$$

In Figure 5 we show an example of $p(L_X, N_H | z, T_i, b_i)$ for a source detected in the hard band (left) and a source detected in the soft band (right) with *Chandra*. In both cases we fix $z = 1.0$ and assume $T_i = 30$ total observed counts with an expected background of $b_i = 5.0$ counts. We restrict the possible column densities to $20 < \log N_H (\text{cm}^{-2}) < 26$. The shading and contours indicate the range of possible values of L_X and N_H for each source. Detection in a single, broad energy band does not place constraints on the value of N_H , but we are able to place constraints on the range of possible values of L_X and how this depends on N_H . For the hard band example, we can constrain L_X to within ~ 0.25 dex, provided the absorption column is $N_H \lesssim 10^{23} \text{ cm}^{-2}$. If the column density is higher, then the same observed counts must correspond to a higher value of L_X . For a source detected in the soft band, absorption effects are apparent for lower column densities ($N_H \gtrsim 10^{22} \text{ cm}^{-2}$). At column densities $N_H \gtrsim 2 \times 10^{23} \text{ cm}^{-2}$ (at $z = 1.0$) the observed flux in the 0.5–2keV band is dominated by the scattered component only; the intrinsic luminosity is poorly constrained but must be a factor ~ 30 greater than if the source had a lower N_H .

We note that the effect of absorption on the observed counts in the hard or soft energy bands will vary significantly with red-

shift. For example, at $z \sim 3$, where the observed 0.5–2 keV energy band probes rest-frame energies $\sim 2 - 8 \text{ keV}$, the inferred L_X is only affected for column densities $\gtrsim 10^{23} \text{ cm}^{-2}$. These redshift-dependent effects are fully accounted for in our calculation of the expected counts, $c_i(z, L_X, N_H, \xi)$, using our X-ray spectral model.

For the analysis in this paper, we treat each source in our hard- or soft-band samples as independent detections. Thus, the only X-ray data we use for an individual source is the detected counts (and the expected background), which as described above does not allow us to constrain the value of N_H for an individual source. Instead, our approach (see Sections 3.2 below) involves determining the overall distribution of X-ray luminosities and absorption column densities (given by the XLAF) that correctly describes both our hard- and soft-band samples. Using this method, certain values of L_X and N_H may be disfavored *a posteriori* (after performing our full analysis of the overall sample). For example, higher values of L_X may be disfavored *a posteriori* based on the shape of the XLF (that generally shows higher L_X sources should be rarer). Conversely, lower values of N_H may be disfavored if our overall XLAF requires a large fraction of absorbed sources.

Nevertheless, it is worth noting that a more sophisticated analysis could place constraints on the value of N_H (or indeed any of the other spectral parameters) for an individual source *before* performing fits to the overall samples. We could combine the information on the counts in both the hard and soft bands for an individual source to place constraints on the underlying X-ray spectrum, based on this hardness ratio information (e.g. Hasinger 2008; Xue et al. 2010). Alternatively, the full X-ray spectrum could be extracted for a source and fitted with our spectral model or any other well-motivated model (e.g. Tozzi et al. 2006; Buchner et al. 2014). However, accurately incorporating this information into our Bayesian analysis—and fully propagating the uncertainties in the estimated spectral parameters, L_X and N_H —is beyond the scope of this paper.

3.2 The likelihood function for the overall sample

Here, we describe how we combine the individual probability distribution functions for each source in our sample with a given model of the XLAF to construct the likelihood function for our overall sample.

Following A10 (see also Loredi 2004), we assume our individual sources are effectively Poisson points drawn from a distribution described by the XLAF and the overall sample selection function. The expected number of sources in our hard-band sample is given by

$$\mathcal{N}_{\text{hard}}(\Theta) = \int d \log L_X \int d \log N_H \int \frac{dV}{dz} dz \int d\xi \quad (11)$$

$$\left[\psi(L_X, z, N_H | \Theta) \pi(\xi) \sum_{j=1}^{N_{\text{fields}}} A_j(L_X, z, N_H, \xi) \right]$$

where $\psi(L_X, z, N_H | \Theta)$ is the XLAF—the differential co-moving number density of AGNs per logarithmic interval in L_X and N_H —and is described by a model with a given set of parameters, Θ . We note that the XLF, $\phi(L_X, z | \Theta) \equiv \frac{d\psi(L_X, z)}{d \log L_X}$ (i.e. the differential co-moving number density of AGNs per logarithmic interval in L_X), can be recovered from the full XLAF by integrating over N_H . Thus,

$$\phi(L_X, z | \Theta) = \int d \log N_H \psi(L_X, z, N_H | \Theta). \quad (12)$$

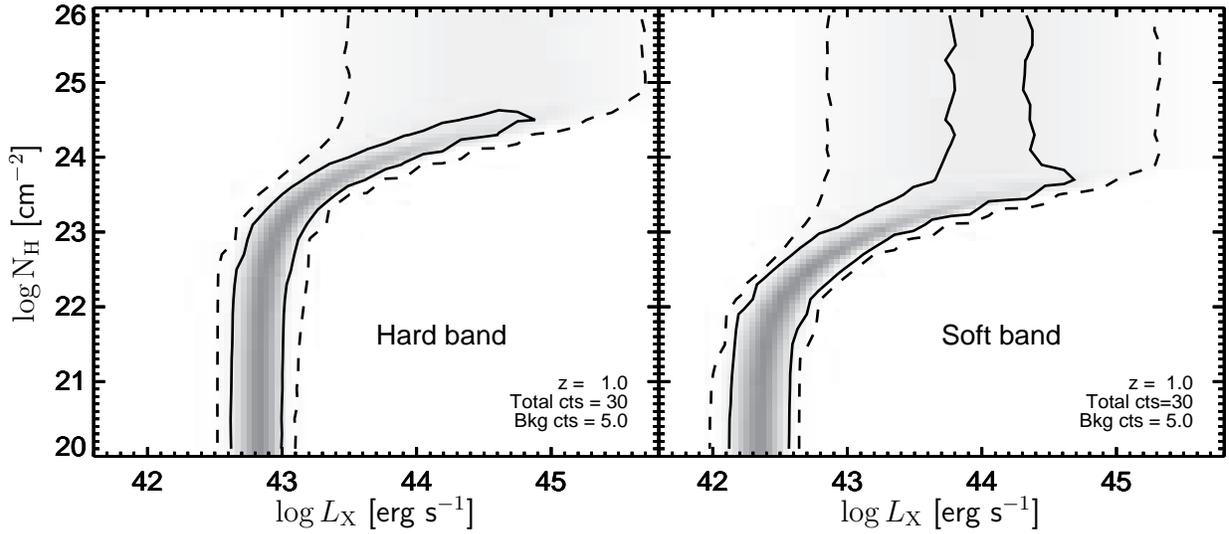


Figure 5. Example constraints on the rest-frame 2–10keV luminosity, L_X , and the line-of-sight absorption column density, N_H , for sources detected in the observed hard band (2–7keV, left) and soft band (0.5–2keV, right) in a typical deep *Chandra* observation. In both cases, we assume 30 total counts detected in the observed energy band with an expected background of 5 counts, typical of our sample in the AEGIS-XD survey region. We fix the redshift at $z = 1.0$. The grey shading indicates the probability distribution function for L_X and N_H , allowing for the Poisson uncertainty from the observed counts and assuming the X-ray spectral model shown in Figure 4, marginalized over the uncertainty in the “nuisance” spectral parameters (Γ , f_{scatt} , and R). The solid and dashed lines indicate the 68.3 per cent and 95.4 per cent (i.e. 1 and 2 σ equivalent) confidence intervals on the joint parameter space. Detection in a single band does not constrain the value of N_H but does place constraints on the range of possible values of L_X . At low values of N_H , the X-ray luminosity is constrained to within ~ 0.25 dex. However, for higher N_H values the X-ray luminosity must be correspondingly higher to produce the same number of observed counts. Absorption affects the X-ray spectrum (and thus the constraints on L_X) for lower column densities in the soft band ($N_H \gtrsim 10^{22} \text{ cm}^{-2}$) than the hard band ($N_H \gtrsim 10^{23} \text{ cm}^{-2}$) at $z = 1.0$.

We investigate and compare various different parametrizations of the XLAF (or the XLF directly), which are described in Sections 4 and 5 below.

The distribution of possible spectral parameters is given by $\pi(\xi)$, the prior on Γ , f_{scatt} and R that we describe in Section 3.1 above. $\frac{dV}{dz}$ is the differential comoving volume per unit area, as a function of z . $A_j(L_X, z, N_H, \xi)$ is the area for field j that is sensitive to a source of luminosity L_X , redshift z , absorption column N_H , and spectral parameters ξ , effectively giving the probability of a source entering our sample. To calculate $A_j(L_X, z, N_H, \xi)$ we use our X-ray spectral model described in Section 3.1 above to convert from a given L_X , z , N_H and ξ to an expected count rate, $c_i(z, L_X, N_H, \xi)$.⁶ We then convert the count rate to an “effective flux” using our standard conversion factors and determine the area sensitive to this flux from our sensitivity curves (see Section 2.1 and Figure 1). This calculation is performed for each of our N_{fields} fields and the total area is summed.

The likelihood function may now be constructed from the product of the probability distribution functions for the individual sources in the sample and the probability that no other sources were detected (A10, Loredo 2004). For a Poisson process, the probability of no detected sources (in the hard-band sample) for a given model XLAF is $e^{-N_{\text{hard}}(\Theta)}$. Thus, for the hard-band sample the

likelihood is

$$\mathcal{L}(\mathbf{D}_{\text{hard}} | \Theta) = e^{-N_{\text{hard}}(\Theta)} \times \prod_{i=1}^{n_{\text{hard}}} \left(\int d \log L_X \int d \log N_H \int \frac{dV}{dz} dz \right. \quad (13)$$

$$\left. \left[p(z, L_X, N_H | D_i) \psi(L_X, z, N_H | \theta) \right] \right)$$

where $\mathcal{L}(\mathbf{D}_{\text{hard}} | \Theta)$ is the likelihood of obtaining the observed data from the entire hard-band sample (\mathbf{D}_{hard}) for a given set of parameters (Θ) for the XLAF, and $p(z, L_X, N_H | D_i)$ is the probability distribution function for z , L_X and N_H for an individual hard-band source with observed data (D_i), marginalized over the nuisance spectral parameters (ξ). The product is taken over all n_{hard} sources in our hard-band sample.

The likelihood function for the soft-band sample can be constructed in a completely analogous manner to that described for the hard-band sample above. As we treat the hard- and soft-band samples as independent samples, the overall likelihood function for both samples is simply given by the product of the two likelihood functions. Thus,

$$\mathcal{L}(\mathbf{D}_{\text{hard}}, \mathbf{D}_{\text{soft}} | \Theta) = \mathcal{L}(\mathbf{D}_{\text{hard}} | \Theta) \mathcal{L}(\mathbf{D}_{\text{soft}} | \Theta) \quad (14)$$

where we assume the same underlying XLAF (with parameters Θ) can describe the observed data from both samples. This approach is taken in Section 5 where we combine the two samples to determine the form of the XLAF that adequately describes our two samples.

⁶ We assume an appropriate Galactic hydrogen column density, $N_{\text{H}}^{\text{Gal}}$, for each survey. For the *ROSAT* surveys and AMSS we assume the median $N_{\text{H}}^{\text{Gal}}$ of the sample.

3.3 Including the contribution of normal galaxies

So far we have implicitly assumed that all of the X-ray point sources in our two samples are actually due to the emission from AGNs, rather than an alternative process. We have already excluded stars from our sample, either by masking known bright stars or according to the criteria given in Section 2.4. However, we have so far made no attempt to exclude normal galaxies from our sample, where the X-ray emission may be predominantly due to the combined emission from X-ray binaries within the galaxy—with the total luminosity tracing the star-formation rate of the galaxy (e.g. Ranalli et al. 2003; Mineo et al. 2014)—rather than due to an AGN. The total luminosity produced by a galaxy is usually much lower than could be produced by an AGN, even in galaxies with extremely high star-formation rates. Thus, a common solution is to simply cut any source with an X-ray luminosity below $L_X \lesssim 10^{42}$ erg s $^{-1}$, and not consider sources below this luminosity any further in measurements of the XLF of AGNs (e.g. Barger et al. 2005; Silverman et al. 2008). Such a cut is fairly conservative as a high fraction of the sources below this L_X cut will be dominated by AGN emission rather than star formation. Including such sources can also improve constraints on the faint end of the XLF.

In our Bayesian analysis we do not assign a single L_X to any individual source; we have a probability distribution describing the range of possible L_X . Thus, we are unable to apply a strict luminosity cut to identify and remove normal galaxies from our sample. Instead, we account for the potential contamination of our AGN sample by including the X-ray luminosity function of normal galaxies (hereafter, the “galaxy luminosity function” or GLF) directly in our analysis. We assume the GLF is described by a Schechter function,

$$\phi_{\text{gal}}(L_X, z | \omega) d \log L_X = A \left(\frac{L_X}{L_*(z)} \right)^{-\alpha} \exp \left(-\frac{L_X}{L_*(z)} \right) d \log L_X \quad (15)$$

where $\phi_{\text{gal}}(L_X, z | \omega)$ is the GLF for a given set of parameters, ω . We allow for pure luminosity evolution of the GLF by allowing the characteristic luminosity, L_* , to vary with redshift as

$$\log L_*(z) = \begin{cases} \log L_0 + \beta \log(1+z) & \text{if } z < z_c \\ \log L_0 + \beta \log(1+z_c) & \text{if } z \geq z_c \end{cases} \quad (16)$$

where L_0 is the characteristic luminosity at $z = 0$ and z_c is a cut-off redshift. The GLF function is thus described by five parameters, $\omega = [A, \alpha, L_0, \beta, z_c]$. We adopt fairly strong (Gaussian or lognormal) priors on the principal parameters A , α , and L_0 that describe the GLF at $z = 0$, based on a previous study using a completely independent dataset (Georgakakis et al. 2006, see Table 3). As the redshift evolution is less well constrained, we apply a weak, Gaussian prior on β (based on Georgakakis et al. 2007) and adopt a constant prior on z_c , allowing any value in the range $0.5 < z_c < 2.5$. Ultimately, we allow our own data to constrain these parameters through our exploration of the posterior parameter space (see Section 3.4 below).

We must now modify the likelihood function for our overall sample to include the contribution of normal galaxies. The expected number of galaxies in our hard-band X-ray sample is given by

$$\mathcal{N}_{\text{gal,hard}}(\omega) = \int d \log L_X \int \frac{dV}{dz} dz \int d\Gamma_{\text{gal}} \left[\phi_{\text{gal}}(L_X, z | \omega) \pi(\Gamma_{\text{gal}}) \sum_{j=1}^{N_{\text{fields}}} A_j(L_X, z, \Gamma_{\text{gal}}) \right]. \quad (17)$$

Table 3. Priors on the galaxy luminosity function parameters, ω .

Parameter	Prior type	Prior specification
$\log A$ (Mpc $^{-3}$ dex $^{-1}$)	lognormal	-3.96 ± 0.25^a
α	Gaussian	0.76 ± 0.10^a
$\log L_0$ (erg s $^{-1}$)	lognormal	41.15 ± 0.14^a
β	Gaussian	2.7 ± 0.5^b
z_c	constant	0.5 – 2.5

^aBased on measurements in Georgakakis et al. (2006), converted to rest-frame 2–10keV luminosities and our parametrization of the galaxy luminosity function (Equation 15).

^bWeak prior, based on Georgakakis et al. (2007).

An analogous expression gives the predicted number of galaxies in the soft-band sample, $\mathcal{N}_{\text{gal,soft}}(\omega)$.

We assume the spectrum for a galaxy’s X-ray emission can be described by a single power law (subjected only to Galactic absorption), and thus is described by a single parameter, the photon index Γ_{gal} . We adopt a Gaussian distribution with mean 1.9 and standard deviation 0.2 (e.g. ?) as the prior, $\pi(\Gamma_{\text{gal}})$, on the range of possible photon indices.

Our likelihood function (for the hard-band sample) must then be modified to include the expected number of normal galaxies and allow for the possibility that an individual detection is associated with a galaxy. Thus, Equation 13 is modified to

$$\mathcal{L}(\mathcal{D}_{\text{hard}} | \Theta, \omega) = e^{-(\mathcal{N}_{\text{hard}}(\Theta) + \mathcal{N}_{\text{gal,hard}}(\omega))} \times \prod_{i=1}^{n_{\text{hard}}} \left(P_i^{[\text{AGN}]} + P_i^{[\text{gal}]} \right) \quad (18)$$

where

$$P_i^{[\text{AGN}]} = \int d \log L_X \int d \log N_H \int \frac{dV}{dz} dz \left[(1 - \eta_i) p(z, L_X, N_H | D_i) \psi(L_X, z, N_H | \Theta) \right] \quad (19)$$

corresponds to the probability that source i is an AGN, for a given realization of the XLF and N_H function, and

$$P_i^{[\text{gal}]} = \int d \log L_X \int \frac{dV}{dz} dz \left[\eta_i p_{\text{gal}}(z, L_X | D_i) \phi_{\text{gal}}(L_X, z | \omega) \right] \quad (20)$$

corresponds to the probability that source i is a galaxy, for a given realization of the GLF. $p_{\text{gal}}(z, L_X | D_i)$ describes the probability distribution function for L_X and z under the assumption that source i is a galaxy and is marginalized over the single nuisance spectral parameter, Γ_{gal} . Thus,

$$p_{\text{gal}}(z, L_X | D_i) \propto p(z | d_i) \times \int d\Gamma_{\text{gal}} \mathcal{L}(T_i | c'_i(z, L_X, \Gamma_{\text{gal}}), b_i) \pi(\Gamma_{\text{gal}}) \quad (21)$$

where $c'_i(z, L_X, \Gamma_{\text{gal}})$ is the expected counts based on our single power-law X-ray spectral model (for a galaxy with redshift z , rest-frame 2–10 keV X-ray luminosity L_X , and photon index Γ_{gal}) and $\mathcal{L}(T_i | c'_i(z, L_X, \Gamma_{\text{gal}}), b_i)$ is the Poisson likelihood (cf. Equation 7). We retain the same redshift probability distribution function, $p(z | d_i)$, for the galaxy as was assumed in the case of an AGN.

The parameter η_i in Equations 19 and 20 represents any prior knowledge as to whether source i is a galaxy or an AGN. This prior

knowledge could be based on some other available data; for example an optical spectrum, morphological information, or the multiwavelength SED. However, for this work we simply assume no *a priori* preference for a galaxy or AGN; thus we set $\eta_i = 0.5$ for all sources. Instead, we allow our prior knowledge of the GLF itself (the relatively strong priors that we set on the parameters, ω) to drive our *a posteriori* inferences on the probability that a source is an AGN or a galaxy. As the GLF drops rapidly above $L_X \approx 10^{41-42} \text{ erg s}^{-1}$, any X-ray source with a higher luminosity is very unlikely to be associated with a galaxy. Due to their unabsorbed spectrum, galaxies will just as easily be seen in the soft band as the hard band; indeed, our approach naturally accounts for the rise in the number of sources detected in the soft-band at the very faintest fluxes (e.g. Georgakakis et al. 2008; Lehmer et al. 2012).

As in Section 3.2 above, the final likelihood function for the full data from both the hard- and soft-band samples, $\mathcal{L}(\mathbf{D}_{\text{hard}}, \mathbf{D}_{\text{soft}} | \Theta, \omega)$, is given by the product of the hard-band and soft-band likelihood functions.

3.4 Parameter estimation and model comparison

Having determined the likelihood function for our overall sample, our knowledge of the XLAF and GLF can be fully described by the posterior probability distribution function,

$$p(\Theta, \omega | \mathbf{D}, M) = \frac{\mathcal{L}(\mathbf{D} | \Theta, \omega) \pi(\Theta, \omega | M)}{p(\mathbf{D} | M)} \quad (22)$$

where \mathbf{D} indicates all of our observed data and we have introduced M to indicate the hypothesis that a particular model parametrization for the XLAF and GLF describes our observed data. The model itself is described by the parameters Θ, ω ; our prior knowledge of these parameters for a given model is described by

$$\pi(\Theta, \omega | M) = \pi(\Theta | M_{\text{XLAF}}) \pi(\omega | M_{\text{gal}}) \quad (23)$$

where M_{XLAF} and M_{gal} indicate our models for the XLAF and GLF respectively.

The denominator in Equation 22, $p(\mathbf{D} | M)$, is known as the Bayesian evidence, \mathcal{Z} . This factor is calculated by integrating the likelihood function over the prior parameter space:

$$\begin{aligned} \mathcal{Z} &= p(\mathbf{D} | M) \\ &= \int d\Theta \int d\omega \mathcal{L}(\mathbf{D} | \Theta, \omega) \pi(\Theta, \omega | M). \end{aligned} \quad (24)$$

The posterior probability for a particular model (given the data) can be calculated from the evidence and any prior knowledge of the probability that the model is correct:

$$p(M | \mathbf{D}) = p(\mathbf{D} | M) \pi(M) \quad (25)$$

The ratio of the posterior probabilities for different models, known as the Bayes factor, can be used for model comparison; if neither model is favored *a priori*, then the Bayes factor is simply given by the ratio of the evidences, \mathcal{Z} , of each model. We adopt this approach to compare between different models for the XLAF. We report the difference in logarithmic evidence between two models,

$$\begin{aligned} \Delta \ln \mathcal{Z} &= \ln \mathcal{Z}_1 - \ln \mathcal{Z}_2 \\ &= \ln \left(\frac{p(M_1 | \mathbf{D})}{p(M_2 | \mathbf{D})} \right) \end{aligned} \quad (26)$$

where \mathcal{Z}_1 is the Bayesian evidence for model 1 (M_1) and \mathcal{Z}_2 is the Bayesian evidence for model 2 (M_2). A difference in the logarithmic evidence of $\Delta \ln \mathcal{Z} > 4.6$ (corresponding to posterior odds

of 100:1) indicates very strong evidence in favour of the model 1, whereas $\Delta \ln \mathcal{Z} < -4.6$ indicates very strong evidence in favor of model 2 (Jeffreys 1961). If $-4.6 < \Delta \ln \mathcal{Z} < 4.6$ then, while we may still favour the model with the higher evidence, we cannot decisively rule out one model in favour of the other.

This Bayesian model comparison approach fully incorporates uncertainties in the underlying parameters and naturally applies Occam's Razor, favouring a simpler model with fewer free parameters over a more complex one, unless the latter is required by the data (e.g. Kass & Raftery 1995). Our different models for the XLF or the full XLAF are described in Sections 4 and 5 below (along with a description of the appropriate priors on the different parameters). Our model for the GLF, described in Section 3.3 above, is kept the same throughout the analysis of this paper.

To perform the integration in Equation 24 and thus calculate the evidence for a given model of the XLAF, we adopt the MULTINEST algorithm (Feroz et al. 2009), which extends the ‘‘nested sampling’’ approach of Skilling (2004) to allow for efficient exploration of large, multimodal parameter spaces. We use the MULTINEST code v3.6 with 400 ‘‘live points’’ and an efficiency factor of 0.3. The algorithm also provides estimates of the posterior probability distribution for the parameters, $p(\Theta, \omega | \mathbf{D}, M)$, which we use to obtain our ‘‘best estimates’’ of the parameters for a given model. We choose to report the posterior mean for a parameter as our best estimate. To characterize the uncertainty in a parameter, we report the 68.3 per cent (i.e. 1σ equivalent) central confidence interval, marginalized over all other parameters.

3.5 Binned estimates

To aid in the visualization of our results and the comparison between different models, it is useful to produce binned estimates of the XLF (including the GLF at low luminosities) in fine bins of L_X and z for both the hard- and soft-band samples. To achieve this, we adopt the $N_{\text{obs}}/N_{\text{mdl}}$ method (Miyaji et al. 2001), as expanded on in A10. This method compares the observed number of sources in a given L_X - z bin to that predicted based on a model fit. The binned estimate is then given by

$$\phi_b \approx \frac{N_{\text{obs}}}{N_{\text{mdl}}} \left[\phi(L_b, z_b | \hat{\Theta}) + \phi_{\text{gal}}(L_b, z_b | \hat{\omega}) \right] \quad (27)$$

where L_b and z_b are the luminosity and redshift of the center of the bin, ϕ_b is the binned estimate, and $\hat{\Theta}$ and $\hat{\omega}$ are our best (*a posteriori*) estimates of the parameters for the XLAF and GLF.

The predicted number of sources in a bin, N_{mdl} , for a given model XLAF and GLF, can be calculated from Equations 12 and 18, restricting the integration to the appropriate range in L_X and z . To calculate the observed number of sources, N_{obs} , we must account for the distribution of possible values of L_X (and in many cases z) for an individual source. A single source can make a partial contribution to the *effective* observed number in various L_X - z bins. To calculate N_{obs} , we sum the partial contributions of the individual sources to each L_X - z bin:

$$N_{\text{obs}} = \sum_{i=0}^n \int_{\log L_{1o}}^{\log L_{hi}} d \log L_X \int_{z_{1o}}^{z_{hi}} dz p(L_X, z | D_i, \hat{\Theta}, \hat{\omega}). \quad (28)$$

where $\log L_{1o}$, $\log L_{hi}$, z_{1o} and z_{hi} indicates the limits of the bin and $p(L_X, z | D_i, \hat{\Theta}, \hat{\omega})$ is the *a posteriori* probability distribution for L_X and z for source i , given the data for that source, D_i , and our best estimate of the overall XLAF and GLF (described by our

best estimates of the parameters, $\hat{\Theta}, \hat{\omega}$). Thus,

$$p(L_X, z | D_i, \hat{\Theta}, \hat{\omega}) \propto \quad (29)$$

$$\left[p_{\text{gal}}(z, L_X | D_i) \phi_{\text{gal}}(L_X, z | \hat{\omega}) \frac{dV}{dz} \right]$$

$$+ \int d \log N_H \left(p(z, L_X, N_H | D_i) \psi(L_X, z, N_H | \hat{\Theta}) \frac{dV}{dz} \right)$$

where we have marginalized over the probability distribution for N_H for source i . We assign an error to each binned estimate based on the approximate Poisson error in the effective number of observed sources in each bin. We only plot points for bins with $N_{\text{obs}} \geq 1$.

Our binned estimates of the XLF calculated in this manner will depend on the underlying XLA model. For example, an XLA that requires a high fraction of heavily absorbed sources would increase the binned estimate of the XLF based on the soft-band sample, to account for absorbed sources that would be missing from our sample. Furthermore, the individual binned estimates are not independent as a single source can enter multiple bins. Nevertheless, the binned estimates serve to visualize and compare how well the data from our soft- and hard-band samples are described by a particular model.

4 X-RAY LUMINOSITY FUNCTIONS FROM THE HARD-BAND AND SOFT-BAND SAMPLES

In this section we investigate and compare different models to describe the evolution of the XLF, neglecting the effects of absorption. For this initial study, we fit the hard-band and soft-band samples separately and determine the model parameters for the XLF from each sample. We adopt a log-constant distribution for N_H over a limited range ($20 < \log N_H < 21$) but otherwise use our Bayesian methodology described in Section 3 above (accounting for the distributions of spectral parameters and including the GLF). Given the lack of absorption corrections, the measurements of the XLF presented in this section can be seen as tracing the ‘‘observed’’ luminosity, corrected to 2–10keV rest-frame values. In Sections 4.1 to 4.3 below we describe our various parametrizations of the XLF. Section 4.4 summarizes our findings.

4.1 Luminosity-dependent density evolution (LDDE)

Following a number of studies of the XLF of AGNs based on both soft and hard X-ray selected samples (e.g. Miyaji et al. 2000; Ebrero et al. 2009; Ueda et al. 2014), we first investigated luminosity-dependent density evolution (LDDE) model parametrizations. These models start with the assumption that the XLF at $z = 0$ can be described by a smoothly connected double power-law,

$$\phi(L_X, z = 0) = \frac{d\Phi(L_X, z = 0)}{d \log L_X}$$

$$= K \left[\left(\frac{L_X}{L_*} \right)^{\gamma_1} + \left(\frac{L_X}{L_*} \right)^{\gamma_2} \right]^{-1} \quad (30)$$

where γ_1 is the faint-end slope, γ_2 is the bright-end slope, L_* is the characteristic break luminosity, and K is the overall normalization. The XLF is then modified by an evolution term, $e(z, L_X)$, which is a function of both redshift and luminosity:

$$\phi(L_X, z) = \phi(L_X, z = 0) e(z, L_X). \quad (31)$$

We adopt two different models for the evolution term. The first, hereafter referred to as the LDDE1 model, is taken from Ueda et al. (2003), and assumes that $e(z, L_X)$ is a power-law function of $(1+z)$, with different indices above and below a cutoff redshift, z_{c1} , which is itself a function of z . Thus,

$$e(z, L_X) = \begin{cases} (1+z)^{e_1} & [z \leq z_{c1}(L_X)] \\ (1+z_{c1}(L_X))^{e_1} \left(\frac{1+z}{1+z_{c1}(L_X)} \right)^{e_2} & [z > z_{c1}(L_X)] \end{cases} \quad (32)$$

where

$$z_{c1}(L_X) = \begin{cases} z_{c1}^* & [L_X \geq L_{a1}] \\ z_{c1}^* \left(\frac{L_X}{L_{a1}} \right)^{\alpha_1} & [L_X < L_{a1}]. \end{cases} \quad (33)$$

For our second parametrization (hereafter LDDE2), we adopt the further refinements of LDDE1 proposed by Ueda et al. (2014). The evolutionary term is modified to allow for a stronger density evolution at the highest redshifts, above a second cutoff redshift, z_{c2} . Thus,

$$e(z, L_X) = \begin{cases} (1+z)^{e_1} & [z < z_{c1}(L_X)] \\ (1+z_{c1}(L_X))^{e_1} \left(\frac{1+z}{1+z_{c1}(L_X)} \right)^{e_2} & [z_{c1}(L_X) < z < z_{c2}(L_X)] \\ (1+z_{c1}(L_X))^{e_1} \left(\frac{1+z_{c2}(L_X)}{1+z_{c1}(L_X)} \right)^{e_2} \left(\frac{1+z}{1+z_{c2}(L_X)} \right)^{e_3} & [z > z_{c2}(L_X)]. \end{cases} \quad (34)$$

where the cutoff redshifts are both functions of L_X , given by

$$z_{c1}(L_X) = \begin{cases} z_{c1}^* & [L_X \geq L_{a1}] \\ z_{c1}^* \left(\frac{L_X}{L_{a1}} \right)^{\alpha_1} & [L_X < L_{a1}] \end{cases} \quad (35)$$

and

$$z_{c2}(L_X) = \begin{cases} z_{c2}^* & [L_X \geq L_{a2}] \\ z_{c2}^* \left(\frac{L_X}{L_{a2}} \right)^{\alpha_2} & [L_X < L_{a2}]. \end{cases} \quad (36)$$

Furthermore, in the LDDE2 model the e_1 parameter is allowed to depend on luminosity, with the form

$$e_1(L_X) = e_1^* + \beta_1 (\log L_X - \log L_p) \quad (37)$$

as proposed in Hasinger et al. (2005).

Both of these LDDE models have a large number of parameters: 9 parameters are needed for the LDDE1 model, whereas the LDDE2 model has 15 parameters. Given the large parameter space, earlier studies often fix some of these parameters (e.g. Ueda et al. 2014, fix 6 of the parameters in the LDDE2 model). In our Bayesian analysis, we instead apply priors for all parameters but otherwise allow them to vary, thus allowing the increased parameter space to be accounted for and appropriately penalized when calculating the Bayesian evidence. The physical meaning of the parameters is somewhat obscure, thus we apply constant (or log-constant) priors over reasonable ranges given the form of the parametrization and the range of our data. Table 4 gives our best *a posteriori* estimates of the parameters based on either the hard-band sample or the soft-band sample, along with our prior limits and the Bayesian evidence for LDDE2 relative to the simpler LDDE1 parametrization. We note that the GLF is also included in each of the fits, introducing additional parameters, although these are partly constrained by the priors given in Table 3.

As can be seen from Table 4, we have very strong evidence in favor of the LDDE2 model over the LDDE1 model for both the hard-band and soft-band samples, indicating that the introduction of the additional parameters in LDDE2 is justified by the data.

Table 4. Prior limits and best *a posteriori* estimates of parameters for the LDDE1 and LDDE2 models for separate fits to the hard-band and soft-band samples (neglecting absorption effects).

Parameter	Lower limit	Upper limit	Hard band		Soft band	
			LDDE1	LDDE2	LDDE1	LDDE2
$\log K$ (Mpc ⁻³)	-7.0	-3.0	-5.63 ± 0.07	-5.72 ± 0.07	-5.87 ± 0.05	-5.97 ± 0.05
$\log L_*$ (erg s ⁻¹)	43.0	46.0	44.10 ± 0.05	44.09 ± 0.05	44.17 ± 0.04	44.18 ± 0.03
γ_1	-1.0	1.5	0.72 ± 0.02	0.73 ± 0.02	0.67 ± 0.02	0.79 ± 0.01
γ_2	1.5	4.0	2.26 ± 0.07	2.22 ± 0.06	2.37 ± 0.06	2.55 ± 0.05
e_1 (or e_1^*)	2.0	6.0	3.97 ± 0.17	4.34 ± 0.18	3.67 ± 0.09	4.35 ± 0.35
β_1	-2.0	2.0	...	-0.19 ± 0.09	...	0.65 ± 0.06
$\log L_p$ (erg s ⁻¹)	43.0	46.0	...	44.48 ± 0.44	...	44.45 ± 0.53
e_2	-5.0	0.0	-2.08 ± 0.17	-0.30 ± 0.13	-2.92 ± 0.14	-0.96 ± 0.12
e_3	-10.0	-5.0	...	-7.33 ± 0.62	...	-7.84 ± 0.51
z_{c1}^*	0.4	2.5	2.02 ± 0.09	1.85 ± 0.08	2.27 ± 0.09	1.80 ± 0.08
z_{c2}^*	2.5	3.5	...	3.16 ± 0.10	...	3.09 ± 0.11
$\log L_{a1}$ (erg s ⁻¹)	43.0	46.0	44.71 ± 0.09	44.78 ± 0.07	0.92 ± 0.11	44.92 ± 0.12
$\log L_{a2}$ (erg s ⁻¹)	43.0	46.0	...	44.46 ± 0.17	...	44.27 ± 0.30
α_1	-1.0	1.0	0.20 ± 0.01	0.23 ± 0.01	0.18 ± 0.01	0.16 ± 0.01
α_2	-1.0	1.0	...	0.12 ± 0.02	...	0.06 ± 0.02
$\Delta \ln \mathcal{Z}$			0.0	+36.1	0.0	+74.8

4.2 Luminosity and density evolution (LADE)

The LDDE model substantially warps the shape of the XLF with redshift, in particular introducing a flattening of the faint-end slope at $z \sim 1$ (as well as a further turn up at the lowest luminosities). Whether the observational evidence supports such changes in the shape of the XLF has been a matter of debate (e.g. A10, Aird et al. 2008; Fiore et al. 2012; Ueda et al. 2014; Buchner et al. 2015). Furthermore, a physical interpretation of this complex parametrization is difficult. Therefore, in A10 we proposed a model where the shape of the XLF (the smoothly-connected double power-law given in Equation 30 above) is kept the same at all redshifts, but undergoes a shift in luminosity with redshift as well as an overall decrease in density. We refer to this model as Luminosity And Density Evolution (LADE, see also e.g. Yencho et al. 2009; Ross et al. 2013, for similar parametrizations). The luminosity evolution is achieved by allowing L_* to change with redshift,

$$\log L_*(z) = \log L_0 - \log \left[\left(\frac{1+z_c}{1+z} \right)^{p_1} + \left(\frac{1+z_c}{1+z} \right)^{p_2} \right] \quad (38)$$

where p_1 and p_2 allow for a different evolution of L_* above and below a transition redshift, z_c . The additional density evolution was introduced by allowing the normalization, K , to evolve as

$$\log K(z) = \log K_0 + d(1+z). \quad (39)$$

In A10, we found that the Bayesian evidence for the LADE model was comparable to the evidence for the LDDE1 model⁷ when considering a hard-band selected sample of X-ray sources at $z < 1.2$ and colour pre-selected samples at higher redshifts.

Here, we have repeated the fitting of the LADE model with our updated samples and improved analysis (allowing for a distribution of spectral parameters and the contribution from the GLF, but assuming all sources are unabsorbed). Our best *a posteriori* estimates of the parameters for the LADE model (along with the prior constraints) are reported in Table 5, along with the Bayesian

evidence, relative to the LDDE1 model fit. For both our samples, we find strong evidence in favour of the LADE model compared to the LDDE1 model, indicating that it provides a better description of the XLF despite having (slightly) fewer free parameters. The luminosity evolution of the best-fit LADE model in this work is somewhat different to the A10 findings, with a slightly weaker luminosity evolution at low redshifts ($p_1 \approx 4$, compared to $p_1 \approx 6$ in A10) that continues to higher redshifts ($z_c \approx 2$, compared to $z_c \approx 0.8$ in A10) and has a much stronger, negative luminosity evolution at higher redshifts ($p_2 \approx -2$, compared to $p_2 \approx -0.2$ in A10). These differences indicate that the LADE provides a better description than LDDE1 for the high-redshift evolution, which is probed by our updated samples. Nevertheless, the evidence for the LDDE2 model—which introduces an additional density evolution at the highest redshifts—is significantly stronger than for the LADE model, despite the larger number of free parameters and additional complexity of LDDE2.

4.3 Flexible double-power law (FDPL)

The LADE model has a number of issues that may limit its ability to accurately describe the observed behaviour of the XLF. Firstly, the forms of the luminosity evolution (in L_*) and the density evolution (in K) are restricted to the specified functional form. Thus, the model is unable to reproduce any more complex evolutionary behaviour across certain redshift ranges, even if required by the data (e.g. at the highest redshifts). Secondly, the LADE model strictly requires that the shape of the XLF, described by the faint-end and bright-end slopes (γ_1, γ_2), remains exactly the same at all redshifts.

We therefore propose a new set of models to parametrize the evolution of the XLF, which we refer to as the Flexible Double Power-Law (FDPL) models. These models assume that the XLF can be described by a double power-law form (Equation 30) at any redshift with four parameters: K, L_*, γ_1 and γ_2 . We then allow any of these parameters to evolve with redshift, parametrizing the evolution by a polynomial function of $\log(1+z)$. By comparing the Bayesian evidence for models with different order polynomi-

⁷ The evidence for LDDE1 was higher but not by enough to decisively favour LDDE1 over LADE.

Table 5. Prior limits and best *a posteriori* estimates of parameters for the LADE model for separate fits to the hard-band and soft-band samples (neglecting absorption effects).

Parameter	Lower limit	Upper limit	Hard band LADE	Soft band LADE
$\log K$ (Mpc ⁻³)	-7.0	-3.0	-4.03 ± 0.08	-4.28 ± 0.05
$\log L_*$ (erg s ⁻¹)	43.0	46.0	44.84 ± 0.05	44.93 ± 0.03
γ_1	-1.0	1.5	0.48 ± 0.03	0.44 ± 0.02
γ_2	1.5	4.0	2.27 ± 0.07	2.18 ± 0.04
p_1	3.0	10.0	3.87 ± 0.17	3.39 ± 0.08
p_2	-4.0	3.0	-2.12 ± 0.39	-3.58 ± 0.26
z_c	0.4	3.0	2.00 ± 0.13	2.31 ± 0.07
d	-1.5	0.5	-0.19 ± 0.02	-0.22 ± 0.01
$\Delta \ln \mathcal{Z}$			+9.6	+38.4

als, we can determine whether our data require a more complicated density or luminosity evolution, or whether there is any evidence for changes in the faint-end or bright-end slopes with redshift.

With the simplest form of the polynomial, it is difficult to place meaningful priors on the free parameters: the various polynomial coefficients. Instead, we use Chebyshev polynomials over the interval $0 < \log(1+z) < \zeta_{\max}$, where $\zeta_{\max} = \log(1+z_{\max})$ and we set $z_{\max} = 7$. For a given order, n , the polynomial form is defined by the value of the parameter (e.g. $\log K$) at each of the $n+1$ Chebyshev nodes, i.e. at redshifts, z_k , where

$$\log(1+z_k) = \frac{1}{2}\zeta_{\max} + \frac{1}{2}\zeta_{\max} \cos\left(\frac{2k+1}{2(n+1)}\pi\right). \quad (40)$$

The value of e.g. $\log K(z)$ is then given by a sum of Chebyshev polynomials,

$$\log K(z) = \sum_{j=0}^n c_j T_j(x) \quad (41)$$

where $T_j(x)$ are the Chebyshev polynomials of the first kind and x is a re-scaled version of $\zeta = \log(1+z)$ over the interval $[-1, 1]$. Thus,

$$x = 2\frac{\log(1+z)}{\zeta_{\max}} - 1. \quad (42)$$

The coefficients, c_j , are given by

$$c_j = \frac{2}{n+1} \sum_{k=1}^{n+1} \log K(z=z_k) T_k(x=x_k) \quad (43)$$

where $\log K(z=z_k)$ is the value of $\log K$ at z_k .

This scheme may appear complicated but it can easily be understood as choosing values for a parameter (e.g. $\log K$) at a number of redshifts, z_k , and then finding a polynomial function to interpolate between them for all z . We can thus set a prior on the value of the parameter at each of the Chebyshev nodes (z_k). The choice of Chebyshev polynomials minimizes Runge’s phenomenon and thus prevents the adoption of highly oscillatory solutions for the redshift dependence.

The FDPL models could take a large number of possible forms, with a polynomial of arbitrary order describing the redshift dependence of each of the four double power-law parameters (K , L_* , γ_1 , γ_2). Fully exploring all possible combinations is computationally prohibitive. To find an appropriate description of the XLF we proceed as follows. First, we try to find the best parametrization for the overall density evolution and luminosity evolution by

allowing $\log K(z)$ and $\log L_*(z)$ to each be described by a polynomial of up to fourth order, but with no redshift-dependence for γ_1 and γ_2 . We evaluate the Bayesian evidence for all possible combinations of polynomials of up to fourth order for $\log K(z)$ and $\log L_*(z)$, corresponding to 25 different models. We find that the model where $\log K(z)$ is described by a second-order polynomial and $\log L_*(z)$ is described by a third-order polynomial has the highest Bayesian evidence for the hard-band sample (although we note that alternative combinations cannot be ruled out). We then fix the polynomial orders for $\log K(z)$ and $\log L_*(z)$ ⁸ and evaluate the evidence for models where $\log \gamma_1(z)$ (the faint-end slope) is described by a polynomial of up to fourth order. We choose to describe $\log \gamma_1(z)$ as a polynomial function, rather than γ_1 directly, so that extrapolation of our model to high redshifts does not result in negative, likely unphysical slopes (see Hopkins et al. 2007). We find strong evidence that $\gamma_1(z)$ changes with redshift, with a dependence that is described by a first-order polynomial (i.e. a linear relation) in $\log(1+z)$; the higher-order polynomial dependences have lower Bayesian evidence but cannot be definitively ruled out. Finally, we fix the form for $\log \gamma_1(z)$ to a first-order polynomial and test models where $\log \gamma_2(z)$ (the bright-end slope) is allowed to vary with redshift. We find strong Bayesian evidence favouring the model with no redshift-dependence for γ_2 .

Our final “best-fit” model thus consists of $\log K(z)$ being described by a second-order polynomial, $\log L_*(z)$ being described by a third-order polynomial, $\log \gamma_1(z)$ being described by a first-order polynomial and γ_2 being constant with redshift. Table 6 gives the best *a posteriori* estimates of the parameters, corresponding to the values of the parameters at each of the Chebyshev nodes. We also provide a simplified polynomial expression for the redshift dependence of each of the double power-law parameters. In addition, we give the Bayesian evidence for this FDPL model, relative to the evidence for the LDDE1 model (from Table 4). We find strong evidence in favour of our final FDPL model compared to any of the previously considered models, including LDDE2.

We then repeat the entire process with the soft-band sample. We find that the soft-band sample requires the same model form as the hard-band sample, consisting of a second-order polynomial for $\log K(z)$, a third-order polynomial for $\log L_*(z)$, a linear function for $\log \gamma_1(z)$ and a constant γ_2 (although the values of these

⁸ The form of the redshift-dependence is thus fixed for $\log K(z)$ and $\log L_*(z)$ but the parameters that describe it are still allowed to vary in subsequent fits.

parameters are different). The results, including the Bayesian evidence relative to the LDDE1 model fit to the soft-band sample, are given in Table 7. Again, we find strong evidence favouring the FDPL model over LDDE1, LDDE2, or the LADE model.

4.4 Summary of results

In this section we have explored a number of different model parametrizations that can be used to describe the XLF, measured using either our hard-band or soft-band sample (neglecting the effects of absorption). The analysis of both samples leads to similar conclusions. We find that there is strong evidence for the updated LDDE2 model of Ueda et al. (2014) compared to the simpler LDDE1 parametrization or the LADE model of A10 (where the shape of the XLF remains the same at all redshifts). We also introduce a more flexible parametrization (FDPL) that models the XLF as a double power-law but allows for arbitrary evolution in the overall luminosity or density, as well as allowing for changes in the overall shape. The final FDPL model requires 10 free parameters to describe the XLF. Our FDPL provides a simpler parametrization that is nevertheless able to reproduce the overall evolution of the XLF. The Bayesian evidence strongly favours our FDPL model over LDDE1, LDDE2 and LADE.

In Figure 6 we show how the double power-law parameters change with redshift based on our final FDPL model for the hard-band sample (black line). We find that the evolution of the XLF is driven by a combination of a) an overall density evolution (traced by the $\log K$ parameter) that peaks at $z \sim 2$, b) a relatively mild luminosity evolution that shifts the overall XLF towards higher luminosities at higher redshifts (which continues out to high redshifts), and c) a mild flattening of the faint-end slope of the XLF with increasing redshift. For comparison, we also show the redshift dependence for the parameters based on our LADE model fit to the hard-band sample (dashed red line, along with uncertainties shown by the orange hatched region). The LADE model is unable to reproduce the form of density evolution required by our FDPL model and does not allow for any flattening of the faint-end slope.

Figure 7 presents our best-fit FDPL model at a range of redshifts based on the hard-band (blue line) and soft-band (red line) samples, along with binned estimates using the method described in Section 3.5. While we find that the form of the evolution of the XLF (described by the FDPL model) is very similar for both the hard-band and soft-band samples, Figure 7 shows that there are significant differences between the XLFs from the two samples. At high luminosities ($L_X \gtrsim L_*$) the space densities based on the hard-band and soft-band samples are broadly in agreement but at lower luminosities the best-fit model and binned estimates from the soft-band sample fall significantly below the hard-band estimates. This discrepancy is most likely due to our neglect of absorption effects and motivates further investigation in Section 5 below.

5 JOINT CONSTRAINTS ON THE X-RAY LUMINOSITY AND ABSORPTION DISTRIBUTION FUNCTION

In this section we attempt to find a parametrization for the full X-ray luminosity and absorption-distribution function (XLAF) that—after accounting for the relevant sensitivity effects and selection biases—can simultaneously describe both our hard-band and soft-band samples. For each individual detection, we adopt a log-constant prior distribution for N_H in the range $20 < \log N_H < 26$ (reflecting our lack of *a priori* knowledge of the absorption) and

the corresponding constraints on L_X (e.g. Figure 5). Below, we describe our model for the XLAF (Section 5.1) and present our results (Section 5.2).

5.1 Model for the XLAF

Most previous studies adopt a parametrization for the overall XLF of all AGNs (or, in many cases, all Compton-thin AGNs) and separately model the distribution of N_H at a given L_X and z (often described as the “ N_H function” e.g. Ueda et al. 2003; La Franca et al. 2005). The overall shape of the N_H function and any luminosity or redshift dependence is defined by f_{abs} , the fraction of absorbed AGNs⁹. Previous studies have shown that f_{abs} is strongly dependent on luminosity and may also increase at higher redshifts (e.g. Hasinger 2008; Ueda et al. 2014, but see also Dwelly & Page 2006). These dependencies may be modeled directly using specific parametrizations to describe f_{abs} as a function of L_X and z (e.g. power-law functions). Further assumptions, such as setting limits on f_{abs} at certain luminosity or redshift thresholds, are often required. Additional parameters are also required to fully describe the overall distribution of N_H for a given value of f_{abs} . Such models can have a very large number of parameters and thus many of the parameters may be fixed based on measurements in the local Universe or prior expectations (e.g. Ueda et al. 2014). The resulting model parametrizations often have sharp breaks and discontinuities at certain luminosities or redshifts that may not be physical.

In this work, we therefore choose to take a different approach that builds on the flexibility of our FDPL model described above. We allow the unabsorbed ($20 < \log N_H < 22$) and absorbed ($22 < \log N_H < 24$) AGN populations to be described by independent luminosity functions— $\phi_{\text{unabs}}(L_X, z | \theta_{\text{unabs}})$ and $\phi_{\text{abs}}(L_X, z | \theta_{\text{abs}})$ —with independent sets of parameters describing their double power-law shape that can evolve differently with redshift. We assume a fixed fraction of unabsorbed AGNs, f_{21-22} , have column densities in the range $21 < \log N_H < 22$. Likewise, we assume that a fixed fraction of absorbed AGNs, f_{23-24} , have column densities in the range $23 < \log N_H < 24$. We assume a log-constant distribution of N_H within each of these 1 dex wide ranges (at a given L_X and z). The addition of f_{21-22} and f_{23-24} ensures we have sufficient freedom to describe the overall distribution of N_H within the unabsorbed or absorbed populations. We also allow for a population of Compton-thick AGNs with $24 < \log N_H < 26$, which for this work are assumed to completely track the evolution of absorbed AGNs. Thus, the XLF of the Compton-thick AGNs is assumed to be simply a factor, β_{Cthick} , times the XLF of absorbed AGNs. Our overall model of the XLAF, $\psi(L_X, z, N_H | \Theta)$, is thus given by

$$\psi(L_X, z, N_H | \Theta) = \begin{cases} (1 - f_{21-22})\phi_{\text{unabs}}(L_X, z | \theta_{\text{unabs}}) & [20 \leq \log N_H < 21] \\ f_{21-22} \phi_{\text{unabs}}(L_X, z | \theta_{\text{unabs}}) & [21 \leq \log N_H < 22] \\ (1 - f_{23-24})\phi_{\text{abs}}(L_X, z | \theta_{\text{abs}}) & [22 \leq \log N_H < 23] \\ f_{23-24} \phi_{\text{abs}}(L_X, z | \theta_{\text{abs}}) & [23 \leq \log N_H < 24] \\ \frac{\beta_{\text{Cthick}}}{2} \phi_{\text{abs}}(L_X, z | \theta_{\text{abs}}) & [24 \leq \log N_H < 26] \end{cases} \quad (44)$$

where $\Theta = \{\theta_{\text{unabs}}, \theta_{\text{abs}}, f_{21-22}, f_{23-24}, \beta_{\text{Cthick}}\}$ and thus encapsulates all parameters required to describe the XLAF. The factor $\frac{1}{2}$ in Equation 44 for Compton-thick AGNs ($24 \leq \log N_H < 26$)

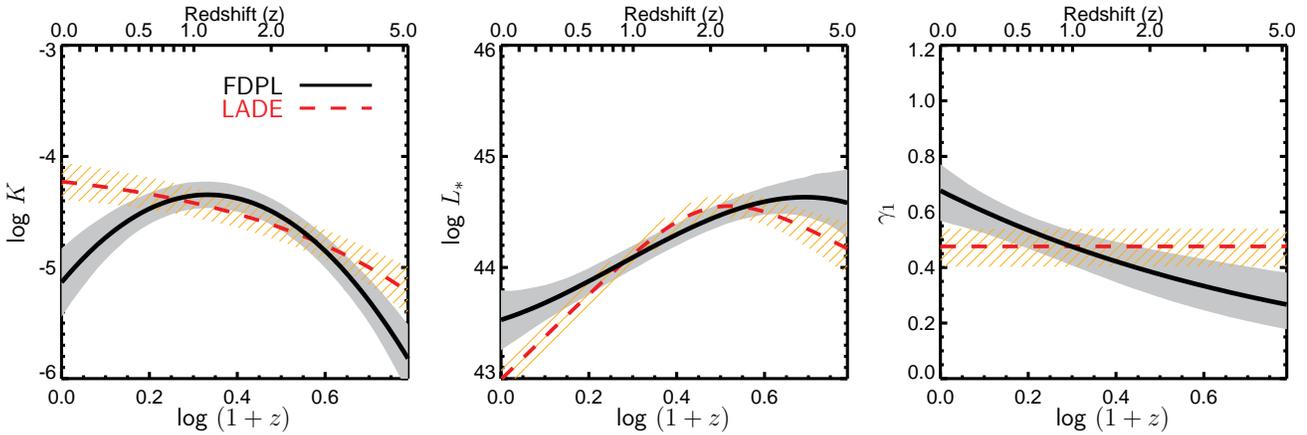
⁹ Defined here as the fraction of AGNs with column densities $22 < \log N_H < 24$ relative to all AGNs with $\log N_H < 24$

Table 6. Prior limits and best *a posteriori* estimates of parameters for the FDPL model for the hard-band sample (neglecting absorption effects).

Parameter	Lower limit	Upper limit	Parameter value at node k			
			$k = 0$	$k = 1$	$k = 2$	$k = 3$
$\log K(z = z_k)$ (Mpc^{-3})	-7.0	-3.0	-6.19 ± 0.14	-4.44 ± 0.05	-4.87 ± 0.09	...
$\log L_*(z = z_k)$ (erg s^{-1})	43.0	46.0	44.45 ± 0.17	44.61 ± 0.05	44.04 ± 0.04	43.57 ± 0.08
$\gamma_1(z = z_k)$	0.01	1.5	0.27 ± 0.04	0.58 ± 0.02
$\gamma_2(z = z_k)$	1.5	4.0	2.31 ± 0.07
$\Delta \ln \mathcal{Z}$			+45.0			
$\log K(z) = -5.13 + 4.73\zeta - 7.10\zeta^2$ $\log L_*(z) = 43.53 + 1.23\zeta + 3.35\zeta^2 - 4.08\zeta^3$ $\log \gamma_1(z) = -0.17 - 0.51\zeta$ $\gamma_2(z) = 2.31$						
where $\zeta = \log(1 + z)$						

Table 7. Prior limits and best *a posteriori* estimates of parameters for the FDPL model for the soft-band sample (neglecting absorption effects).

Parameter	Lower limit	Upper limit	Parameter value at node k			
			$k = 0$	$k = 1$	$k = 2$	$k = 3$
$\log K(z = z_k)$ (Mpc^{-3})	-7.0	-3.0	-6.47 ± 0.09	-4.73 ± 0.03	-5.20 ± 0.07	...
$\log L_*(z = z_k)$ (erg s^{-1})	43.0	46.0	44.18 ± 0.15	44.61 ± 0.04	44.12 ± 0.03	43.81 ± 0.05
$\gamma_1(z = z_k)$	0.01	1.5	0.25 ± 0.02	0.55 ± 0.02
$\gamma_2(z = z_k)$	1.5	4.0	2.34 ± 0.05
$\Delta \ln \mathcal{Z}$			+93.1			
$\log K(z) = -5.47 + 4.88\zeta - 7.20\zeta^2$ $\log L_*(z) = 43.81 - 0.27\zeta + 6.82\zeta^2 - 6.94\zeta^3$ $\log \gamma_1(z) = -0.18 - 0.55\zeta$ $\gamma_2(z) = 2.34$						
where $\zeta = \log(1 + z)$						

**Figure 6.** Dependence of the parameters that specify the double power-law form of the XLF on redshift based on best fits to the hard-band sample (neglecting absorption effects) with our FDPL model (black solid line) and LADE model (red dashed line). The solid grey and hatched orange regions indicate the 99 per cent confidence interval on the parameters for FDPL and LADE respectively based on the posterior distributions of the model parameters from our Bayesian analysis. The LADE model as proposed by A10 is restricted to a monotonic overall density evolution (traced by the $\log K$ parameter) and does not allow for any flattening of the faint-end slope (γ_1). Thus LADE is unable to produce the evolutionary behaviour described by the more flexible FDPL model, which provides a better description of our data and is thus favoured according to the Bayesian evidence.

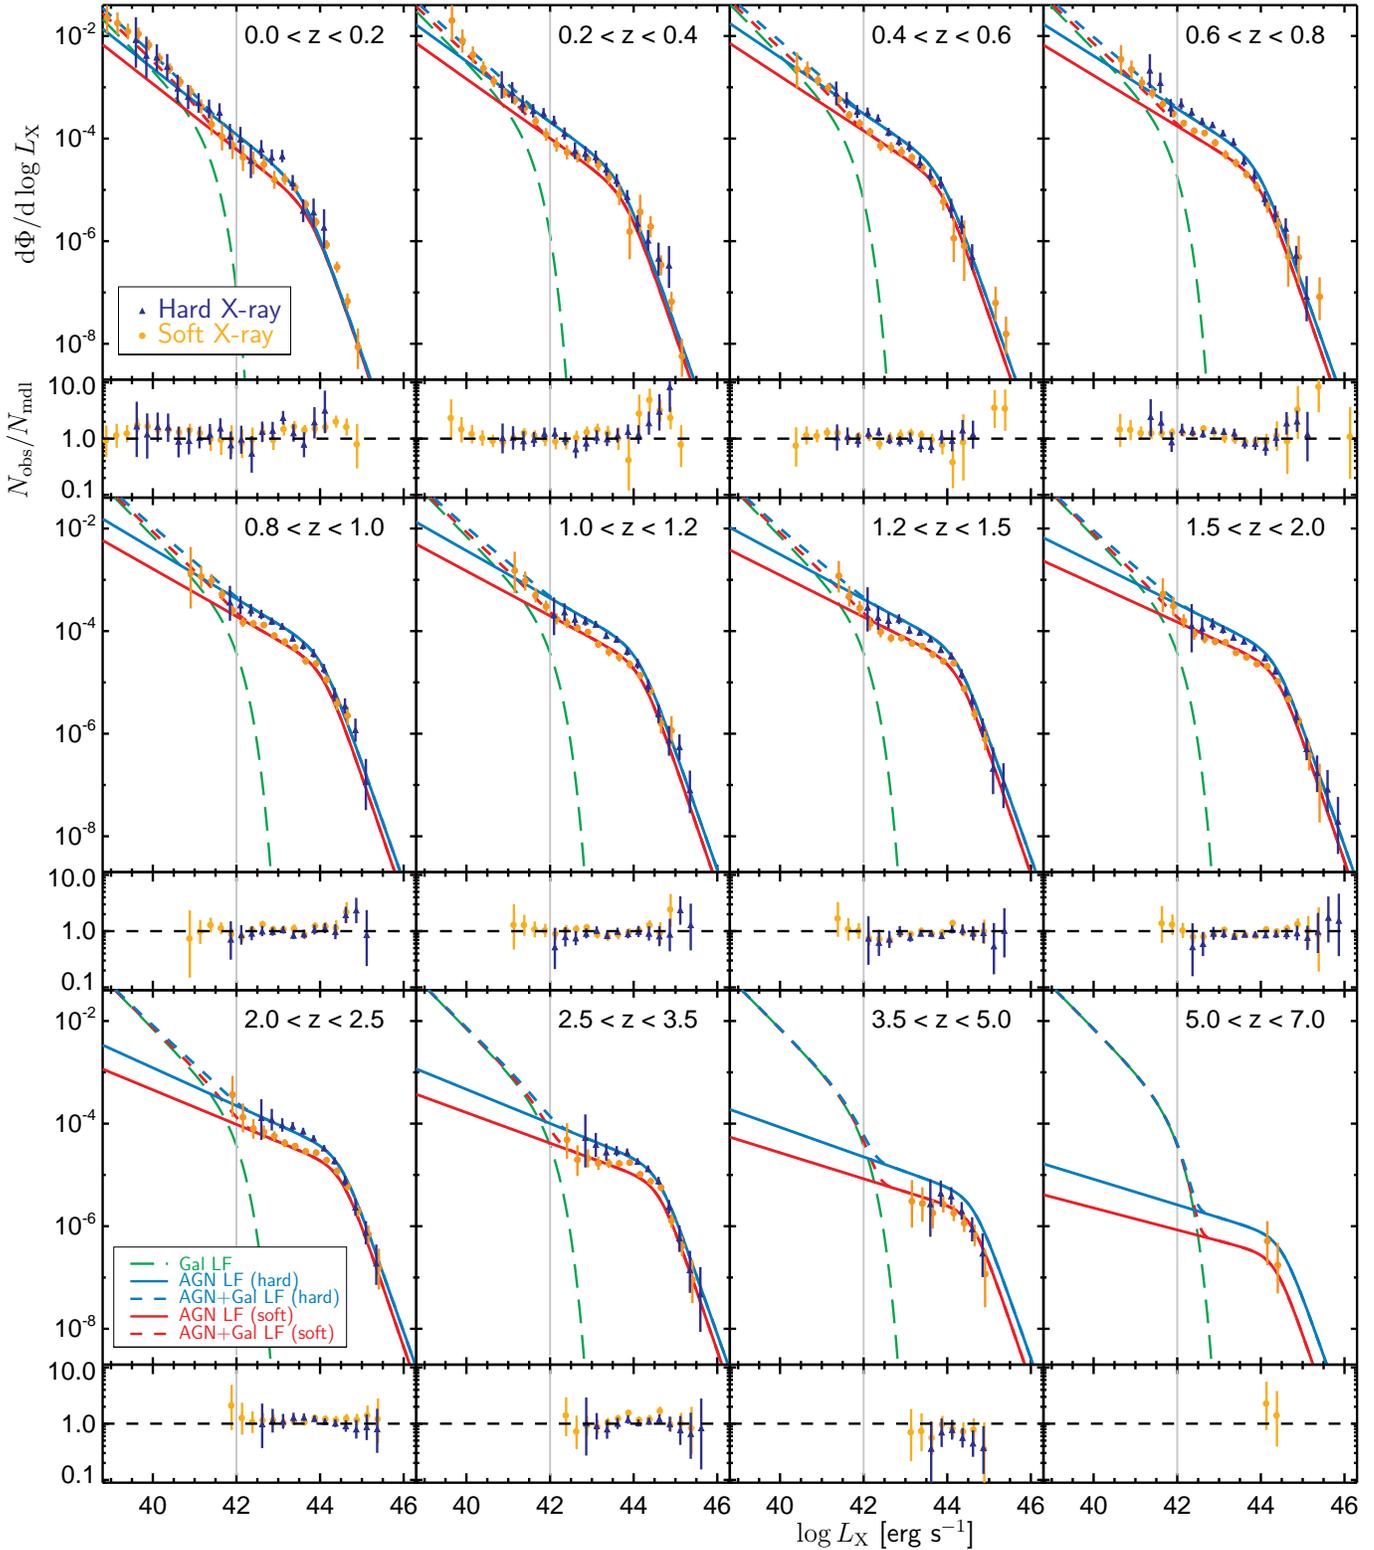


Figure 7. XLFs calculated using the hard-band and soft-band X-ray samples, neglecting absorption effects and fitting each band individually. The solid blue and solid red lines show our best-fit FDPL models for the AGN XLF based on the hard-band and soft-band samples respectively. The green dashed line shows the GLF (soft-band fit only for clarity) and the dashed blue and dashed red lines show the total luminosity function, including both AGNs and galaxies, which should be compared to the binned estimates (dark blue triangles and orange circles for the hard-band and soft-band samples respectively). The sub-panels show the residuals between the data and model in terms of the $N_{\text{obs}}/N_{\text{mdl}}$ ratio. While the form of the evolution based on the two samples is very similar, there are significant discrepancies in the space densities below the break in the XLF across all redshifts. To reconcile these discrepancies we must include the distribution of absorption column densities in our modeling and the effects on the inferred luminosities and space densities for both samples.

represents our assumption that the column densities are evenly distributed over the 2 dex wide bin in $\log N_{\text{H}}$.

It is worth noting that in our analysis both the hard- and soft-band samples will contain AGN from the unabsorbed, absorbed, and Compton-thick AGN populations, although the soft-band sample is likely to contain a lower fraction of absorbed sources than the hard-band sample (due to the detection biases against absorbed sources at soft energies). We do not know if an individual hard- or soft-band detection corresponds to an unabsorbed, absorbed, or Compton-thick source; we are only able to make statements about the population as a whole, based on our model for the XLAF.

Marginalizing over $\log N_{\text{H}}$, we can recover the total XLF for all AGN,

$$\begin{aligned} \phi(L_{\text{X}}, z | \Theta) &= \int_{20}^{26} d \log N_{\text{H}} \psi(L_{\text{X}}, z, N_{\text{H}} | \Theta) \\ &= \phi_{\text{unabs}}(L_{\text{X}}, z | \theta_{\text{unabs}}) \\ &+ \phi_{\text{abs}}(L_{\text{X}}, z | \theta_{\text{abs}}) \\ &+ \beta_{\text{Cthick}} \phi_{\text{abs}}(L_{\text{X}}, z | \theta_{\text{abs}}) \end{aligned} \quad (45)$$

We assume that ϕ_{unabs} and ϕ_{abs} are each individually described by an FDPL model parameterization where the four XLF parameters— K , L_* , γ_1 , γ_2 —can all vary with redshift according to polynomial functions of $\log(1+z)$. Fully exploring all possible combinations of polynomials of different orders for the redshift-dependence of all eight of these shape parameters is computationally prohibitive. Instead, we initially assume that ϕ_{unabs} and ϕ_{abs} are both described by the form found in Section 4.3 when analyzing the hard- and soft-band samples separately. Thus, we adopt a second-order polynomial for $\log K(z)$, a third-order polynomial for $\log L_*(z)$, a first-order polynomial for $\log \gamma_1(z)$, and a constant γ_2 as our baseline model. We then increase or decrease (if possible) the polynomial order for each of the eight shape parameters in turn and evaluate the Bayesian evidence relative to our baseline model.

5.2 Results

In Table 8 we give the change in the logarithmic Bayesian evidence, $\Delta \ln \mathcal{Z}$, found when increasing or decreasing the polynomial order for the redshift dependence of each of the FDPL parameters in turn, relative to the evidence for our baseline model. In the majority of cases, the altered model has a *lower* Bayesian evidence, thus favouring our baseline model. For the normalization (K) of both the unabsorbed and absorbed AGN XLF and the L_* of the absorbed AGN XLF, we can decisively rule out the simpler models (with lower order polynomials used to describe the redshift dependence) as $\Delta \ln \mathcal{Z} < -4.6$. For the L_* of the absorbed AGN, a simpler (second-order polynomial) model is weakly favoured by our data ($\Delta \ln \mathcal{Z} > 0$), but the evidence for this simplification is not strong. In all other cases, the baseline model is very weakly favoured ($\Delta \ln \mathcal{Z} < 0$).

We also tested a model with no change in the faint-end slope for *both* the unabsorbed and absorbed AGN XLFs (but otherwise retaining the baseline model). While such a model is weakly favored over the baseline ($\Delta \ln \mathcal{Z} = 0.71$), we are unable to choose decisively between these models based purely on our data. We therefore choose to retain the baseline model, which has greater flexibility to describe the faint end of the XLF. Nevertheless, as we are unable to fully explore all possible combinations of different polynomials, our final model must be seen as one possible, but adequate, way of describing our observational data. The evolution of

Table 10. Best *a posteriori* estimates of parameters describing the galaxy luminosity function.

Parameter	Value
$\log A$ ($\text{Mpc}^{-3} \text{ dex}^{-1}$)	-3.59 ± 0.08
α	0.81 ± 0.03
$\log L_0$ (erg s^{-1})	41.12 ± 0.07
β	2.66 ± 0.24
z_c	0.82 ± 0.10

N.B. All parameters are partially constrained by our informative priors, specified in Table 3

the faint-end slope of the XLF is discussed in more detail in Section 6.2 below.

In Table 9 we give our best estimates of the XLAF parameters: the shape parameters for the unabsorbed and absorbed XLFs and their redshift dependencies in terms of Chebyshev polynomials and in a simplified form, as well as f_{21-22} , f_{23-24} and β_{Cthick} . In Table 10 we also provide our best *a posteriori* estimates of the parameters that describe the GLF, which are constrained along with the AGN XLAF.

Figure 8 shows our model for the total AGN XLF (found by marginalizing over N_{H} , see Equation 46) compared to binned estimates based on our hard- and soft-band samples. The binned estimates from both the hard- and soft-band samples are consistent with the total XLF (with an additional contribution from the GLF at low luminosities), showing that our overall model of the XLAF is able to describe both of our samples (cf. Figure 7, where absorption effects were neglected).

Figure 9 compares our model XLFs of the unabsorbed (blue) and absorbed (red) AGNs at a range of redshifts, as well as showing the total XLF of the full population (black). The XLF of absorbed AGNs generally has a lower break luminosity (L_*), a higher normalization, and a steeper faint-end slope than the XLF of unabsorbed AGNs. Hence, the absorbed AGN population is dominant at low luminosities, whereas unabsorbed AGN are increasingly important at higher luminosities. Both the absorbed and unabsorbed AGN XLFs undergo a strong luminosity evolution, shifting to higher luminosities from $z \sim 0$ out to $z \sim 3$. Both XLFs also evolve in overall density, with the absorbed AGNs evolving more strongly and thus making up a larger fraction of the total AGN population at $z = 1$, even at high luminosities, but dropping away more rapidly at higher redshifts. The evolution of the XLF shape parameters with redshift is presented in Figure 10. The differing evolution of the absorbed and unabsorbed AGN XLFs and thus their relative contributions to the total space density of AGNs results in a complex evolution in the shape of the total XLF of AGNs (solid black line in Figure 9).

In Figure 11 we plot the fraction of absorbed AGNs (relative to the total Compton-thin population), f_{abs} , as a function of L_{X} at various redshifts (i.e. the ratio of the absorbed AGN XLF to the sum of the absorbed and unabsorbed XLFs). The absorbed fraction roughly plateaus at high and low luminosities, with a fairly rapid transition in between due to the difference in the break luminosities of the XLFs. The strong luminosity evolution of both XLFs causes this transition to shift to higher luminosities at higher redshifts. The stronger overall density evolution of the absorbed AGNs also leads to an increase in f_{abs} at high luminosities ($L_{\text{X}} \gtrsim 10^{45} \text{ erg s}^{-1}$), where absorbed AGNs increasingly dominate, although our constraints on the bright end of the absorbed

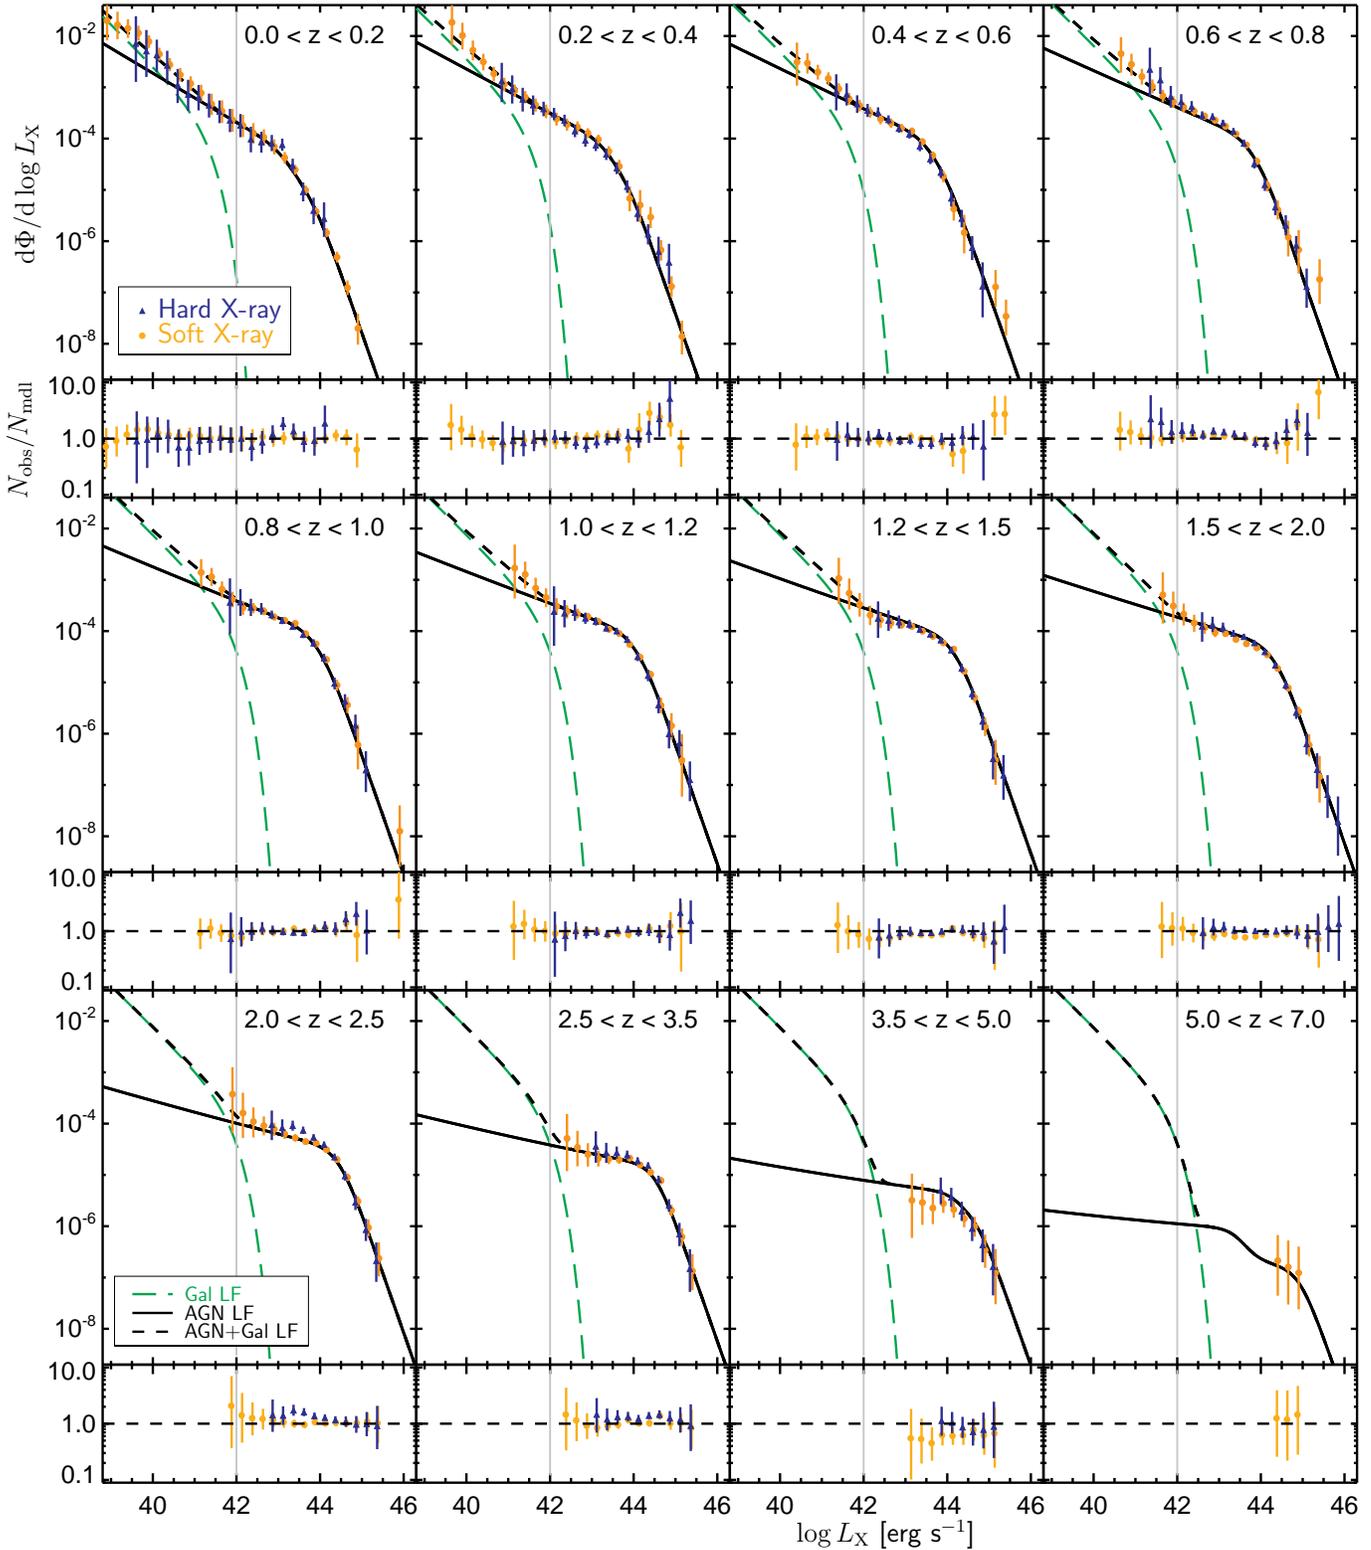


Figure 8. XLF of AGNs across the full range of absorption ($20 < \log N_{\text{H}} < 26$) calculated by fitting both the hard- and soft-band X-ray samples to our full X-ray luminosity and absorption-distribution function (XLAf) model described in Section 5.1 (solid black line). We also show the GLF (green dashed line) which is fitted simultaneously with the AGN XLAf. Binned estimates are based on either the hard-band (dark blue triangles) or soft-band (orange circles) X-ray samples and should be compared to the total XLF, including both AGNs and galaxies (black dashed line). The sub-panels show the residuals between the data and model in terms of the $N_{\text{obs}}/N_{\text{mdl}}$ ratio. Our final model of the full XLAf of AGNs allows us to reconcile measurements based on the hard- and soft-band X-ray samples as ultimately coming from the same AGN population (but subject to different selection biases and incompleteness as a function of N_{H} , z and L_X). Hence, both the hard- and soft-band estimates are consistent with the total XLF (cf. Figure 7).

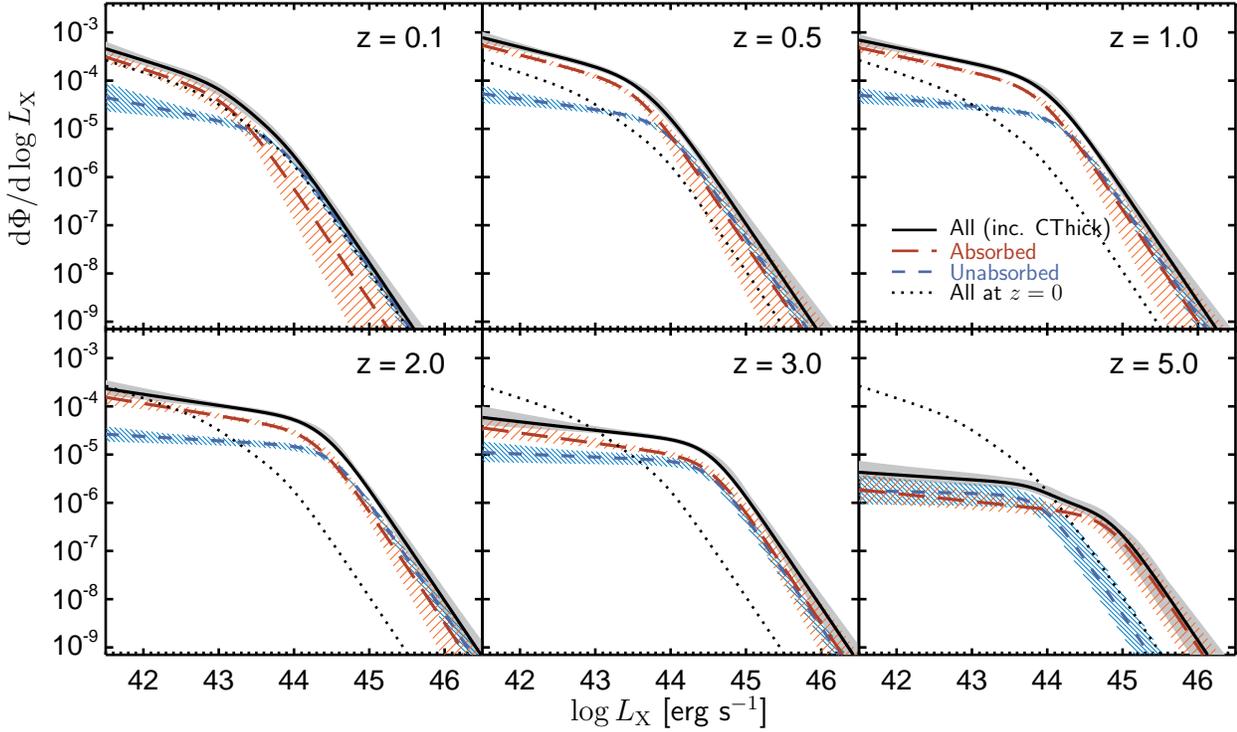


Figure 9. Our best model for the total XLF of AGNs (black solid line) and the constituent XLFs of the unabsorbed ($20 < \log N_{\text{H}} < 22$: blue short-dashed lines) and absorbed ($22 < \log N_{\text{H}} < 24$: red long-dashed lines) AGNs evaluated at a number of redshifts, based on our simultaneous fitting of both the hard- and soft-band X-ray samples to our XLAf model. The shaded regions indicate the 99 per cent confidence interval based on the posterior distribution of the model parameters. At $z \lesssim 2$, unabsorbed AGNs tend to dominate at high luminosities, whereas absorbed AGNs—due to the lower L_* of their XLF—dominate at lower luminosities. Both populations undergo a strong luminosity evolution (generally shifting to higher luminosities as redshift increases) and an overall density evolution, changing the ratio of absorbed to unabsorbed AGNs at different luminosities and resulting in a complex evolution in the shape of the total XLF. Extrapolating our model to the highest redshifts ($z \sim 5$) indicates that absorbed AGNs may start to dominate at the highest luminosities; however, our observational constraints are weak in this regime and thus this behaviour may be an artefact of our particular model parametrization.

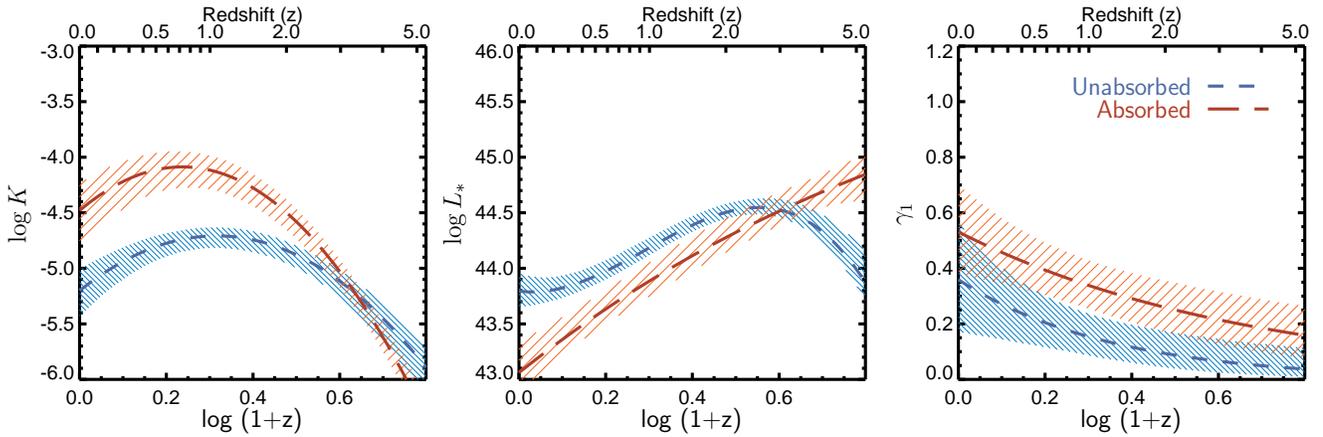


Figure 10. Shape parameters (normalization K , characteristic break luminosity L_* , and faint-end slope γ_1) for the XLFs of unabsorbed (blue short-dashed line) and absorbed (red long-dashed line) as a function of redshift, calculated by fitting both the hard- and soft-band X-ray samples to our full XLAf model. The shaded regions indicate the 99 per cent confidence interval based on the posterior distribution of the model parameters and thus may under-represent the uncertainty in the observations at a particular redshift.

Table 8. Bayesian evidence (relative to our baseline XLAF model) for increasing or decreasing the polynomial order for the redshift dependence of each parameter in turn in the independent FDPL models (for the unabsorbed and absorbed populations) adopted for our modeling of the XLAF.

AGN population	Parameter	Order (increased)	$\Delta \ln \mathcal{Z}$	Order (decreased)	$\Delta \ln \mathcal{Z}$
Unabsorbed	$\log K$ (Mpc^{-3})	3	-0.86	1	-22.74
	$\log L_*$ (erg s^{-1})	4	-3.12	2	-10.26
	γ_1	2	-3.24	0	-0.25
	γ_2	1	-2.45
Absorbed	$\log K$ (Mpc^{-3})	3	-4.18	1	-27.23
	$\log L_*$ (erg s^{-1})	4	-3.54	2	+2.85
	γ_1	2	-4.23	0	-2.15
	γ_2	1	-2.96

Table 9. Prior limits and best *a posteriori* estimates of parameters for our full XLAF model, consisting of independent XLFs (described by the FDPL model) for the unabsorbed and absorbed AGN populations, and a scalar times the absorbed AGN XLF for the Compton-thick population.

Parameter	Lower limit	Upper limit	Parameter value at node k			
			$k = 0$	$k = 1$	$k = 2$	$k = 3$
Unabsorbed AGNs ($20 < \log N_{\text{H}} < 22$)						
$\log K(z = z_k)$ (Mpc^{-3})	-7.0	-3.0	-6.17 ± 0.11	-4.81 ± 0.04	-5.03 ± 0.06	...
$\log L_*(z = z_k)$ (erg s^{-1})	43.0	46.0	43.26 ± 0.14	44.49 ± 0.03	44.15 ± 0.03	43.79 ± 0.05
$\gamma_1(z = z_k)$	0.01	1.5	0.04 ± 0.02	0.25 ± 0.04
$\gamma_2(z = z_k)$	1.5	4.0	2.32 ± 0.04
f_{21-22}	0.1	0.9	0.43 ± 0.04
$\log K(z) = -5.21 + 0.72\zeta - 0.26\zeta^2$ $\log L_*(z) = 43.56 + 0.32\zeta + 0.32\zeta^2 - 0.12\zeta^3$ $\log \gamma_1(z) = -0.44 - 0.52\zeta$ $\gamma_2(z) = 2.32$						
where $\zeta = \log(1 + z)$						
Absorbed AGNs ($22 < \log N_{\text{H}} < 24$)						
$\log K(z = z_k)$ (Mpc^{-3})	-7.0	-3.0	-6.81 ± 0.09	-4.44 ± 0.05	-4.30 ± 0.09	...
$\log L_*(z = z_k)$ (erg s^{-1})	43.0	46.0	44.96 ± 0.13	44.57 ± 0.06	43.85 ± 0.06	43.17 ± 0.08
$\gamma_1(z = z_k)$	0.01	1.5	0.16 ± 0.04	0.43 ± 0.04
$\gamma_2(z = z_k)$	1.5	4.0	2.33 ± 0.18
f_{23-24}	0.1	0.9	0.70 ± 0.02
$\log K(z) = -4.48 + 0.76\zeta - 0.37\zeta^2$ $\log L_*(z) = 43.06 + 0.98\zeta - 0.14\zeta^2 + 0.01\zeta^3$ $\log \gamma_1(z) = -0.28 - 0.28\zeta$ $\gamma_2(z) = 2.33$						
where $\zeta = \log(1 + z)$						
Compton-thick AGNs ($24 < \log N_{\text{H}} < 26$)						
β_{Cthick}	0.2	2.0	0.34 ± 0.08

AGN XLF are fairly poor. Our model reproduces a luminosity and redshift dependence of the absorbed fraction and is generally consistent with direct estimates based on X-ray spectral classifications (e.g. Ueda et al. 2014; Buchner et al. 2015). We do not find any evidence for the *decrease* in f_{abs} with decreasing luminosity, which is seen by Buchner et al. (2015) and some other studies (e.g. Burlon et al. 2011; Brightman & Nandra 2011b), although this may reflect the limited diagnostic power of our method to distinguish absorbed and unabsorbed AGNs at the lowest luminosities.

Our model indicates that the positive luminosity evolution of the unabsorbed AGN population slows down by $z \sim 3$ and appears

to start a *negative* luminosity evolution to higher redshifts (see central panel of Figure 10). At $z \gtrsim 3$, the negative luminosity evolution of the unabsorbed AGN XLF combined with the density evolution of the absorbed AGN XLF leads to a change in the pattern of f_{abs} . At $z = 3$, f_{abs} is approximately constant as a function of L_X^{10} , whereas at higher redshifts the absorbed AGNs appear to domi-

¹⁰ Such a lack of luminosity dependence in the absorbed fraction at $z \gtrsim 3$ has also been found by Vito et al. (2014), albeit classifying column densities of $N_{\text{H}} > 10^{23} \text{ cm}^{-2}$ as absorbed.

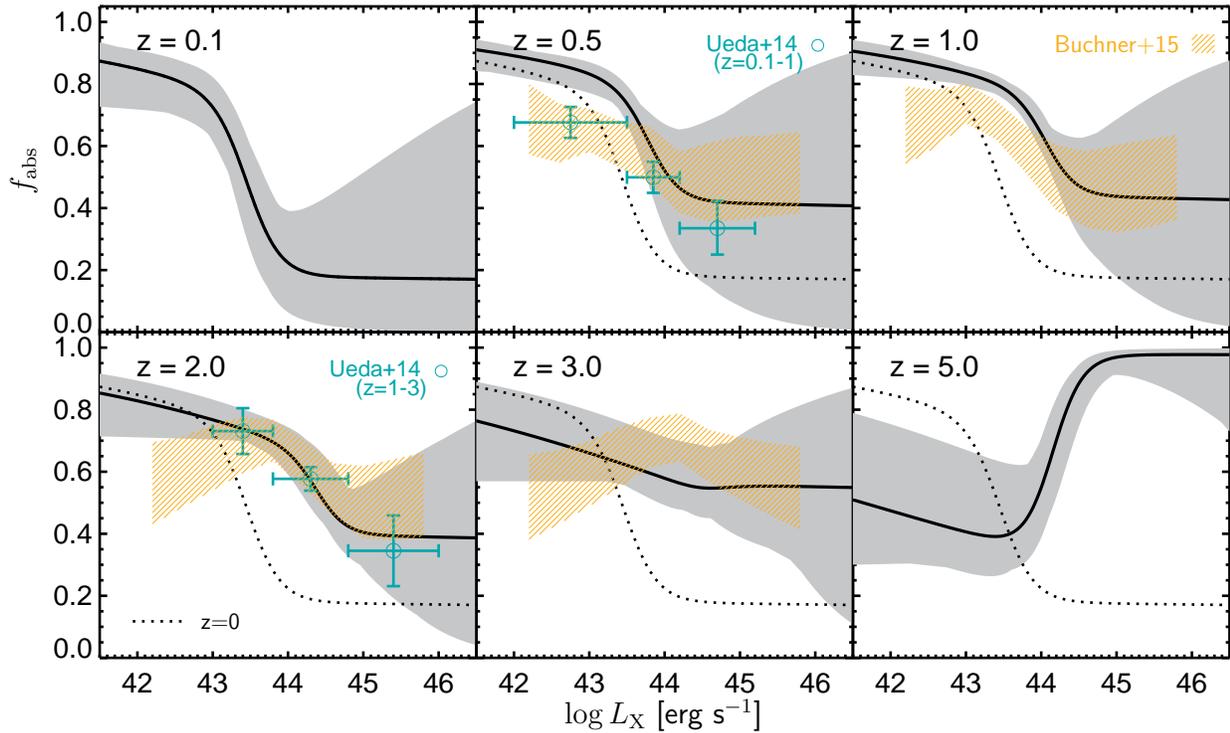


Figure 11. Fraction of absorbed AGNs, f_{abs} (defined here as the fraction of AGNs with $22 < \log N_{\text{H}} < 24$ relative to all Compton-thin, $\log N_{\text{H}} < 24$, AGNs) based on our best-fit model of the XLAF evaluated at a number of redshifts (solid black lines). The grey regions indicate the 99 per cent confidence interval based on the posterior distribution of the model parameters. The model evaluated at $z = 0$ is shown by the dotted line in all panels. We also show direct estimates of the absorbed fraction as a function of luminosity from Ueda et al. (2014, green circles, for the indicated redshift ranges) and Buchner et al. (2015, orange shaded regions indicate their 90 per cent confidence intervals for bins that approximately span each redshift). Our flexible model reproduces a luminosity dependence of f_{abs} , which plateaus at low and high luminosities, with a rapid transition in between. The position of this transition shifts to higher luminosities at higher redshifts due to the luminosity evolution of both the unabsorbed and absorbed AGN XLFs. This pattern is roughly consistent with the Ueda et al. (2014) and Buchner et al. (2015) measurements, although our estimates are slightly higher at low luminosities. At $z = 3$ we find that the absorbed fraction is roughly constant as a function of luminosity. At higher redshifts ($z = 5$), absorbed AGNs may dominate at high luminosities, although this behaviour may be driven by our assumed model XLFs and the extrapolation to these high redshifts.

nate at the highest luminosities. These results hint at a complex and differing evolution of absorbed and unabsorbed AGNs in the early Universe. However, our observational constraints on the total XLF at $z \sim 4 - 7$ are generally fairly poor (see Figure 9). Furthermore, our diagnostic power to determine the absorbed fraction of AGN—comparing the hard- and soft-band X-ray samples—is substantially reduced at these redshifts. The hard-band sample is small at these redshifts and both the hard and soft bands probe high rest-frame energies and are thus only weakly affected by absorption. Thus, the behaviour at high redshifts may be driven by the functional form that best describes the lower redshift data and the extrapolation of this model to high redshifts.

Nonetheless, our model is able to reproduce the overall evolution of the XLF and by accounting for absorption can reconcile our hard- and soft-band X-ray samples. Furthermore, we naturally recover the luminosity dependence of the absorbed fraction at $z \sim 0 - 2$ and its evolution.

6 DISCUSSION

6.1 Comparison with previous studies of the X-ray luminosity function

Our work allows us to interpret the evolution of the total XLF of AGNs as due to the combination of the XLFs of unabsorbed and absorbed (as well as Compton-thick) AGNs. We find that both XLFs can be described by our FDPL model, albeit with different sets of parameters. Our approach differs from most prior studies of the evolution of the XLF, which instead attempt to find a model that directly describes the *total* XLF, either using an LDDE parametrization or some form of luminosity and density evolution (e.g. Ebrero et al. 2009; Yencho et al. 2009; Aird et al. 2010).

In Figure 12 we compare the total XLF based on our final model (black solid line) to models from three recent studies: the LADE model from A10; the LDDE2 model put forward by Ueda et al. (2014, hereafter U14); and the LDDE parametrization used in recent work by ?, hereafter M15. We also compare estimates from Buchner et al. (2015, hereafter B15) who used a non-parametric method to estimate the space density of AGNs in fixed bins of L_{X} , N_{H} and z with only simple smoothness assumptions used to link values between bins.

At $z \lesssim 3$ there is good agreement between our work and the

previous studies, although there are a number of important differences. The bright end of our XLF is generally in good agreement with previous studies, although our updated work predicts a substantially higher space density than the A10 model at both the lowest ($z = 0.1$) and highest ($z \gtrsim 2$) redshifts. The B15 estimates are higher than our model at the bright end, especially at lower redshifts. However, this work lacked the very large-area ($\gtrsim 50 \text{ deg}^2$) samples included in our study and other work. Thus, the B15 estimates at high luminosities and low redshifts may be driven by the smoothness assumptions inherent to their method.

At $z \approx 0.5 - 1$ the faint end of our total XLF generally lies slightly above the A10 model. The differences between A10 and our updated study are likely due to a combination of absorption effects (neglected in A10), the limitations of the LADE model, and the smaller sample in the A10 work. The LDDE models from U14 and M15 are in better agreement with our work, although there are slight differences in the shapes that may reflect differences in the parametrizations. For example, the M15 XLF is also somewhat higher than our model at the faintest luminosities at $z \lesssim 0.5$, which may be due to the LDDE parametrization requiring a steep faint-end slope (that is altered by the luminosity-dependent evolution). Our study probes to lower luminosities than M15, accounting for normal galaxies at low luminosities, and can thus rule out such a steep slope. We also predict a slightly higher space density than U14 or M15 at $L_X \sim 10^{43-44} \text{ erg s}^{-1}$ and $z = 0.5$, and a slightly lower space density at moderate luminosities ($L_X \lesssim 10^{44} \text{ erg s}^{-1}$) at $z = 3$. The B15 estimates are generally in good agreement with our results at fainter luminosities out to $z = 3$, although our model tends to lie towards the lower end of their confidence intervals. There are small differences close to L_* at $z = 0.5$ and $z = 1.0$ whereby B15 predicts slightly lower space densities, which may also be driven by the assumed functional form and differences in the methodology.

At $L_X \sim 10^{44-45} \text{ erg s}^{-1}$ and $z = 5.0$ our model is below the extrapolation of the A10 work or the B15 estimates but is consistent with the recent work of U14 and M15 (which did include some sources at these redshifts and luminosities). We find a strong decline in the space density in of AGNs to high redshifts in this moderate luminosity range, which is consistent with previous studies (e.g. Civano et al. 2011; Vito et al. 2013). However, the high-redshift behaviour of our model may, in part, be driven by the extrapolation of the functional form that best describes the lower redshift XLF, where the bulk of our sample lies. We will present a focused study of the evolution of the XLF at high redshifts in future work (Georgakakis et al. in preparation).

As an additional check, in Figure 13 we compare the extrapolation of our model XLF to $z = 0$ with measurements based on local AGN populations. Our total XLF (left panel) does not undergo a sharp break, but flattens gradually below $L_X \sim 10^{44} \text{ erg s}^{-1}$ due to the mixing of the unabsorbed and absorbed AGN XLFs. This mixing results in a relatively steep slope in the $L_X \sim 10^{43-44} \text{ erg s}^{-1}$ luminosity range, consistent with the Ueda et al. (2011) measurements shown by the black squares. The right panel compares direct measurements of the XLFs of unabsorbed ($N_H < 10^{22} \text{ cm}^{-2}$) and absorbed ($N_H > 10^{22} \text{ cm}^{-2}$) from Burlon et al. (2011) using the *Swift*/BAT sample, where the populations were divided based on the results of X-ray spectral fitting for individual sources (cf. our statistical method). The extrapolation based on our baseline model for unabsorbed and absorbed AGNs is in good agreement with these measurements, which is reassuring given that our model is primarily constrained by much higher redshift sources and we do not directly measure N_H for individual sources. We

note that the lowest two luminosity bins for the unabsorbed AGNs from Burlon et al. (2011) lie above our model (although only by $\sim 1 - 2\sigma$). Thus, Burlon et al. (2011) found that the absorbed fraction decreases at both high and low luminosities (see also Brightman & Nandra 2011b; Buchner et al. 2015). Our model does not generally predict such a behaviour in the absorbed fraction¹¹. We find that f_{abs} plateaus at ~ 85 per cent at low luminosities (see Figure 11). In Figure 13 (right panel) we also show our final best fit GLF. Normal galaxies have much lower luminosities than the unabsorbed AGN found by Burlon et al. (2011). Thus, the lack of a decrease in f_{abs} at low luminosities does not appear to be due to us (statistically) misclassifying some fraction of our X-ray sources as galaxies rather than unabsorbed AGNs.

In conclusion, our work is mostly consistent with recent estimates of the total XLF of AGNs at $z \approx 0 - 4$. However, our new model provides a simpler way of interpreting the evolution of the XLF and the absorbed fraction as due to the combination of the unabsorbed and absorbed AGN XLFs and their slightly different evolution. We discuss the possible physical origin of this behaviour in Section 6.5 below.

6.2 Does the faint end of the XLF flatten at high redshifts?

How the shape of the XLF evolves with redshift, and in particular whether the faint-end slope becomes progressively flatter at higher redshifts, remains a major question in studies of the evolution of the AGN population (e.g. Hopkins et al. 2006; Aird et al. 2010). In our study, we do not directly constrain the faint-end slope of the total XLF of AGNs – instead we describe the total XLF as a combination of the XLFs of the unabsorbed and absorbed AGNs and infer different faint-end slopes for each population (hereafter referred to as γ_{unabs} and γ_{abs} respectively). Our baseline model allows both γ_{unabs} and γ_{abs} to change with redshift and our results indicate that both slopes become flatter as redshift increases (see Figure 10). At all redshifts, γ_{unabs} is flatter than γ_{abs} . However, based purely on our data and analyses, we are unable to rule out a simpler model where neither γ_{unabs} or γ_{abs} change with redshift (hereafter the “no-change” model). We still require that γ_{unabs} is flatter than γ_{abs} in the no-change model.

Despite these issues, we find that the faint-end slope of the total XLF does appear to change with redshift. This can be seen in Figure 9. At $z = 0.1$ there is a gradual transition in the $L_X \sim 10^{43-44} \text{ erg s}^{-1}$ luminosity range between the unabsorbed AGN XLF (dominant at high luminosities) and the absorbed AGN XLF (dominant at low luminosities) which results in a relatively steep slope at these intermediate luminosities. It is only at fainter luminosities that the absorbed AGN XLF starts to dominate and a flatter overall slope is seen¹². At higher redshifts, the absorbed AGN XLF is increasingly dominant at moderate luminosities and thus a clearer break and a flatter faint-end slope is seen in the total XLF at $L_X \sim 10^{42-44} \text{ erg s}^{-1}$ by $z \sim 1$. The underlying flattening of γ_{abs} in our baseline model contributes to this effect but the mixing of the populations appears to be more important.

¹¹ We note that fitting the unabsorbed and absorbed XLFs as independent broken double power laws could reproduce such a pattern, as in Burlon et al. (2011), thus the lack of such behaviour is not due to a limitation of our XLAF model.

¹² Indeed, an advantage of our own study is that we probe down to very faint luminosities ($L_X \sim 10^{39} \text{ erg s}^{-1}$) at low redshifts using deep *Chandra* fields, carefully considering the contribution from normal galaxies, and improve the overall constraints on the faint end of the AGN XLF.

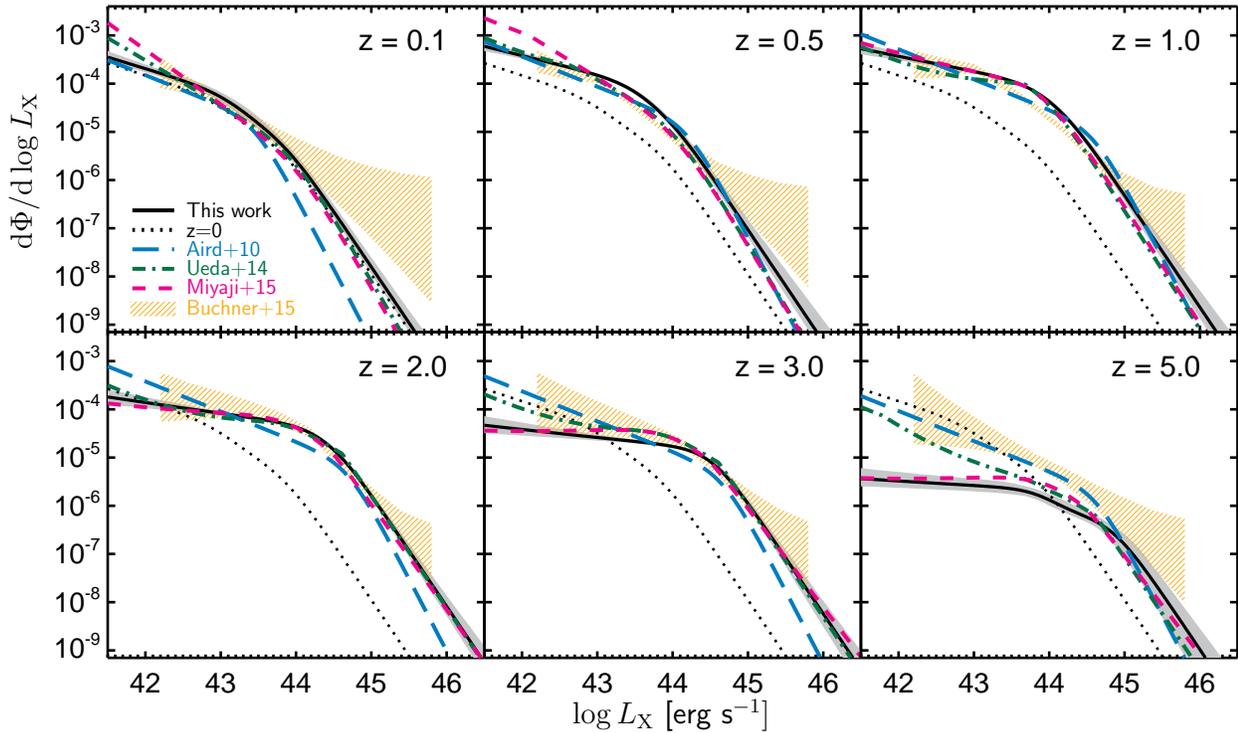


Figure 12. Comparison of our model for the total XLF of all Compton-thin ($N_{\text{H}} < 10^{24} \text{ cm}^{-2}$) AGNs (black solid line, grey region indicates 99 per cent confidence interval in model parameters) to prior model fits: the LADE model from Aird et al. (2010, blue long-dashed line); the LDDE2 model from Ueda et al. (2014, green dot-dashed line); and the LDDE model from ?, pink short-dashed line. We also show estimates from Buchner et al. (2015) who used a non-parametric method with simple smoothness assumptions to estimate the space density in fixed bins of L_X , N_{H} and z (orange hatched region indicate their 90 per cent confidence intervals for bins that approximately span the indicated redshift and include all Compton-thin AGNs).

To further examine this effect, in Figure 14 we show estimates of the *apparent* shape of the total XLF as a function of redshift. We estimate double power-law shape parameters by taking our overall model at various redshifts, creating fake data points in 0.25 dex wide redshift bins for $42.5 < \log L_X < 46$, and performing a simple least squares fit to the standard double power law form. The solid black lines show the shape parameters of the total XLF for our baseline model, whereas the dashed green line shows the parameters for the no-change model. In both cases, the evolution of the total XLF is very similar – described by a combination of density evolution, positive luminosity evolution, and an *apparent* flattening of the faint-end slope as redshift increases. The apparent flattening is driven by the mixing of the constituent unabsorbed and absorbed AGN XLFs and their relative evolution. Both the baseline and no-change models predict a similar evolution of the total XLF, including this apparent flattening. We also show estimates of the shape parameters of the total XLF from M15, which they calculated by directly fitting the XLF with a double power-law in thin redshift shells. Their observed evolutionary pattern, including a strong flattening in the faint-end slope of the total XLF, is consistent with our findings.

Additional evidence in support of our baseline model rather than the no-change scenario may be provided by the data shown in Figure 13 (right). These data are in good agreement with our baseline model, extrapolated to $z = 0$ (although the lowest luminosity bins for the absorbed XLF are systematically above our model). However, in the no-change model, γ_{abs} and γ_{unabs} are driven towards the values required by our higher redshift data. Thus, the

extrapolation of the no-change model to $z = 0$ does not provide a good agreement with the Burlon et al. (2011) measurements, lending support to the baseline model. We note that in both cases our total XLF is in good agreement with the Ueda et al. (2011) $z = 0$ measurements (see Figure 13 left). Incorporating these data sets into our Bayesian analysis remains beyond the scope of this paper.

In conclusion, our own analysis requires that γ_{unabs} is flatter than γ_{abs} at any given redshift but does not provide definitive evidence that the individual slopes change with redshift. However, direct measurements of the XLFs of unabsorbed and absorbed AGNs at $z = 0$ indicate that redshift evolution of γ_{unabs} and γ_{abs} is required when compared to our higher redshift study. Regardless of whether γ_{unabs} and γ_{abs} are changing, the apparent faint-end slope of the *total* XLF does flatten with increasing redshift. This behaviour is primarily driven by the changing mix of unabsorbed and absorbed AGNs.

6.3 Evolution of the absorbed fraction of AGNs

An increase in the fraction of absorbed AGNs with redshift (at a fixed luminosity) is now well established (e.g. La Franca et al. 2005; Hasinger 2008; Treister et al. 2009). Our model reproduces such a behaviour (see Figure 15, top). However, we are able to attribute the underlying cause of this evolution to the luminosity evolution the XLFs of both the unabsorbed and absorbed AGNs. This luminosity evolution shifts the XLFs of both populations to higher luminosities at higher redshifts, and thus the transition between the lower luminosity regime where absorbed AGNs dominate and the

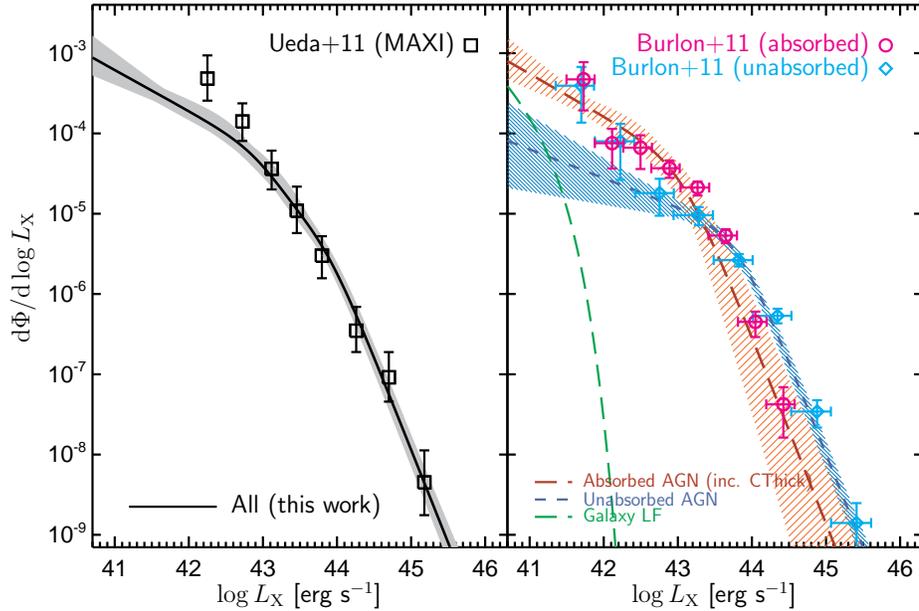


Figure 13. *Left:* Our best fit model for the total XLF of AGNs (black line) extrapolated to $z = 0$, compared to measurements of the local 2–10 keV XLF of AGNs from Ueda et al. (2011) using MAXI (black squares). *Right:* Our best fit model for the XLF of unabsorbed (blue short-dashed line) and absorbed (red long-dashed line) AGNs extrapolated to $z = 0$, compared to measurements from Burlon et al. (2011) from *Swift*/BAT (converted to 2–10 keV luminosities assuming $\Gamma = 1.9$) where the populations were divided based on X-ray spectral fits (cyan diamonds and pink circles for unabsorbed and absorbed AGNs, respectively). The grey, blue and red regions indicate the 99 per cent confidence interval in our model extrapolation based on the posterior distribution of model parameters. Our model extrapolation, based on predominantly higher redshift samples and indirectly inferring the absorption distribution, is generally in good agreement with these local measurements. Our results indicate that the steep slope of the total XLF at $z \sim 0$ observed by Ueda et al. (2011, see also Ballantyne 2014) may be due the gradual change in the mix of absorbed and unabsorbed AGNs. However, our model does not predict the increase in the unabsorbed AGN XLF at $L_X \lesssim 10^{43} \text{ erg s}^{-1}$ found by Burlon et al. (2011), which leads to a decreased absorbed fraction at low luminosities.

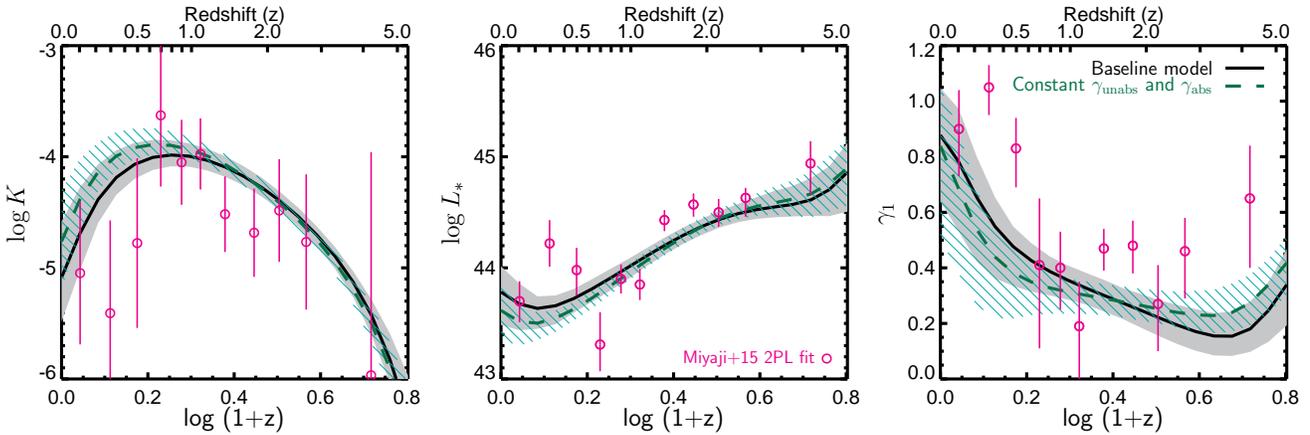


Figure 14. The apparent double power-law shape parameters of the *total* XLF of AGN as a function of redshift. These estimates are based on a simple least-squares fit at various redshifts and in the luminosity range $42.5 < \log L_X < 46.5$ to the predictions of our model (which describes the XLF as the combination of independent XLFs of the unabsorbed and absorbed AGNs; see Figure 10 for the redshift dependence of the parameters of the individual XLFs that are directly constrained by our data). The solid black line indicates our baseline model where both γ_{unabs} and γ_{abs} of the underlying unabsorbed and absorbed AGN XLFs flatten with redshift, whereas the dashed green line uses our “no-change” model where γ_{unabs} and γ_{abs} remain constant. The shaded regions represent 99 per cent confidence regions in the underlying models. In both cases, the apparent faint-end slope (γ_1) undergoes a strong flattening with increasing redshift. The similarity in the parameters describing the total XLF for both models illustrates how distinguishing between these models is difficult without additional information. Pink data points show estimates of the double power-law shape parameters of the total XLF in thin redshift shells from ? and reveal a similar evolutionary pattern.

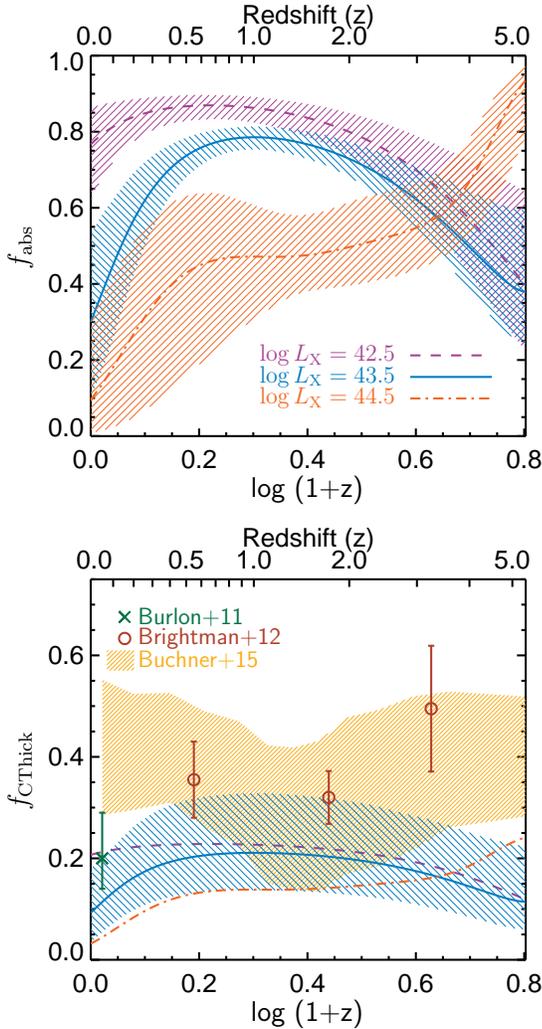


Figure 15. *Top:* Absorbed fraction of Compton-thin AGNs (f_{abs}) versus redshift based on our model, evaluated at three different luminosities. Shaded regions indicate 99 per cent confidence interval in the model parameters. The absorbed fraction changes most rapidly at $\log L_X = 43.5$ due to the evolution of L_* for the unabsorbed and absorbed AGN XLFs. At high redshifts ($z \gtrsim 3$) we have tentative evidence of a drop in f_{abs} at low and moderate luminosities and a rise at higher luminosities. *Bottom:* Fraction of Compton-thick AGNs (relative to all AGNs at a given luminosity) versus redshift from our model. For clarity, we omit the confidence regions for the low and high luminosities. We compare to the estimates of the Compton-thick fraction from Burlon et al. (2011, green crosses), Brightman & Ueda (2012, red circles), and Buchner et al. (2015, orange shaded region indicates 90 per cent confidence interval) that used X-ray spectral fitting to directly identify Compton-thick AGNs. While our estimate of f_{CThick} at $\log L_X = 43.5$ is systematically lower than prior works, we generally agree within the uncertainties.

higher luminosity regime where unabsorbed AGNs dominate shifts to higher luminosities at higher redshifts (see Figure 11). Thus, close to the break luminosities of the XLFs ($L_X \sim 10^{43.5-44.5}$ erg s $^{-1}$) a strong evolution in the absorbed fraction is observed, whereas at higher and lower luminosities the evolution is weaker (see Figure 15, top).

Our results also suggest (albeit at low significance) that the absorbed fraction at high luminosities may increase slightly between $z \sim 0$ and $z \sim 1$ as the overall normalization of the absorbed AGN XLF increases more rapidly than the normalization

of the unabsorbed XLF. While such a redshift dependence is much weaker than the strong evolution around L_* , it may explain why some studies have found substantial populations of obscured, luminous AGNs (Type-2 QSOs) at these high redshifts (e.g. Brusa et al. 2010; Stern et al. 2012; Banerji et al. 2015). The normalization of the absorbed AGN XLF peaks at $z \approx 1$ and declines towards higher redshifts. By $z \sim 3$ the continuing, strong luminosity evolution of the absorbed AGNs appears to take over, again leading to a high absorbed fraction at the highest luminosities (conversely the absorbed fraction of lower luminosity AGNs may decline at $z \gtrsim 2$). It should be noted, however, that our best-fitting model is primarily constrained by lower redshift data, where the bulk of our sample lies. The high-redshift behaviour may simply reflect an extrapolation of this model to higher redshifts where our observational constraints are poor. Nevertheless, our findings do hint at a potentially complex evolution in the absorbed fraction of high luminosity AGNs.

6.4 The N_{H} distribution, the fraction of Compton-thick AGNs and their contribution to the cosmic X-ray background

While we do not directly estimate absorption column densities for individual sources, our study does predict an overall distribution of N_{H} based on our statistical approach. In Figure 16 we show our inferred distribution of N_{H} (the N_{H} function) evaluated at $z = 0.05$ and $\log L_X = 43.5$. In the left panel, we compare to estimates of the N_{H} function from U14, based on direct measurements of N_{H} in the *Swift*/BAT local AGN sample (Tueller et al. 2008; Ichikawa et al. 2012) and corrected for absorption-dependent biases. Our inferred N_{H} function has a similar overall shape to the U14 measurements, dropping between $\log N_{\text{H}} \sim 20 - 21$ and $\log N_{\text{H}} \sim 22 - 23$ and peaking again at $\log N_{\text{H}} \approx 23 - 24$. Our absorbed fraction (f_{abs}) is slightly lower at this specific luminosity and redshift, thus our N_{H} function predicts slightly more unabsorbed ($\log N_{\text{H}} = 20 - 22$) sources and fewer absorbed ($\log N_{\text{H}} = 22 - 24$) sources. Furthermore, our prediction of the N_{H} function in the Compton-thick regime ($\log N_{\text{H}} > 24$) is lower than the U14 measurements and corresponds to an overall Compton-thick fraction (f_{CThick}) at $z = 0.05$ and $\log L_X = 43.5$ of just $11.4^{+7.7}_{-5.6}$ per cent. Even so, the differences between our model and the *Swift*/BAT measurements from U14 are not significant.

The right panel of Figure 16 compares our results to the N_{H} functions assumed in various AGN population synthesis models, used to estimate the integrated cosmic X-ray background (CXB). While the overall shape of our N_{H} -function is similar, there are significant differences from prior studies. The ?, hereafter G07 and Treister et al. (2009, hereafter T09) models predict fewer AGNs with $N_{\text{H}} \approx 10^{21-22}$ cm $^{-2}$, although neither our study nor the CXB models are particularly sensitive to the exact distribution in this range. More importantly, our N_{H} -function is significantly different in the Compton-thick regime. The normalisation is significantly lower than the G07 or U14 models across the $N_{\text{H}} = 10^{20-26}$ cm $^{-2}$ range, ruling out these distributions (and the correspondingly higher f_{CThick}) at the > 99 per cent confidence level. The T09 N_{H} -function, on the other hand, is higher than our model at $N_{\text{H}} = 10^{24-25}$ cm $^{-2}$ but assumes no sources with $N_{\text{H}} > 10^{25}$ cm $^{-2}$; thus, their overall f_{CThick} is consistent with our estimates where we assume a constant distribution up to $N_{\text{H}} = 10^{26}$ cm $^{-2}$. The exact shape of the N_{H} function in the Compton-thick regime, the relative number of $N_{\text{H}} \gtrsim 10^{25}$ cm $^{-2}$ AGNs, and the precise X-ray spectral properties of such sources remain major uncertain-

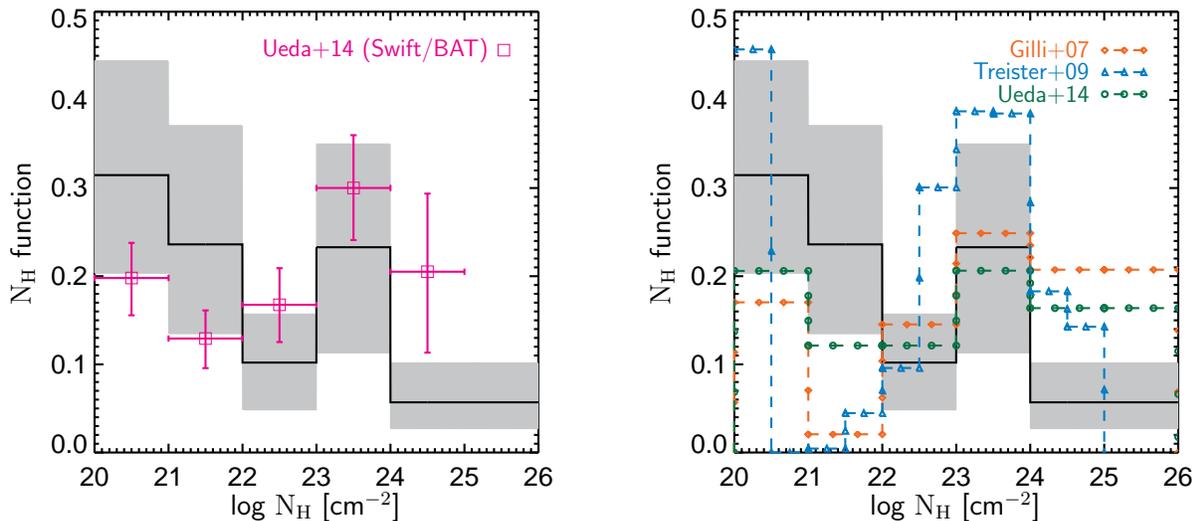


Figure 16. Our inferred distribution of absorption column densities (N_{H}) evaluated at $z = 0.05$ and $\log L_{\text{X}} = 43.5$ (black line, shaded region indicates 99 per cent confidence interval in the model) and compared to direct measurements based on X-ray spectral fits of the *Swift*/BAT AGNs (*left*: Ueda et al. 2014) and that assumed in previous CXB synthesis models (*right*: ?Treister et al. 2009; Ueda et al. 2014).

ties in estimates of f_{CThick} and could account for the differences between our estimates and the prior studies shown in Figure 16 (right). For example, if the most heavily Compton-thick sources ($\log N_{\text{H}} \gtrsim 25$) are also deeply buried and do not exhibit the same level of soft, scattered emission as less absorbed sources (as suggested by some studies e.g. Brightman et al. 2014) then our measurements would be biased to below the true value.

Figure 15 (bottom) also compares our model for f_{CThick} to estimates from Burlon et al. (2011), Brightman & Ueda (2012) and B15 that used X-ray spectral fitting to directly identify Compton-thick AGNs in samples from *Swift*/BAT, *Chandra*, or *XMM-Newton* respectively. While our model is systematically lower than all of these measurements, it is generally consistent given the uncertainties in our model and the data. The B15 estimates are significantly higher at low ($z \lesssim 0.2$) and high ($z \gtrsim 3$) redshifts, although these regimes are not well probed by their samples and thus the estimates may be partly driven by their priors. In the key $0.2 < z < 3$ range, where the B15 estimates are most reliable, our estimates are statistically consistent, although the B15 confidence intervals do allow for the potential of a higher f_{CThick} .

It is worth noting that our inferences regarding the Compton-thick population are *indirect*. Our estimates of the underlying distribution of N_{H} rely on global comparisons of the hard- and soft-band samples and the changing sensitivities of these energy bands to different levels of absorption as a function of redshift. With this approach, our ability to distinguish between the effects of Compton-thick sources and heavily absorbed, yet Compton-thin ($N_{\text{H}} \approx 10^{23-24} \text{ cm}^{-2}$) sources is limited, which could affect our estimates of f_{CThick} . We also assume that the XLF of Compton-thick AGNs traces the absorbed AGN XLF exactly (apart from a single scale factor) at all redshifts and assume a completely flat distribution over the $N_{\text{H}} = 10^{24-26} \text{ cm}^{-2}$ range. Given these issues, our estimates of the Compton-thick fraction should be treated with caution.

As an additional consistency check of our results, in Figure 17 we use our model to estimate the overall contribution of AGNs to the CXB, retaining our assumed X-ray spectral model described in Section 3.1. Our model prediction for the CXB is in good agree-

ment with the measurements at energies > 10 keV that we do not directly probe with our source samples (we discuss the low energy CXB and the contribution of galaxies in Section 6.6 below). Despite our low Compton-thick fraction compared to U14 or G07, we are still able to reproduce the observed peak in the CXB at $\sim 20 - 30$ keV. At these energies, we find that Compton-thick AGNs make a small, but important contribution to the integrated CXB emission.

Part of the reason for our successful reproduction of the CXB may be the comparatively strong reflection that we allow for in unabsorbed ($\log N_{\text{H}} < 22$) and moderately absorbed ($\log N_{\text{H}} \approx 22 - 23$) sources. The extent and strength of reflection in such sources, especially those with $\gtrsim L_*$ luminosities and at $z \gtrsim 1$ that make a large contribution to the CXB flux, is still a matter of debate (e.g. Vasudevan et al. 2013; Del Moro et al. 2014). Indeed, the CXB can be reproduced with a wide variety of models that assume different X-ray spectral models for the constituent AGNs, as well as different XLFs, N_{H} distributions, and—perhaps most crucially—Compton-thick fractions (e.g. Draper & Ballantyne 2009; Ballantyne et al. 2011; Akylas et al. 2012). While fully exploring CXB synthesis models for different spectral models and Compton-thick fractions is beyond the scope of this work, Figure 17 does indicate that our relatively low Compton-thick fraction should not be ruled out.

6.5 What drives the absorption-dependent evolution of the AGN population?

The shape and evolution of the XLF of AGNs provides important constraints on the evolution of the underlying physical properties of growing SMBHs, namely their masses and Eddington ratios (a tracer of the rate of accretion relative to the SMBH mass). Our study—revealing the independent XLFs of unabsorbed and absorbed AGNs out to high redshifts—provides further insights into how the evolution of the AGN population relates to their absorption properties.

We find that, at least to $z \sim 2$, the evolution of the unabsorbed XLF and the evolution of the absorbed XLF are fairly sim-

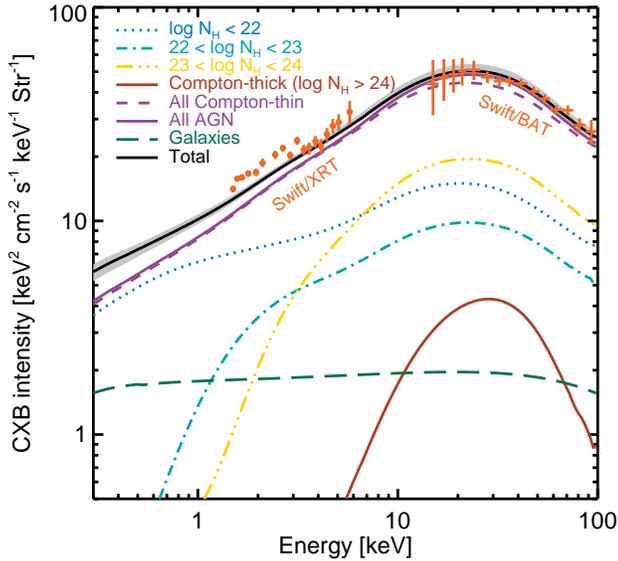


Figure 17. Predicted Cosmic X-ray Background (CXB) spectrum based on our final model of the XLAF and GLF (black solid line, grey region indicates 99 per cent confidence interval) compared to recent measurements from *Swift*/XRT (orange circles at < 7 keV, Moretti et al. 2009) and *Swift*/BAT (orange crosses at > 10 keV, Ajello et al. 2008). The blue dotted, light-green dot-dashed, yellow triple-dot-dashed, and dark-red solid lines show the contributions from AGNs with different absorption column densities. The dashed purple line indicates the contribution of all Compton-thin ($N_{\text{H}} < 10^{24} \text{ cm}^{-2}$) AGNs, while the solid purple line is the total AGN contribution including Compton-thick sources. The contribution of Compton-thick AGNs in our model to the peak of the CXB at ~ 30 keV is minimal but does bring the model closer to the observational data. The dark green long-dashed line indicates the contribution of normal galaxies to the CXB (based on our estimates of the GLF) and accounts for $\gtrsim 10$ per cent of the total CXB flux in our model at $\lesssim 3$ keV.

ilar. The bulk of the change in both XLFs is driven by moderately strong luminosity evolution, which could be attributed to a change in the underlying distribution of SMBH masses, whereby the average mass of an accreting SMBH is higher at earlier cosmic times. Alternatively, the luminosity evolution may reflect an increase in the average accretion rates with increasing redshift. Both XLFs also undergo an overall density evolution, whereby their normalizations increase between $z = 0$ and $z \sim 1$ then drop rapidly again at higher redshifts. This evolutionary pattern may reflect changes in the rate at which AGN activity is triggered within the galaxy population. Finally, the faint-end slopes of both XLFs may also be changing with redshift. Such changes could indicate that the underlying distribution of Eddington ratios is changing, with more rapidly accreting sources become relatively more predominant at higher redshifts. However, the observed flattening of the total XLF may be more closely related to the changing mix of unabsorbed and absorbed AGNs (see Section 6.2 above).

While the overall evolution of the XLFs of unabsorbed and absorbed AGNs follow similar patterns, the individual XLFs are significantly different at a given redshift. The most important differences (at least at $z \lesssim 2$) appear to be 1) the lower L_* of the absorbed AGNs, and 2) the steeper faint-end slope of the absorbed AGN XLF. These effects combine with the overall evolution of the XLFs to produce the luminosity and redshift dependence of the absorbed fraction seen in Figure 11.

The differences in L_* may reflect the life cycle of AGN activ-

ity. Luminous, unabsorbed AGNs could correspond to short periods in high accretion rate phases that have the power to eject material from the galactic nucleus. Exhaustion of the fuel supply then leads to a comparatively rapid fading of the AGN, that is traced by the faint end of the unabsorbed AGN XLF. Absorbed AGN, conversely, may correspond to longer periods when the SMBH grows at lower accretion rates but has little impact on the surrounding obscuring material. The two phases could correspond to separate AGN events with different triggering mechanisms (e.g. major mergers versus secular processes within a host galaxy). Changes in the extent and efficiency of different triggering mechanisms could then lead to the slight differences in the evolution of the unabsorbed and absorbed XLFs at $z \lesssim 2$. Alternatively, the unabsorbed and absorbed phases could reflect long-term variations related to a single AGN triggering event.

At $z \gtrsim 2$, our results hint at a more complex evolutionary behaviour. The luminosity evolution of the unabsorbed XLF appears to slow at $z \sim 2 - 3$ and then decline to higher redshifts, whereas the absorbed XLF continues to evolve positively in L_* but the normalization (K) is rapidly declining (see Figure 10). This pattern suggests that absorbed growth phases continue to shift to higher accretion rates and/or SMBH masses at high redshifts, yet the triggering of such events becomes increasingly rarer. Conversely, the most luminous, unabsorbed growth phases may be suppressed at the earliest times. However, our power to distinguish between the unabsorbed and absorbed populations is limited at the highest redshifts.

The combined evolution of the unabsorbed and absorbed XLFs across our entire redshift range results in strong, luminosity-dependent evolution in the overall space densities of AGNs (see Figure 18). Higher luminosity AGNs peak, in terms of their space densities, at higher redshifts than lower luminosity AGNs. A similar pattern has been seen in a number of previous works, most recently by M15 (data points in Figure 18). The minor differences between the M15 estimates and our results are likely due to the different modelling of the underlying XLF and the correction for absorption.¹³ This pattern in space densities, described as “downsizing” of AGN luminosities (e.g. Ueda et al. 2003; Barger et al. 2005), is often attributed purely to changes in the average SMBH mass with redshift. However, our constraints on the underlying evolution of the XLFs indicate that the observed pattern in the evolution of total space densities (Figure 18) may be due to a combination of changes in SMBH mass, the distributions of Eddington ratios, and the life cycles of obscured and unobscured AGN phases.

6.6 The contribution of normal galaxies to the observed luminosity function and the cosmic X-ray background

An important advance in our work is the inclusion of normal, X-ray detected galaxies in our analysis of the XLF. Instead of applying a strict luminosity cut (e.g. $L_X > 10^{42} \text{ erg s}^{-1}$), we develop an approach to statistically classify X-ray sources as galaxies or AGNs. We find that AGNs still dominate the space density of X-ray selected sources at $L_X \approx 10^{42} \text{ erg s}^{-1}$ at low redshifts. At $z \approx 0.1$ the space density of galaxies is a factor ~ 1000 lower than the space density of AGNs at $L_X = 10^{42} \text{ erg s}^{-1}$. Galaxies only begin to make a substantial contribution (> 10 per cent in space density)

¹³ M15 use the N_{H} -function from U14 to correct for absorption effects (as a function of L_X and z), which differs from our results.

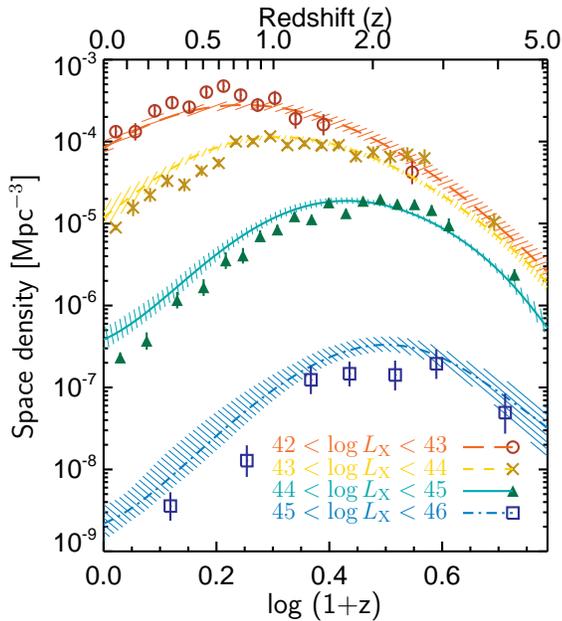


Figure 18. Total space density of AGNs for different ranges of X-ray luminosity based on our model (coloured lines). Shaded regions indicate the 99 per cent confidence interval in our model parameters. We see clear “downsizing” in terms of AGN luminosity, whereby higher luminosity AGNs peak in terms of their space density at higher redshifts than lower luminosity sources. The data points are taken from the recent work of ? and reveal the same pattern; the small discrepancies with our model are most likely due to differences in the modelling of the XLF and the N_{H} distribution.

at $L_X \lesssim 2.5 \times 10^{41} \text{ erg s}^{-1}$ and start to dominate (> 50 per cent in space density) at $L_X \lesssim 2 \times 10^{40} \text{ erg s}^{-1}$.

We find that the GLF undergoes strong luminosity evolution up to $z \approx 0.8$, whereby the characteristic luminosity increases as approximately $(1+z)^{2.7}$ (consistent with previous studies of the evolution of the GLF at X-ray wavelengths, e.g. Ptak et al. 2007; Georgakakis et al. 2007; Tzanavaris & Georgantopoulos 2008). This evolution increases the contribution of galaxies at low luminosities. At $z \approx 1$, the space density of galaxies corresponds to ~ 10 per cent of the AGN space density at $L_X = 10^{42} \text{ erg s}^{-1}$ and galaxies dominate at $L_X \lesssim 1.6 \times 10^{41} \text{ erg s}^{-1}$. Careful consideration of contamination by normal galaxies is thus vital in studies of faint AGN at these redshifts.

At higher redshifts, our data are not deep enough to accurately probe the GLF and its evolution. However, our model indicates that by $z \approx 3$ galaxies may dominate over AGNs at $L_X = 10^{42} \text{ erg s}^{-1}$ (see also Laird et al. 2006; Aird et al. 2008). Thus, galaxies may make a substantial contribution to the X-ray population at high redshifts and should be carefully considered in future studies of the high-redshift XLF of AGNs.

Galaxies also make a significant contribution to the CXB at low energies. Previous studies have shown that normal galaxies can account for $\sim 4 - 20$ per cent of the total CXB emission at $\sim 1 - 6 \text{ keV}$ energies, based on stacking of the X-ray data at the optical positions of known galaxies (e.g. Worsley et al. 2006; Hickox & Markevitch 2007; Xue et al. 2012). Here, we can account for the contribution of normal galaxies to the CXB based on our modeling of the resolved X-ray populations. Figure 17 shows that galaxies contribute $\gtrsim 10$ per cent of the total CXB emission at energies $\lesssim 3 \text{ keV}$. However, our estimate of the total CXB remains $\sim 10 - 20$ per cent lower than the recent *Swift*/XRT mea-

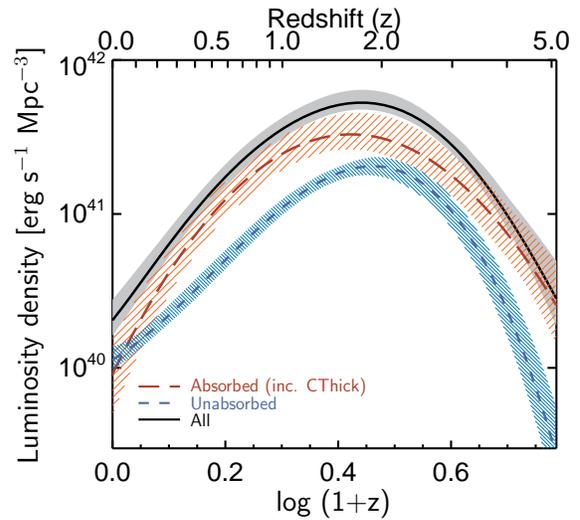


Figure 19. Bolometric luminosity density, a tracer of the total accretion density, for unabsorbed, absorbed (including Compton-thick) and all AGNs as a function of redshift, based on our final model of the XLAF. Shaded regions indicate the 99 per cent confidence interval in our model parameters. The total rate of SMBH growth peaks at $z \sim 2$ and is dominated by absorbed and Compton-thick AGNs.

surements at $\lesssim 3 \text{ keV}$, although measurements of the absolute flux of the CXB vary substantially between missions in this regime, probably due to calibration uncertainties (e.g. Barcons et al. 2000; Moretti et al. 2009). Indeed, our model of the total CXB is in excellent agreement with earlier measurements of the CXB flux with ASCA (Gendreau et al. 1995) over the full $\sim 1 - 7 \text{ keV}$ range. Thus, the significance of any discrepancy is uncertain.

6.7 The accretion history of the Universe

Our new measurements of the XLF of AGNs out to $z \sim 5$ enable us to place improved constraints on the overall accretion history of the Universe. In Figure 19 we plot the total luminosity density of AGNs,

$$\int d \log L_{\text{bol}} \phi_{\text{bol}}(L_{\text{bol}}, z | \Theta) \quad (46)$$

where L_{bol} is the bolometric luminosity and $\phi_{\text{bol}}(L_{\text{bol}}, z | \Theta)$ is the bolometric luminosity function of AGNs, found by integrating our XLAF model over $\log N_{\text{H}}$ and converting the resulting XLF to bolometric values. We adopt the luminosity-dependent bolometric corrections from Hopkins et al. (2007). We find a clear peak in the luminosity density at $z = 1.76 \pm 0.05$, consistent with prior studies (e.g. U14, ?). We also show the luminosity density for the absorbed (including Compton-thick) and unabsorbed populations. At $z \gtrsim 0.1$, the majority of SMBH growth is taking place in absorbed or Compton-thick AGNs. The proportion increases rapidly between $z = 0$ and the overall peak at $z \sim 1 - 2$. The contribution from unabsorbed AGNs increases more gradually and peaks at a slightly higher redshift ($z = 1.89 \pm 0.05$).

By integrating the AGN luminosity density across our redshift range, we can estimate the total SMBH mass density that has been built up by accretion (Soltan 1982). Adopting a radiative efficiency of $\eta = 0.1$, we estimate a relic SMBH mass density of $4.20^{+0.29}_{-0.14} \times 10^5 \mathcal{M}_{\odot} \text{ Mpc}^{-3}$ at $z = 0$, where the error represents the uncertainty in our XLF model parameters. Of this mass

density, ~ 30 per cent has been built up in unabsorbed growth phases, ~ 50 per cent in absorbed AGNs, and a further ~ 20 per cent in Compton-thick sources. Our estimate is higher than in A10 ($2.2 \pm 0.2 \times 10^5 \mathcal{M}_\odot \text{Mpc}^{-3}$), likely due to our improved accounting for absorbed and Compton-thick sources, and is in much better agreement with the estimate by Marconi et al. (2004) based on the velocity dispersions of local galaxies and the $\mathcal{M}_{\text{BH}} - \sigma$ relation ($4.6_{-1.4}^{+1.9} \times 10^5; \mathcal{M}_\odot \text{Mpc}^{-3}$). However, uncertainties in the bolometric correction and radiative efficiency dominate over our XLF model uncertainties. A higher radiative efficiency (e.g. if a large fraction of SMBHs have high spins) would reduce our estimate of the relic SMBH mass density. Alternatively, a large population of deeply buried, highly Compton-thick sources could bring our estimate of the relic SMBH mass density more in line with recent revisions to the $\mathcal{M}_{\text{BH}} - \sigma$ relation (e.g. Comastri et al. 2015).

In Figure 20 we compare estimates of the total star formation rate (SFR) density (from a recent review by Madau & Dickinson 2014) to our estimates of the SMBH accretion density, converted to $\mathcal{M}_\odot \text{yr}^{-1} \text{Mpc}^{-3}$ using the same bolometric corrections and radiative efficiency assumptions described above and scaled up by an additional factor of 1500. We extrapolate our model to $z > 5$ to enable comparisons with estimates of the SFR density in the first ~ 1 Gyr of cosmic time. The SFR density and SMBH accretion rate both peak at $z \sim 2$. However, the total rate of SMBH growth rises slightly more rapidly from $z = 0$ to $z \sim 2$, indicating that the relative efficiency of SMBH growth relative to galaxy growth may change over this period. The SMBH accretion density drops much more rapidly to higher redshifts than the SFR density. Thus, galaxy growth may precede the build up of their central SMBHs in the early Universe, or the relative efficiency with which SMBHs form and grow may be substantially lower. However, our estimates are extrapolations at the highest redshifts. New and ongoing surveys with *Chandra* (e.g. COSMOS-Legacy, P.I. Civano; CDFS-7Ms, P.I. Brandt) may improve constraints in this regime. Nonetheless, obtaining direct and accurate measurements of the XLF of AGNs at the highest redshifts ($z \gtrsim 6$) requires ~ 2 orders of magnitude improvement in survey speed and is one of the principal aims of ESA's next large mission, the *Athena* X-ray observatory (Nandra et al. 2013; Aird et al. 2013b).

7 SUMMARY

In this paper we present new measurements of the XLF of AGNs out to $z \sim 5$. In this section we first summarize our work and then present our conclusions.

We compile large samples of uniformly selected X-ray sources in the hard (2–7 keV) and soft (0.5–2 keV) energy bands from a number of deep and wide *Chandra* surveys (CDFS, CDFN, EGS, COSMOS and Bootes), supplemented by wider area surveys from *ASCA* and *ROSAT*. We also compile multiwavelength photometry across our fields that we use to robustly identify counterparts to our X-ray sources and calculate photometric redshifts. Next, we present a Bayesian methodology that allows us to incorporate a range of X-ray spectral shapes and statistically account for the effects of absorption, as well as accounting for photometric redshift uncertainties, flux uncertainties, and the Eddington bias. We also introduce an approach to statistically account for the contamination of our samples by normal, X-ray detected galaxies. This allows us to place improved constraints on the XLF of AGNs down to much lower luminosities than prior studies.

We present initial measurements of the XLF based on our

hard-band and soft-band samples separately, neglecting absorption effects. We introduce a new, flexible model to describe the XLF of AGNs as an evolving double power-law, where the break luminosity (L_*), normalization (K) and faint-end slope (γ_1) all evolve with redshift. There is significant evidence (based on Bayesian model comparison) favouring this model over luminosity-dependent density evolution (LDDE) parametrizations of the XLF based on either the hard-band or soft-band samples. However, differences in the space densities at low luminosities based on the hard-band compared to the soft-band samples indicate that absorption has a substantial effect and requires further consideration.

To investigate the extent and distribution of absorption, we model the XLFs of unabsorbed ($N_{\text{H}} < 10^{22} \text{cm}^{-2}$) and absorbed ($N_{\text{H}} > 10^{22} \text{cm}^{-2}$) AGNs as independent, flexible double power-law functions. We use the combination of the hard- and soft-band samples to constrain our flexible model, seeking a model that can reconcile both samples assuming no *a priori* knowledge of the extent of absorption for any individual source. We find that the absorbed AGN population has an XLF that generally has a lower break luminosity, a higher normalization, and a steeper faint-end slope than the unabsorbed AGNs at all redshifts out to $z \sim 2$. Both the unabsorbed and absorbed AGN XLFs evolve in luminosity and density. Differences in the extent of the luminosity and density evolution lead to a comparatively complex evolution in the shape of the total XLF of AGNs.

Our work provides new insight into the nature and evolution of the AGN population. Our main conclusions are as follows:

1. Differences between the break luminosities (L_*), normalizations (K), and faint-end slopes (γ_1) of the unabsorbed and absorbed AGN XLFs explain the strong luminosity dependence in the fraction of absorbed AGNs at $z \lesssim 2$. Absorbed AGNs dominate at low luminosities, with the absorbed fraction falling rapidly as luminosity increases above a characteristic value.
2. Both the unabsorbed and absorbed XLFs undergo strong luminosity evolution up to $z \sim 2$, shifting the transition in the absorbed fraction to higher luminosities at higher redshifts. A strong evolution in the absorbed fraction is thus observed at $L_{\text{X}} \sim 10^{43.5-44.5} \text{erg s}^{-1}$, close to the break luminosities of the XLFs. At higher redshifts ($z \gtrsim 2$), the luminosity evolution of the unabsorbed AGN XLF slows and may become negative, although this behaviour may be driven by the model parametrization at lower redshifts.
3. The normalization of the absorbed AGN XLF evolves more rapidly than the normalization of the unabsorbed AGN XLF, peaking at $z \sim 1$ and declining to higher redshifts. This leads to a mild increase in the absorbed fraction at high luminosities ($L_{\text{X}} \sim 10^{45} \text{erg s}^{-1}$) between $z \sim 0$ and $z \sim 1$.
4. The total XLF of AGNs (combining unabsorbed and absorbed populations) undergoes a complicated evolution with redshift, including substantial flattening of the faint-end slope. *This behaviour is primarily driven by the changing mix of unabsorbed and absorbed AGNs.*
5. Our inferred distribution of N_{H} generally exhibits two peaks at $N_{\text{H}} \sim 10^{20-21} \text{cm}^{-2}$ and $N_{\text{H}} \sim 10^{23-24}$, although the exact shape depends on luminosity and redshift. We infer a relatively low fraction of Compton-thick AGNs ($\sim 10 - 20$ per cent at $L_{\text{X}} = 10^{43.5} \text{erg s}^{-1}$). Nevertheless, our model reproduces the peak of the cosmic X-ray background at $\sim 20 - 30$ keV.
6. The majority of SMBH growth takes place in absorbed or

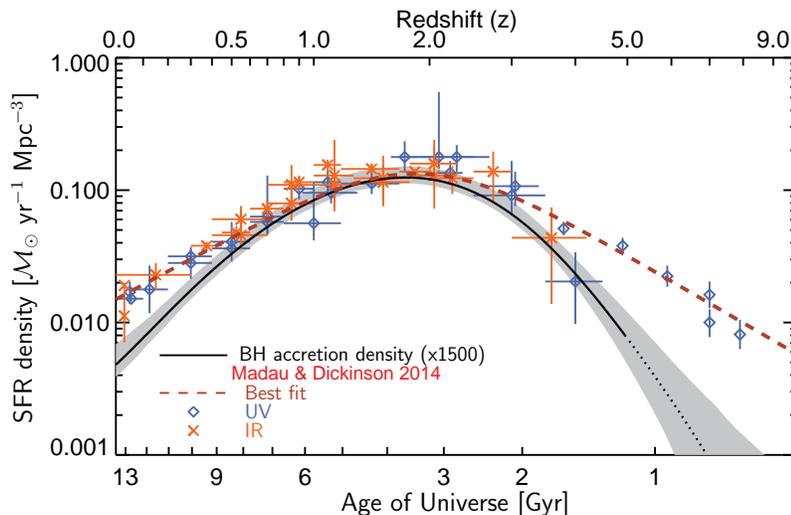


Figure 20. Total star-formation rate (SFR) density compared to our estimate of the total SMBH accretion density. We show the best overall fit (dashed dark-red line) for the evolution of the SFR density from the recent review by Madau & Dickinson (2014) along with their compilation of measurements based on rest-frame ultraviolet (blue diamonds) and infrared (orange crosses) observations. Our estimate of the SMBH accretion density (solid black line) is scaled up by an arbitrary factor of 1500. Shaded regions indicate the 99 per cent confidence interval in our model parameters and may under-represent the overall uncertainties in the bolometric correction, radiative efficiency and observational data at a given redshift. The dotted black line indicates where we are substantially extrapolating our model outside the coverage of our data. Both galaxy and SMBH growth peak at $z \sim 2$. However, the SMBH accretion density appears to evolve more rapidly, with a much stronger decline to higher redshifts.

Compton-thick AGNs at all redshifts $z \gtrsim 0.1$, although the proportion increases rapidly between $z = 0$ and $z \sim 2$. The contribution from unabsorbed AGNs increases more gradually and peaks at a higher redshift ($z \sim 2$). Unabsorbed phases of activity constitute a substantial fraction (~ 30 per cent) of the overall accretion density.

7. At low redshift ($z \approx 0.1$) the space density of X-ray selected galaxies is negligible compared to AGNs at $L_X = 10^{42} \text{ erg s}^{-1}$. At $z \approx 1$, however, galaxies correspond to ~ 10 per cent of the AGN space density at $L_X = 10^{42} \text{ erg s}^{-1}$ and dominate at $L_X \lesssim 1.6 \times 10^{41} \text{ erg s}^{-1}$. Normal galaxies contribute $\gtrsim 10$ per cent of the total integrated CXB emission at energies $\lesssim 3 \text{ keV}$.
8. The total AGN accretion density generally tracks the rise in the overall SFR density from $z \sim 0$ to $z \sim 2$, indicating that SMBH growth and galaxy growth proceed together over this period. The accretion density appears to fall off substantially faster than the SFR density to higher redshifts, indicating that galaxy growth may precede the build up of their central SMBHs in the early Universe or the relative efficiency of SMBH growth may be lower.

Our work provides new insight into the nature of the evolution of the AGN population and the underlying cause of the luminosity and redshift dependence of the absorbed fraction. However, to fully reveal the underlying physical processes that drive both unabsorbed and absorbed growth phases requires further study to connect the distributions of AGN accretion rates and their absorption properties to the key physical properties of the evolving galaxy population.

ACKNOWLEDGEMENTS

We thank Johannes Buchner for helpful discussions and providing machine-readable results. We also acknowledge helpful discus-

sions with Ranjan Vasudevan. The scientific results reported in this article are based to a significant degree on observations made by the *Chandra* X-ray Observatory. This work has made use of the Rainbow Cosmological Surveys Database, which is operated by the Universidad Complutense de Madrid (UCM). This study makes use of data from AEGIS, a multiwavelength sky survey conducted with the *Chandra*, *GALEX*, *Hubble*, Keck, CFHT, MMT, Subaru, Palomar, *Spitzer*, VLA, and other telescopes and supported in part by the NSF, NASA, and the STFC. Funding for the DEEP2 Galaxy Redshift Survey has been provided by NSF grants AST-95-09298, AST-0071048, AST-0507428, and AST-0507483 as well as NASA LTSA grant NNG04GC89G. Funding for the DEEP3 Galaxy Redshift Survey has been provided by NSF grants AST-0808133, AST-0807630, and AST-0806732. This work made use of images and/or data products provided by the NOAO Deep Wide-Field Survey (Jannuzi & Dey 1999), which is supported by the National Optical Astronomy Observatory (NOAO). NOAO is operated by AURA, Inc., under a cooperative agreement with the National Science Foundation. Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org/>. SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration. Based in part on data collected at Subaru Telescope, which is operated by the National Astronomical Observatory of Japan. This work is based in part on observations made with the *Spitzer* Space Telescope, which is operated by the Jet Propulsion Laboratory, California Institute of Technology under a contract with NASA. Some of the data used in this paper were obtained from the Mikulski Archive for Space Telescopes (MAST). STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555. Support for MAST for non-HST data is provided by the NASA Office of Space Science via grant NNX13AC07G and by other grants and contracts. Based in part on observations

made with ESO Telescopes at the La Silla or Paranal Observatories. JA acknowledges support from a COFUND Junior Research Fellowship from the Institute of Advanced Study, Durham University, and ERC Advanced Grant FEEDBACK at the University of Cambridge. ALC acknowledges support from NSF CAREER award AST-1055081. AG acknowledges the THALES project 383549 that is jointly funded by the European Union and the Greek Government in the framework of the programme “Education and lifelong learning”. PGP-G acknowledges support from the Spanish Government through MINECO Grant AYA2012-31277.

REFERENCES

- Ahn, C. P., et al. 2012, *ApJS*, 203, 21
Aird, J., et al. 2012, *ApJ*, 746, 90
— 2013a, *ApJ*, 775, 41
— 2013b, An Athena+ Supporting Paper (arXiv:1306.2325)
Aird, J., Nandra, K., Georgakakis, A., Laird, E. S., Steidel, C. C., & Sharon, C. 2008, *MNRAS*, 387, 883
Aird, J., et al. 2010, *MNRAS*, 401, 2531
Ajello, M., et al. 2008, *ApJ*, 689, 666
Akiyama, M., et al. 2000, *ApJ*, 532, 700
Akiyama, M., Ueda, Y., Ohta, K., Takahashi, T., & Yamada, T. 2003, *ApJS*, 148, 275
Akylas, A., Georgakakis, A., Georgantopoulos, I., Brightman, M., & Nandra, K. 2012, *A&A*, 546, A98
Akylas, A., Georgantopoulos, I., Georgakakis, A., Kitsionas, S., & Hatziminaoglou, E. 2006, *A&A*, 459, 693
Alexander, D. M., et al. 2003, *AJ*, 126, 539
Antonucci, R. 1993, *ARA&A*, 31, 473
Appenzeller, I., et al. 1998, *ApJS*, 117, 319
Ashby, M. L. N., et al. 2009, *ApJ*, 701, 428
Assef, R. J., et al. 2011, *ApJ*, 728, 56
Ballantyne, D. R. 2014, *MNRAS*, 437, 2845
Ballantyne, D. R., Draper, A. R., Madsen, K. K., Rigby, J. R., & Treister, E. 2011, *ApJ*, 736, 56
Banerji, M., Alaghband-Zadeh, S., Hewett, P. C., & McMahon, R. G. 2015, *MNRAS*, 447, 3368
Barcons, X., Mateos, S., & Ceballos, M. T. 2000, *MNRAS*, 316, L13
Barger, A. J., et al. 2003, *AJ*, 126, 632
Barger, A. J., Cowie, L. L., Mushotzky, R. F., Yang, Y., Wang, W.-H., Steffen, A. T., & Capak, P. 2005, *AJ*, 129, 578
Barger, A. J., Cowie, L. L., & Wang, W.-H. 2008, *ApJ*, 689, 687
Barro, G., et al. 2011a, *ApJS*, 193, 13
— 2011b, *ApJS*, 193, 30
Bauer, F. E., Alexander, D. M., Brandt, W. N., Schneider, D. P., Treister, E., Hornschemeier, A. E., & Garmire, G. P. 2004, *AJ*, 128, 2048
Berger, J., & Jefferys, W. 1992, *Statistical Methods and Applications*, 1, 17
Bongiorno, A., et al. 2012, *MNRAS*, 427, 3103
Brammer, G. B., van Dokkum, P. G., & Coppi, P. 2008, *ApJ*, 686, 1503
Brand, K., et al. 2006, *ApJ*, 641, 140
Brightman, M., & Nandra, K. 2011a, *MNRAS*, 413, 1206
— 2011b, *MNRAS*, 414, 3084
Brightman, M., Nandra, K., Salvato, M., Hsu, L.-T., Aird, J., & Rangel, C. 2014, *MNRAS*, 443, 1999
Brightman, M., & Ueda, Y. 2012, *MNRAS*, 423, 702
Brusa, M., et al. 2010, *ApJ*, 716, 348
— 2009, *ApJ*, 693, 8
— 2007, *ApJS*, 172, 353
Buchner, J., et al. 2015, *ApJ*, in press
— 2014, *A&A*, 564, A125
Burlon, D., Ajello, M., Greiner, J., Comastri, A., Merloni, A., & Gehrels, N. 2011, *ApJ*, 728, 58
Calzetti, D., Armus, L., Bohlin, R. C., Kinney, A. L., Koornneef, J., & Storchi-Bergmann, T. 2000, *ApJ*, 533, 682
Capak, P., et al. 2007, *ApJS*, 172, 99
— 2004, *AJ*, 127, 180
Cardamone, C. N., Urry, C. M., Schawinski, K., Treister, E., Brammer, G., & Gawiser, E. 2010a, *ApJ*, 721, L38
Cardamone, C. N., et al. 2010b, *ApJS*, 189, 270
Ciliegi, P., Zamorani, G., Hasinger, G., Lehmann, I., Szokoly, G., & Wilson, G. 2003, *A&A*, 398, 901
Civano, F., et al. 2011, *ApJ*, 741, 91
— 2012, *ApJS*, 201, 30
Coil, A. L., et al. 2011, *ApJ*, 741, 8
— 2009, *ApJ*, 701, 1484
Comastri, A., Gilli, R., Marconi, A., Risaliti, G., & Salvati, M. 2015, *ArXiv e-prints*
Comastri, A., et al. 2011, *A&A*, 526, L9
Cooper, M. C., et al. 2011, *ApJS*, 193, 14
— 2012a, *MNRAS*, 419, 3018
— 2012b, *MNRAS*, 425, 2116
Cowie, L. L., Barger, A. J., Hu, E. M., Capak, P., & Songaila, A. 2004, *AJ*, 127, 3137
Davis, M., et al. 2007, *ApJ*, 660, L1
Del Moro, A., et al. 2014, *ApJ*, 786, 16
Dickey, J. M., & Lockman, F. J. 1990, *ARA&A*, 28, 215
Donley, J. L., Rieke, G. H., Pérez-González, P. G., & Barro, G. 2008, *ApJ*, 687, 111
Draper, A. R., & Ballantyne, D. R. 2009, *ApJ*, 707, 778
Dwelly, T., & Page, M. J. 2006, *MNRAS*, 372, 1755
Ebrero, J., et al. 2009, *A&A*, 493, 55
Elston, R. J., et al. 2006, *ApJ*, 639, 816
Elvis, M., et al. 2009, *ApJS*, 184, 158
Fan, X., et al. 2006, *AJ*, 132, 117
Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, 398, 1601
Fioc, M., & Rocca-Volmerange, B. 1997, *A&A*, 326, 950
Fiore, F., et al. 2012, *A&A*, 537, A16
Fischer, J.-U., Hasinger, G., Schwobe, A. D., Brunner, H., Boller, T., Trümper, J., Voges, W., & Neizvestnyj, S. 1998, *Astronomische Nachrichten*, 319, 347
García, J., Dausser, T., Reynolds, C. S., Kallman, T. R., McClintock, J. E., Wilms, J., & Eikmann, W. 2013, *ApJ*, 768, 146
Gawiser, E., et al. 2006, *ApJS*, 162, 1
Gendreau, K. C., et al. 1995, *PASJ*, 47, L5
Georgakakis, A., Nandra, K., Laird, E. S., Aird, J., & Trichas, M. 2008, *MNRAS*, 388, 1205
Georgakakis, A., et al. 2014, *MNRAS*, 440, 339
Georgakakis, A., Rowan-Robinson, M., Babbedge, T. S. R., & Georgantopoulos, I. 2007, *MNRAS*, 377, 203
Georgakakis, A. E., Chavushyan, V., Plionis, M., Georgantopoulos, I., Koulouridis, E., Leonidaki, I., & Mercado, A. 2006, *MNRAS*, 367, 1017
Georgantopoulos, I., et al. 2013, *A&A*, 555, A43
Georgantopoulos, I., Georgakakis, A., Rowan-Robinson, M., & Rovilos, E. 2008, *A&A*, 484, 671
Glikman, E., Helfand, D. J., White, R. L., Becker, R. H., Gregg, M. D., & Lacy, M. 2007, *ApJ*, 667, 673
Goulding, A. D., et al. 2012, *ApJS*, 202, 6
Grazian, A., et al. 2006, *A&A*, 449, 951
Hasinger, G. 2008, *A&A*, 490, 905
Hasinger, G., Miyaji, T., & Schmidt, M. 2005, *A&A*, 441, 417
Hickox, R. C., & Markevitch, M. 2007, *ApJ*, 661, L117
Hickox, R. C., Mullaney, J. R., Alexander, D. M., Chen, C.-T. J., Civano, F. M., Goulding, A. D., & Hainline, K. N. 2014, *ApJ*, 782, 9
Hiroi, K., Ueda, Y., Akiyama, M., & Watson, M. G. 2012, *ApJ*, 758, 49
Hopkins, P. F., & Hernquist, L. 2006, *ApJS*, 166, 1
Hopkins, P. F., Hernquist, L., Cox, T. J., Robertson, B., Di Matteo, T., & Springel, V. 2006, *ApJ*, 639, 700
Hopkins, P. F., Richards, G. T., & Hernquist, L. 2007, *ApJ*, 654, 731
Hsu, L.-T., et al. 2014, *ApJ*, 796, 60
Ichikawa, K., Ueda, Y., Terashima, Y., Oyabu, S., Gandhi, P., Matsuta, K., & Nakagawa, T. 2012, in *Torus Workshop*, 2012, ed. R. Mason, A. Alonso-Herrero, & C. Packham, 109

Ilbert, O., et al. 2009, *ApJ*, 690, 1236
 Jannuzi, B. T., & Dey, A. 1999, in *Astronomical Society of the Pacific Conference Series*, Vol. 191, *Photometric Redshifts and the Detection of High Redshift Galaxies*, ed. R. Weymann, L. Storrie-Lombardi, M. Sawicki, & R. Brunner, 111
 Jeffreys, H. 1961, *Theory of probability* (Clarendon Press)
 Kass, R. E., & Raftery, A. E. 1995, *Journal of the American Statistical Association*, 90, 773
 Kocevski, D. D., et al. 2012, *ApJ*, 744, 148
 Kochanek, C. S., et al. 2012, *ApJS*, 200, 8
 La Franca, F., et al. 2005, *ApJ*, 635, 864
 Lacy, M., Ridgway, S. E., Sajina, A., Petric, A. O., Gates, E. L., Urrutia, T., & Storrie-Lombardi, L. J. 2015, *ArXiv e-prints*
 Laird, E. S., et al. 2009, *ApJS*, 180, 102
 Laird, E. S., Nandra, K., Hobbs, A., & Steidel, C. C. 2006, *MNRAS*, 373, 217
 Lasker, B. M., et al. 2008, *AJ*, 136, 735
 Lawrence, A. 1991, *MNRAS*, 252, 586
 Lehmer, B. D., et al. 2005, *ApJS*, 161, 21
 —. 2012, *ApJ*, 752, 46
 Lilly, S. J., et al. 2009, *ApJS*, 184, 218
 Lilly, S. J., Le Fevre, O., Crampton, D., Hammer, F., & Tresse, L. 1995, *ApJ*, 455, 50
 Lored, T. J. 2004, in *American Institute of Physics Conference Series*, Vol. 735, *American Institute of Physics Conference Series*, ed. R. Fischer, R. Preuss, & U. V. Toussaint, 195–206
 Luo, B., et al. 2010, *ApJS*, 187, 560
 Madau, P., & Dickinson, M. 2014, *ARA&A*, 52, 415
 Magdziarz, P., & Zdziarski, A. A. 1995, *MNRAS*, 273, 837
 Marconi, A., Risaliti, G., Gilli, R., Hunt, L. K., Maiolino, R., & Salvati, M. 2004, *MNRAS*, 351, 169
 Matsuoka, M., et al. 2009, *PASJ*, 61, 999
 McCracken, H. J., et al. 2010, *ApJ*, 708, 202
 Mendez, A. J., et al. 2013, *ApJ*, 770, 40
 Mineo, S., Gilfanov, M., Lehmer, B. D., Morrison, G. E., & Sunyaev, R. 2014, *MNRAS*, 437, 1698
 Miyaji, T., Hasinger, G., & Schmidt, M. 2000, *A&A*, 353, 25
 —. 2001, *A&A*, 369, 49
 Moretti, A., et al. 2009, *Astronomy & Astrophysics*, 493, 501
 Moretti, A., Vattakunnel, S., Tozzi, P., Salvaterra, R., Severgnini, P., Fugazza, D., Haardt, F., & Gilli, R. 2012, *A&A*, 548, A87
 Murphy, K. D., & Yaqoob, T. 2009, *MNRAS*, 397, 1549
 Murray, S. S., et al. 2005, *ApJS*, 161, 1
 Nandra, K., et al. 2013, (arXiv:1306.2307)
 —. 2005, *MNRAS*, 356, 568
 —. 2014, *ApJ*, submitted
 Nandra, K., O’Neill, P. M., George, I. M., & Reeves, J. N. 2007, *MNRAS*, 382, 194
 Polletta, M., et al. 2007, *ApJ*, 663, 81
 Ptak, A., Mobasher, B., Hornschemeier, A., Bauer, F., & Norman, C. 2007, *ApJ*, 667, 826
 Puccetti, S., et al. 2009, *ApJS*, 185, 586
 Ranalli, P., Comastri, A., & Setti, G. 2003, *A&A*, 399, 39
 Reddy, N. A., Steidel, C. C., Erb, D. K., Shapley, A. E., & Pettini, M. 2006, *ApJ*, 653, 1004
 Ross, N. P., et al. 2013, *ApJ*, 773, 14
 Ross, R. R., & Fabian, A. C. 2005, *MNRAS*, 358, 211
 Salvato, M., et al. 2009, *ApJ*, 690, 1250
 —. 2011, *ApJ*, 742, 61
 Sanders, D. B., et al. 2007, *ApJS*, 172, 86
 Schulze, A., et al. 2015, *MNRAS*, 447, 2085
 Schwoppe, A., et al. 2000, *Astronomische Nachrichten*, 321, 1
 Shankar, F., Weinberg, D. H., & Miralda-Escudé, J. 2013, *MNRAS*, 428, 421
 Silverman, J. D., et al. 2008, *ApJ*, 679, 118
 —. 2010, *ApJS*, 191, 124
 Simpson, C. 2005, *MNRAS*, 360, 565
 Skilling, J. 2004, in *AIP Conference Proceedings of the 24th International*

Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Vol. 735, 395–405
 Soltan, A. 1982, *MNRAS*, 200, 115
 Steidel, C. C., Adelberger, K. L., Shapley, A. E., Pettini, M., Dickinson, M., & Giavalisco, M. 2003, *ApJ*, 592, 728
 Stern, D., et al. 2012, *ApJ*, 753, 30
 Tozzi, P., et al. 2006, *A&A*, 451, 457
 Treister, E., & Urry, C. M. 2006, *ApJ*, 652, L79
 Treister, E., Urry, C. M., & Virani, S. 2009, *ApJ*, 696, 110
 Trichas, M., et al. 2012, *ApJS*, 200, 17
 Trouille, L., Barger, A. J., Cowie, L. L., Yang, Y., & Mushotzky, R. F. 2008, *ApJS*, 179, 1
 Tueller, J., Mushotzky, R. F., Barthelmy, S., Cannizzo, J. K., Gehrels, N., Markwardt, C. B., Skinner, G. K., & Winter, L. M. 2008, *ApJ*, 681, 113
 Tzanavaris, P., & Georgantopoulos, I. 2008, *A&A*, 480, 663
 Ueda, Y., Akiyama, M., Hasinger, G., Miyaji, T., & Watson, M. G. 2014, *ApJ*, 786, 104
 Ueda, Y., Akiyama, M., Ohta, K., & Miyaji, T. 2003, *ApJ*, 598, 886
 Ueda, Y., et al. 2011, *PASJ*, 63, 937
 Ueda, Y., Ishisaki, Y., Takahashi, T., Makishima, K., & Ohashi, T. 2001, *ApJS*, 133, 1
 Ueda, Y., et al. 1999, *ApJ*, 518, 656
 Vasudevan, R. V., Mushotzky, R. F., & Gandhi, P. 2013, *ApJ*, 770, L37
 Vito, F., Gilli, R., Vignali, C., Comastri, A., Brusa, M., Cappelluti, N., & Iwasawa, K. 2014, *MNRAS*, 445, 3557
 Vito, F., et al. 2013, *MNRAS*, 428, 354
 Winter, L. M., Mushotzky, R. F., Reynolds, C. S., & Tueller, J. 2009, *ApJ*, 690, 1322
 Wirth, G. D., et al. 2004, *AJ*, 127, 3121
 Worsley, M. A., Fabian, A. C., Bauer, F. E., Alexander, D. M., Brandt, W. N., & Lehmer, B. D. 2006, *MNRAS*, 368, 1735
 Xue, Y. Q., et al. 2010, *ApJ*, 720, 368
 —. 2011, *ApJS*, 195, 10
 —. 2012, *ApJ*, 758, 129
 Yencho, B., Barger, A. J., Trouille, L., & Winter, L. M. 2009, *ApJ*, 698, 380
 York, D. G., et al. 2000, *AJ*, 120, 1579

APPENDIX A: MULTIWAVELENGTH PHOTOMETRIC DATASETS

In this appendix we provide details of the photometric datasets that are used to identify counterparts and calculate photometric redshifts for X-ray sources in our *Chandra* fields. The table below gives the name of all bandpasses and their effective wavelengths. The prioritization in our likelihood ratio counterpart matching (where the band with priority 1 is used first and then the subsequent priority bands are used to identify additional, robust counterparts) is given in the third column. The final column gives the number of counterparts identified in each band and the corresponding fraction of the total X-ray sample.

Band	λ_{eff} (Å)	LR priority	N_{cntprpt}
CDFS			
IRAC [3.6μm]	35416.6	1	593 (95.6%)
IRAC [4.5μm]	44826.2
IRAC [5.8μm]	56457.2
IRAC [8.0μm]	78264.8
VIMOS <i>U</i>	3711.4
VIMOS <i>R</i>	6414.0	2	7 (1.1%)
WFC3 F125W	12425.8
WFC3 F160W	15324.7	3	0 (0.0%)
ACS <i>B</i>	4297.6
ACS <i>V</i>	5840.7

Band	λ_{eff} (Å)	LR priority	N_{entrprt}
ACS <i>I</i>	7668.3	4	0 (0.0%)
ACS <i>Z</i>	9021.7
ACS F814W	7993.3
ESO WFI <i>U</i>	3414.4
ESO WFI <i>B</i>	4558.5
ESO WFI <i>V</i>	5355.9
ESO WFI <i>R</i>	6451.0
ESO WFI <i>I</i>	8568.1	5	5 (0.8%)
CTIO <i>z</i>	8966.4
HAWK-I <i>Y</i>	10194.5
HAWK-I <i>J</i>	12556.2
HAWK-I <i>K</i>	21423.2
HAWK-I 1.061 μm	10619.2
SOFI <i>J</i>	12444.9
ISAAC <i>J</i>	12459.9
ISAAC <i>H</i>	16453.2
ISAAC <i>K</i>	21594.9	6	0 (0.0%)
SOFI K_S	21591.4
ISAAC K_S	21594.9
GALEX NUV	2271.1
GALEX FUV	1528.1
COMBO17 <i>B</i>	4554.4
COMBO17 <i>V</i>	5358.1
COMBO17 <i>R</i>	6432.4
COMBO17 <i>I</i>	8523.6
COMBO17 F420M	4177.7
COMBO17 F464M	4616.8
COMBO17 F485M	4858.5
COMBO17 F518M	5188.1
COMBO17 F571M	5716.1
COMBO17 F604M	6044.3
COMBO17 F646M	6450.9
COMBO17 F696M	6958.9
COMBO17 F753M	7530.5
COMBO17 F815M	8157.2
COMBO17 F855M	8556.4
COMBO17 F915M	9140.0

CDFN

IRAC [3.6 μm]	35416.6	1	446 (90.1%)
IRAC [4.5 μm]	44826.2
IRAC [5.8 μm]	56457.2
IRAC [8.0 μm]	78264.8
KPNO <i>U</i>	3566.7
SUBARU <i>B</i>	4366.4
SUBARU <i>V</i>	5429.6
SUBARU <i>R</i>	6486.4
SUBARU <i>I</i>	7952.0	2	13 (2.6%)
SUBARU <i>z'</i>	9146.7
ACS <i>B</i>	4297.6
ACS <i>V</i>	5840.7
ACS <i>I</i>	7668.3	3	0 (0.0%)
ACS <i>z</i>	9021.7
QUIRC <i>HK</i>	18264.0
CFHT <i>J</i>	12518.7
CAHA <i>J</i>	12029.7
CAHA <i>K</i>	22070.8
CFHT <i>K</i>	21530.8	4	5 (1.0%)
MOIRCS <i>K</i>	21354.4
GALEX FUV	1528.1
GALEX NUV	2271.1

Band	λ_{eff} (Å)	LR priority	N_{entrprt}
WFC3 F105W	10582.4
WFC3 F125W	12425.8
WFC3 F160W	15324.7

EGS

IRAC [3.6 μm]	35416.6	1	961 (92.5%)
IRAC [4.5 μm]	44826.2
IRAC [5.8 μm]	56457.2
IRAC [8 μm]	78264.8
Subaru R_c	6486.4	2	51 (4.9%)
ACS <i>V</i>	5796.7
ACS <i>I</i>	8234.0	3	0 (0.0%)
Subaru K_S	21354.4	4	1 (0.1%)
CFHTLS u^*	3805.6
CFHTLS g'	4833.7
CFHTLS r'	6234.1
CFHTLS i'	7659.1	5	1 (0.1%)
CFHTLS z'	8820.9
MMT u'	3604.1
MMT g'	4763.5
MMT i'	7770.5	6	1 (0.1%)
MMT z'	9030.9
DEEP <i>B</i>	4402.0
DEEP <i>R</i>	6595.1	7	10 (1.0%)
DEEP <i>I</i>	8118.7
Palomar <i>J</i>	12435.0
Palomar K_S	21353.0	8	4 (0.4%)
GALEX FUV	1528.1
GALEX NUV	2271.1
NICMOS F110W	10622.4
NICMOS F160W	15819.6
CAHA <i>J</i>	12029.7
WFC3 F125W	12425.8
WFC3 F160W	15324.7
NEWFIRM J_1	10441.4
NEWFIRM J_2	11930.0
NEWFIRM J_3	12764.0
NEWFIRM H_1	15585.4
NEWFIRM H_2	17048.8
NEWFIRM <i>K</i>	21643.9

COSMOS

CFHT u^*	3805.0
Subaru B_J	4427.2
Subaru V_J	5454.9
Subaru g^+	4728.6
Subaru r^+	6249.1
Subaru i^+	7646.0	1	1338 (85.2%)
Subaru z^+	9011.0
IRAC [3.6 μm]	35416.6	2	200 (12.7%)
IRAC [4.5 μm]	44826.2
IRAC [5.8 μm]	56457.2
IRAC [8 μm]	78264.8
CFHT i^*	7582.6
Subaru IB427	4262.0
Subaru IB464	4633.7
Subaru IB505	5060.9
Subaru IB574	5762.8
Subaru IB709	7071.5
Subaru IB827	8242.5
CFHT K_S	21530.8
UKIRT <i>J</i>	12464.8

Band	λ_{eff} (Å)	LR priority	N_{entrprt}
Subaru IB484	4847.6
Subaru IB527	5259.4
Subaru IB624	6230.9
Subaru IB679	6778.6
Subaru IB738	7359.5
Subaru IB767	7682.4
GALEX NUV	2271.1
GALEX FUV	1528.1
Bootes			
NDWFS B_W	4184.9
NDWFS R	6501.8
NDWFS I	8030.0	1	763 (91.7%)
SDSS u	3546.0
SDSS g	4669.6
SDSS r	6156.2
SDSS i	7471.6	2	6 (0.7%)
SDSS z	8917.4
NDWFS K	22078.4	3	9 (1.1%)
FLAMINGOS J	12442.5
FLAMINGOS K_S	21483.9
IRAC [3.6 μm]	35416.6	4	40 (4.8%)
IRAC [4.5 μm]	44826.2
IRAC [5.8 μm]	56457.2
IRAC [8 μm]	78264.8
GALEX FUV	1528.1
GALEX NUV	2271.1	5	3 (0.4%)