

OBSCURATION-DEPENDENT EVOLUTION OF ACTIVE GALACTIC NUCLEI

JOHANNES BUCHNER¹, ANTONIS GEORGAKAKIS^{1,2}, KIRPAL NANDRA¹, MURRAY BRIGHTMAN^{1,3}, MARIE-LUISE MENZEL¹, ZHU LIU⁴, LI-TING HSU¹, MARA SALVATO¹, CYPRIAN RANGEL⁵, JAMES AIRD⁶, ANDREA MERLONI¹, AND NICHOLAS ROSS⁷

¹ Max Planck Institut für Extraterrestrische Physik Giessenbachstrasse, D-85748 Garching, Germany; johannes.buchner.acad@gmx.com

² National Observatory of Athens, V. Paulou and I. Metaxa, 11532 Athens, Greece

³ Cahill Center for Astrophysics, California Institute of Technology, 1216 East California Boulevard, Pasadena, CA 91125, USA

⁴ National Astronomical Observatories, Chinese Academy of Sciences, 100012 Beijing, China

⁵ Astrophysics Group, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK

⁶ Department of Physics, Durham University, Durham DH1 3LE, UK

⁷ Department of Physics, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104, USA

Received 2014 October 7; accepted 2014 December 21; published 2015 March 26

ABSTRACT

We aim to constrain the evolution of active galactic nuclei (AGNs) as a function of obscuration using an X-ray-selected sample of ~ 2000 AGNs from a multi-tiered survey including the CDFS, AEGIS-XD, COSMOS, and XMM-XXL fields. The spectra of individual X-ray sources are analyzed using a Bayesian methodology with a physically realistic model to infer the posterior distribution of the hydrogen column density and intrinsic X-ray luminosity. We develop a novel non-parametric method that allows us to robustly infer the distribution of the AGN population in X-ray luminosity, redshift, and obscuring column density, relying only on minimal smoothness assumptions. Our analysis properly incorporates uncertainties from low count spectra, photometric redshift measurements, association incompleteness, and the limited sample size. We find that obscured AGNs with $N_{\text{H}} > 10^{22} \text{ cm}^{-2}$ account for $77_{-5}^{+4}\%$ of the number density and luminosity density of the accretion supermassive black hole population with $L_{\text{X}} > 10^{43} \text{ erg s}^{-1}$, averaged over cosmic time. Compton-thick AGNs account for approximately half the number and luminosity density of the obscured population, and $38_{-7}^{+8}\%$ of the total. We also find evidence that the evolution is obscuration dependent, with the strongest evolution around $N_{\text{H}} \approx 10^{23} \text{ cm}^{-2}$. We highlight this by measuring the obscured fraction in Compton-thin AGNs, which increases toward $z \sim 3$, where it is 25% higher than the local value. In contrast, the fraction of Compton-thick AGNs is consistent with being constant at $\approx 35\%$, independent of redshift and accretion luminosity. We discuss our findings in the context of existing models and conclude that the observed evolution is, to first order, a side effect of anti-hierarchical growth.

Key words: galaxies: active – quasars: supermassive black holes – surveys – X-rays: galaxies

Supporting material: machine-readable table

1. INTRODUCTION

Supermassive black holes (SMBH) are abundant in the local universe—every nearby massive galaxy harbors one (Richstone et al. 1998; Kormendy & Ho 2013). The majority of the growth of these SMBHs must have been through accretion processes (Soltan 1982; Merloni & Heinz 2008). The material accreted by the SMBHs likely originates from galactic scales and therefore one might expect some relation between the properties of galaxies and their central black holes. This idea was reinforced by the discovery of tight relationships between SMBHs and the properties of the stellar spheroid of galaxies (Magorrian et al. 1998; Ferrarese & Merritt 2000; Gebhardt et al. 2000; Tremaine et al. 2002; Ferrarese & Ford 2005; Kormendy & Ho 2013), which can be interpreted as evidence for co-evolution of active galactic nuclei (AGNs) and their host galaxies. For instance, the growth of AGNs through accretion may be linked to the host galaxies star formation, either weakly through a shared gas reservoir, or strongly via so-called “feedback” processes. In the latter, the energetic output from the AGNs into the galactic environment can strongly disturb the galaxy (Silk & Rees 1998; Fabian & Iwasawa 1999).

To understand the growth of SMBHs over cosmic time and the relationship of that growth to galaxy evolution, it is first necessary to estimate the space density of accreting SMBHs, and its evolution over cosmic time in an unbiased way. This is challenging because it is known that many AGNs are

enshrouded in gas and dust, making them difficult to detect directly. X-ray emission is an efficient way to reveal these obscured SMBHs, at least for moderate column densities. On the other hand when the column of equivalent neutral hydrogen exceeds unit optical depth corresponding to the Thomson cross section ($N_{\text{H}} \approx 1.5 \times 10^{24} \text{ cm}^{-2}$), AGNs become difficult to find even at X-ray wavelengths. In this so-called “Compton-thick” regime the AGNs can nonetheless still be identified by hard X-ray emission, which can emerge from the thick covering and/or via radiation that is reflected and scattered into the line of sight. These processes result in a characteristic shape of the X-ray spectrum, including a flat continuum and intense iron $K\alpha$ emission lines that can be identified using spectral analysis. The contribution of these most heavily obscured AGNs represents the major remaining uncertainty in our knowledge of the accretion history, and placing limits on their contribution is therefore vital.

Previous work has attempted to constrain the number of Compton-thick AGNs using X-ray background (XRB) synthesis (e.g., Gilli et al. 2007), but while this population is needed to reproduce the shape of the XRB spectrum, the method is relatively insensitive to the precise Compton-thick AGN fraction (Akylas et al. 2012). Multi-wavelength data have also been exploited to identify Compton-thick AGNs and constrain their number. These include optical (Risaliti et al. 1999; Cappi et al. 2006; Panessa et al. 2006; Akylas & Georgantopoulos 2009; Gilli et al. 2010; Vignali et al. 2010, 2014; Jia et al. 2013;

Mignoli et al. 2013), infrared (e.g., Fiore et al. 2008, 2009; Alexander et al. 2011; Brightman & Nandra 2011a, 2011b), and hard X-ray (e.g., Sazonov et al. 2008; Burlon et al. 2011; Alexander et al. 2013; Lanzuisi et al. 2014) diagnostics applied to local and non-local AGN samples. These studies estimate Compton-thick fractions relative to the overall obscured AGN population in the range of 30%–50%, thereby demonstrating that such heavily obscured sources represent a sizable fraction of the AGN population in the nearby universe. Following a handful of early discoveries of Compton-thick AGNs in deep X-ray surveys (Norman et al. 2002; Tozzi et al. 2006; Comastri et al. 2011), an important recent development has been the identification of significant samples of Compton-thick AGNs at moderate to high redshifts (Brightman & Ueda 2012; Georgantopoulos et al. 2013; Buchner et al. 2014; Brightman et al. 2014). This has been enabled by the combination of extremely deep X-ray data, sufficient to constrain the X-ray spectra, along with extensive multi-wavelength coverage of X-ray survey regions, and new techniques able to determine accurate photometric redshifts for X-ray emitting AGNs (Salvato et al. 2009, 2011). This then offers the exciting possibility of starting to constrain the evolution of the obscured AGN populations, including Compton-thick AGNs, provided that the selection function can be sufficiently well understood.

Further interest in the obscured AGN population lies in the nature of the obscuration itself, and its possible relationship with other source properties. In the standard unification picture (Antonucci & Miller 1985; Antonucci 1993), all AGNs are surrounded by an optically thick toroidal structure relatively close (parsec scale) to the central engine, which can obscure the line of sight depending on the viewing angle. Alternatively, or additionally, obscuration can occur at galactic scales (Maiolino & Rieke 1995). Observations show that AGN host galaxies are massive and lie, at least in an average sense, on the main sequence of star formation (Santini et al. 2012). At moderate redshifts, $z \approx 1$ –2, such galaxies are known to be gas-rich, with gas contents 3–10 times larger than local samples (Tacconi et al. 2013). It is therefore possible that obscuration in moderate redshift AGNs is associated with the same gas fueling both star formation and the accretion process itself. In some scenarios (Hopkins et al. 2006a, 2012), obscured AGNs represent a distinct phase in the co-evolution of the galaxy and its central black hole, with energy output from accretion sweeping up gas from the surroundings with potentially profound effects on star formation (Silk & Rees 1998; Fabian 1999; King 2003). Determining accurately the fraction of obscured AGNs, including Compton-thick AGNs, as a function of other parameters such as luminosity and redshift is critical for both unification and co-evolution models. Previous work has provided evidence that the obscured fraction depends on luminosity, with obscuration being less common in more powerful sources (Ueda et al. 2003, 2014; Treister et al. 2004; Akylas et al. 2006; Hasinger 2008; Ebrero et al. 2009), with local samples suggesting that the obscured fraction may peak at around $L_X \sim 10^{42} \text{ erg s}^{-1}$ (Brightman & Nandra 2011b; Burlon et al. 2011). More controversially, it has been suggested that the obscured fraction increases out to high redshift (La Franca et al. 2005; Hasinger 2008; Brightman & Ueda 2012; Iwasawa et al. 2012; Vito et al. 2014; but see Ueda et al. 2003; Akylas et al. 2006; Ebrero et al. 2009).

The same deep X-ray data and multi-wavelength supporting data needed to identify Compton-thick AGNs also enable an accurate determination of the obscured AGN fraction and its evolution, again with the proviso that the selec-

tion function must be well understood. This is, however, very challenging, given the limited photon statistics in deep X-ray data, and associated uncertainties in X-ray spectral analysis, combined with additional uncertainties, e.g., in photometric redshifts. Moreover, studies of nearby objects show that the X-ray spectra of AGNs are complex, including multiple emission components, e.g., absorption, reflection, or scattering of direct AGN emission.

In this paper, we develop a novel non-parametric method for determining the space density of AGNs as a function of accretion luminosity, redshift, and hydrogen column density. We build on the X-ray spectral analysis of Buchner et al. (2014), applying their Bayesian spectral analysis technique of a realistic, physically motivated model to a multi-layered survey, to determine the luminosity and level of obscuration in a large sample of X-ray-selected AGNs across a wide range of redshifts. Important features of our approach are that, first, all sources of uncertainty are consistently factored into the analysis and second, unlike previous studies, no functional form is imposed. Instead, we use a smoothness assumption to allow the data to determine the space density of AGNs and its dependence on luminosity, redshift, and the level of obscuration along the line of sight. The methodology is designed to be informative at regions of the parameter space where data are sparse and therefore constraints are expected to be loose.

We adopt a cosmology of $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_M = 0.3$, and $\Omega_\Lambda = 0.7$. Solar abundances are assumed. The galactic photo-electric absorption along the line-of-sight direction is modeled with $N_H \approx 8.8 \times 10^{19}$, 1.3×10^{20} , 2.7×10^{20} , and $2.2 \times 10^{21} \text{ cm}^{-2}$ for the CDFS, AEGIS-XD, COSMOS, and XMM-XXL fields, respectively (Dickey & Lockman 1990; Stark et al. 1992; Kalberla et al. 2005). In this work, luminosity (L) always refers to the intrinsic (absorption-corrected) luminosity in the 2–10 keV rest-frame energy range.

2. DATA

The determination of the obscured fraction of AGNs as a function of redshift and accretion luminosity requires good coverage of the L_X – z plane. We therefore combine X-ray survey fields with different characteristics in terms of depth and areal coverage. These include the *Chandra* Deep Field South (CDFS; Xue et al. 2011), the All Wavelength Extended Groth strip International Survey (AEGIS; Davis et al. 2007), the Cosmological evolution Survey (COSMOS; Scoville et al. 2007), and the equatorial region of the XMM-XXL survey (PI: Pierre).

In our analysis we focus on the region of the CDFS that is covered by the 4 Ms *Chandra* observations and the part of the AEGIS field that has been surveyed by *Chandra* to an exposure of 800 ks (AEGIS-XD; Nandra et al. 2015). In the COSMOS field we use the region covered by the *Chandra* observations performed between 2006 November and 2007 June (C-COSMOS; Elvis et al. 2009). Table 1 presents information on the individual X-ray fields used in this paper.

2.1. *Chandra* Observations

The data reduction, source detection, and source catalogs for two of the *Chandra* fields, the CDFS (Rangel et al. 2013) and the AEGIS-XD (Nandra et al. 2015), follow the methodology described by Laird et al. (2009). Briefly, hot pixels, cosmic ray afterglows, and times of anomalously high backgrounds are removed to produce clean level 2 event files. These are then

Table 1
Survey Fields

Survey	CDFS	AEGIS-XD	C-COSMOS	<i>XMM</i> -XXL
Survey area	464 arcmin ²	1010 arcmin ²	0.9 deg ²	20 deg ²
Total/central exposure time	4 Ms	2.4 Ms/800 ks	1.8 Ms/160 ks	10 ks

Table 2
Sample Statistics for the Individual Data Extraction Steps

Section	Extraction Step	CDFS	AEGIS-XD	C-COSMOS	<i>XMM</i> -XXL	Total ^a
2.1, 2.2	X-ray hard-band detected	326	574	1016	206	2122
2.3	Association with optical/IR	315	559	1016	206	2096
2.4	Redshift information: spec-z/photo-z/no-z/removed stars	180/131/11/4	227/322/15/10	491/519/0/6	174/0/32/0	1072/986/58/20
2.5	X-ray spectral extraction and data analysis: successful/failed extraction	321/1	564/0	1010/0	206/0	2101/1
2.6	Galaxies removed (based on $L_{2-10\text{keV}} < 10^{42} \text{ erg s}^{-1}$)	20	11	14	1	46
2.7	Objects used for LF analysis	302	553	996	205	2056

Note. ^a Total number of sources selected.

aligned using bright sources and subsequently merged. Images and exposure maps are constructed in four energy bands, 0.5–2, 2–7, 5–7, and 0.5–7 keV. A candidate source list is created using `wavdetect` at a low significance threshold (10^{-4}). Source and background counts are then extracted at each candidate source position. The source region corresponds to the 70% encircled energy fraction (EEF) of the point-spread function (PSF). The background count region is an annulus with inner radius of 1.5 times the 90% EEF of the PSF and width of 100 pixels 0'.5 in size. For the determination of the background, candidate sources are masked out. For each candidate source position, the Poisson probability that the observed counts are a background fluctuation is computed. Sources are accepted if that probability is $< 4 \times 10^{-6}$ (Nandra et al. 2005). In this paper, we use the hard band (2–7 keV) selected sample. The X-ray sensitivity curves are estimated by extrapolating the background counts and exposure maps in the 2–7 keV band to the limiting flux of a source in the 0.5–10 keV energy range by adopting the methods described in Georgakakis et al. (2008). For the C-COSMOS survey we use the 2–7 keV selected X-ray source catalog presented by Elvis et al. (2009). The number of X-ray sources detected are broken down by field in Table 2. The corresponding sensitivity curves are shown in Figure 2.

2.2. *XMM* Observations

The *Chandra* surveys above need to be complemented by a shallower and wider X-ray field to place constraints at the bright end of the luminosity function (LF). For this we use the *XMM*-XXL survey, which consists of 10 ks *XMM* pointings that cover a total area of 50 deg² split into two equal size fields. In this paper we focus on the equatorial sub-region of the *XMM*-XXL. The data reduction, source detection, and sensitivity map construction follow the methods described in Georgakakis & Nandra (2011). A full description of the *XMM*-XXL X-ray source catalog generation are presented by Z. Liu et al. (in preparation). The most salient details of those steps are outlined here.

The *XMM* observations were reduced using the Science Analysis System (SAS) version 12. The first step is to produce event files from the observation data files using the `EPCHAIN` and `EMCHAIN` tasks of SAS for the European Photon Imaging Camera (EPIC; Strüder et al. 2001; Turner et al. 2001) PN and

MOS detectors, respectively. Pixels along the edges of the CCDs of the PN and MOS detectors are removed because their inclusion often results in spurious detections. Flaring background periods are identified and excluded using a methodology similar to that described in Nandra et al. (2007). Images and exposure maps in celestial coordinates with pixel size of 4'.35 are constructed in five energy bands: 0.5–8, 0.5–2, 2–8, 5–8, and 7.5–12 keV. All overlapping EPIC images are merged prior to source detection to increase the sensitivity to point sources. The detection algorithm is applied independently to each of the five spectral bands defined above.

The source detection methodology is similar to that described in Laird et al. (2009) in the case of *Chandra* data. Source candidates are identified using the wavelet-based `EWAVELET` source detection task of SAS at a low threshold of 4σ above the background, where σ is the rms of the background counts. For each candidate source, the Poisson probability of a random background fluctuation is estimated. This step involves the extraction of the total counts at the position of the source and the determination of the local background value. To match the asymmetric PSF of *XMM*, especially off-axis, elliptical apertures were used from the *XMM*/EPIC PSF parameterization of Georgakakis & Nandra (2011). The count extraction region is obtained by scaling the elliptical apertures to contain 70% of the PSF EEF. The total of the counts at a candidate source position, T , is the sum of the extracted counts from individual EPIC cameras. For each source, the local background is estimated by first masking out all detections within 4' of the source position using an elliptical aperture that corresponds to the 80% EEF ellipse. The counts from individual EPIC cameras are then extracted using elliptical annuli centered on the source with inner and outer semi-major axes of 5 and 15 pixels (0'.36 and 1'.09) respectively, while keeping the same shape as the elliptical aperture in terms of rotation and ellipticity. The mean local background, B , is then estimated by summing up the background counts from individual EPIC cameras after scaling them down to the area of the source count extraction region. The Poisson probability $P(T, B)$ that the extracted counts at the source position, T , are a random fluctuation of the background is calculated. We consider as sources the detections with $P(T, B) < 4 \times 10^{-6}$. The above methodology is optimized for the detection of point sources. The final

catalog, however, includes extended X-ray sources associated with hot gas from galaxy clusters or groups. Also, the extended X-ray emission regions of bright clusters are often split into multiple spurious detections by our source detection pipeline. The EMLDETECT task of SAS is used to identify extended sources (i.e., groups or clusters) and spurious detections. Point sources for which EMLDETECT failed to determine a reliable fit and hence are considered spurious and excluded from the analysis. The EPOSCORR task of SAS is used to correct for systematic errors in the astrometric positions of X-ray sources by cross-correlating with optically sources in the Sloan Digital Sky Survey (SDSS)-DR8 catalog (Aihara et al. 2011) with magnitudes of $r < 22$ mag.

The flux of each source in different spectral bands is estimated by assuming a power-law X-ray spectrum with $\Gamma = 1.4$, i.e., similar to the XRB, absorbed by the appropriate Galactic hydrogen column density. The latter is derived from the H I map of Kalberla et al. (2005) using the right ascension and declination of the aim point of each XMM observation and the NH task of FTOOLS. The energy to flux conversion factors are such that the counts from the 0.5–2, 0.5–8, 2–8, 5–8, and 7.5–12 keV bands are transformed to fluxes in the 0.5–2, 0.5–10, 2–10, 5–10, and 7.5–12 keV bands, respectively.

In this work, we use the XMM-XXL point source subsample that is selected in the 2–8 keV band. We minimize optical identification and spectroscopic redshift determination incompleteness (see next sections) by applying a bright flux cut, $f_X(2\text{--}10\text{ keV}) > 7 \times 10^{-14} \text{ erg s}^{-1} \text{ cm}^{-2}$. The total number of XMM-XXL X-ray sources used in the analysis is 206 (see Table 2).

The construction of the sensitivity maps follows the methodology of Georgakakis et al. (2008). The 2–10 keV sensitivity curve of the sample is plotted in Figure 2. The window function of the spectroscopic follow-up observations in the XMM-XXL field (see Section 2.4) is taken into account in this calculation. Also, in addition to the Poisson false detection probability threshold of $< 4 \times 10^{-6}$ in the sensitivity map calculation, we take into account the flux cut $f_X(2\text{--}10\text{ keV}) > 7 \times 10^{-14} \text{ erg s}^{-1} \text{ cm}^{-2}$, by estimating at each survey position the probability of measuring a flux above this limit.

The number of X-ray sources detected—about 2000 in total—are broken down by field in Table 2. The corresponding sensitivity curves are shown in Figure 2. We associate these detected X-ray source positions to optical counterparts in order to obtain redshifts.

2.3. Association with Optical/IR Counterparts

The identification of the X-ray sources with optical or infrared counterparts in the AEGIS-XD, COSMOS, and XMM-XXL used the likelihood ratio method of Sutherland & Saunders (1992). Specific details on the association of X-ray sources with optical/infrared counterparts are presented by Nandra et al. (2015). They used the multi-waveband photometric catalog provided by the Rainbow Cosmological Surveys Database (Pérez-González et al. 2008; Barro et al. 2011a, 2011b).

The counterparts of C-COSMOS X-ray sources are taken from Civano et al. (2012). For the identification they used the I -band-selected optical sample of Capak et al. (2007), the K -band photometry of McCracken et al. (2010), and the IRAC 3.6 μm catalog of Sanders et al. (2007).

X-ray sources in the XMM-XXL survey were matched to the SDSS-DR8 photometric catalog Aihara et al. (2011) following the methods described in Georgakakis & Nandra (2011).

Hsu et al. (2014) presents the counterparts of the CDFS X-ray sources. They used a sophisticated Bayesian version of the likelihood ratio method, which is based on the probabilistic formalism of Budavári & Szalay (2008). The photometric catalogs used to identify X-ray sources include the CANDELS H -band-selected multi-wavelength catalog in Guo et al. (2013), the MUSYC catalog presented by Cardamone et al. (2010) and the TENIS near-infrared-selected source catalog described by Hsieh et al. (2012).

2.4. Redshift Estimation

The X-ray survey fields used in this paper benefit from extensive spectroscopic campaigns that also specifically target X-ray sources. In the CDFS, we used the spectroscopic redshifts compiled by N. Hathi (private communication; see Hsu et al. 2014). Spectroscopic redshift measurements of X-ray sources in the AEGIS field are extracted from the compilation presented in Nandra et al. (2015), which also included the DEEP2 (Newman et al. 2012) and DEEP3 galaxy redshift surveys (Cooper et al. 2011, 2012) as well as observations carried out at MMT using the Hectospec fiber spectrograph (Coil et al. 2009). Redshifts in the C-COSMOS are used from the compilation of Civano et al. (2012), which includes the public releases of the VIMOS/zCOSMOS bright project (Lilly et al. 2009) and the Magellan/IMACS observation campaigns (Trump et al. 2009).

In the case of XMM-XXL, optical spectroscopy is from Stalin et al. (2010) and the Baryon Oscillation Spectroscopic Survey (Dawson et al. 2013; Bolton et al. 2012; Smee et al. 2013), as well as dedicated SDSS III (York et al. 2000; Gunn et al. 2006; SDSS-III: Eisenstein et al. 2011) ancillary science observations, which specifically targeted X-ray sources in the equatorial XMM-XXL field (PI: A. Merloni, A. Georgakakis). Targets were selected to have $f_X(0.5\text{--}10\text{ keV}) > 10^{-14} \text{ erg s}^{-1} \text{ cm}^{-2}$, and $17 < r < 22.5$, where r corresponds to either the PSF magnitude in the case of optical unresolved sources (SDSS type = 6) or the model magnitude for resolved sources (see M.-L. Menzel et al., in preparation). At the flux limit $f_X(2\text{--}10\text{ keV}) > 7 \times 10^{-14} \text{ erg s}^{-1} \text{ cm}^{-2}$, 84% (174/207) of the XMM-XXL sources have secure redshift measurements.

The CDFS, AEGIS-XD, and COSMOS fields have multi-wavelength photometric observations that allow photometric redshift estimates for sources that lack spectroscopy. The methods developed by Salvato et al. (2009, 2011) are adopted to achieve photometric redshift accuracies for X-ray AGNs comparable to galaxy samples. The photometric redshifts catalogs are presented in Hsu et al. (2014) for the CDFS, Nandra et al. (2015) for the AEGIS-XD field, and (Salvato et al. 2011) for the COSMOS field. The photometric redshifts are included in the analysis in the form of probability distribution functions, incorporating systematic uncertainties (see below). A by-product of the photometric redshift determination is the characterization of the spectral energy distribution (SED) of X-ray sources. This information is used to identify candidate Galactic stars in the X-ray sample and exclude them from the analysis. Following Salvato et al. (2011), if stellar templates provide an improved fit to the SED of a source, as measured by the reduced χ^2 , i.e., $\chi_{\text{star}}^2 < \chi_{\text{gal}}^2/1.5$, and the source is point like in the optical images, then it is considered to be a Galactic star candidate. The number of removed stars is indicated in Table 2.




Figure 1. Luminosity–redshift plot of the full sample. The intrinsic 2–10 keV luminosity in erg s^{-1} and redshift is plotted. For this visualization, we color sources based on $\log N_{\text{H}}$ (in units of cm^{-2} as follows: 20–22, “unobscured,” red; 22–24, “Compton-thin obscured,” blue; 24–26, “Compton-thick,” green). If more than 50% of the N_{H} posterior probability lies within one of the intervals above, the source is color-coded accordingly. Due to the heavy suppression of the flux of Compton-thick objects by absorption, they are typically only detectable at higher intrinsic luminosity compared to unobscured or Compton-thin sources. Some objects lie below the $L = 10^{42} \text{ erg s}^{-1}$ limit and are not used in this work. The inset shows the N_{H} histogram of our sample, while the top and left axes show histograms of the redshift and luminosity distributions in gray. These plots are constructed by drawing a random posterior sample for each object, and thus includes the uncertainty in the parameter estimation as well as the intrinsic distribution. Sources above redshift $z = 3$ that have spectroscopic redshift estimates are indicated with large black circles.

For the *XMM-XXL* field, multi-wavelength photometry from the UV to the infrared which homogeneously covers the surveyed area is not available. This is essential for reliable photometric redshifts, especially in the case of bright X-ray samples like the *XMM-XXL* (e.g., Salvato et al. 2011). Therefore, for the number of X-ray sources without spectroscopic redshift measurement in that field (32/206), we choose not to determine photometric redshifts. These sources are still included in the analysis by assigning them a flat redshift prior (see below for details). In the *XMM-XXL* bright subsample used here, there are no stars based on inspection of the images.

To summarize, the redshift determination falls in one of three cases for each X-ray source.

1. *No redshift information available (58 sources)*. This is the case if no spectroscopic redshifts are available, and photometric redshift estimation is not possible due to limited photometry. This case also applies when no secure association was found. Such a source is associated with a flat redshift distribution in the interval $z = 0.001\text{--}7$.
2. *Spectroscopic redshifts (1072 sources)*. Wherever secure spectroscopic redshifts are available, they are used directly, i.e., do not associate uncertainty with them. Different surveys use different conventions to define the reliability of the redshift measurement. In this paper, we only consider spectroscopic redshifts from the top two quality classes of any study, which typically corresponds to a probability better than 95% of being correct.
3. *Photometric redshifts (986 sources)*. These are included in the analysis in the form of probability distribution functions. We also incorporate systematic uncertainties, e.g., due to incorrect association with an optical counterpart (see Appendix B).

Table 2 shows a breakdown of how many sources fall into each category. The distribution of redshifts is plotted in Figure 1.

2.5. Extraction and Analysis of X-Ray Spectra

For the *Chandra* surveys, the ACIS EXTRACT (AE) software package (Broos et al. 2010) was used to extract spectra for each source. The extraction followed the same methodology described in Brightman et al. (2014). Initially, each source and each pointing is dealt with separately. AE simulates the PSFs at each source position. Regions enclosing 90% PSF at 1.5 keV were used to extract source spectra. The background regions are constructed around the sources such that they contain at least 100 counts, with other sources masked out. AE also constructs local response matrix files (RMFs) and auxiliary matrix files (ARFs) using ray tracing. As a final step, AE merges the extracted spectra so that each source has a single source spectrum, a single local background spectrum, ARF and RMF.

In the *XMM-XXL* field, the X-ray source and background spectra were extracted separately for each EPIC camera using *SAS* version 12.0.1. The source extraction regions were chosen to maximize the signal-to-noise ratio using the *EREGIONANALYSE* task of *SAS*. In the case of EPIC–PN camera, the background region has a radius of $90''$ and is placed, after visual inspection, on a source-free region on the same CCD and raw *Y*-axis as the corresponding source. In the case of EPIC–MOS, because of the larger size of individual CCDs it is possible to define the background region as an annulus with inner and outer radii of $40''$ and $110''$, respectively. If that was not possible, because, e.g., a source might lie close to CCD gaps then on the background region a circle of $90''$ radius was placed on a source-free region on the same CCD as the corresponding source. ARFs and RMFs are generated using the *ARFGEN* and *RMFGEN* tasks of *SAS*. Finally, the ARFs, RMFs, source, and background spectra of the same object observed by different EPIC cameras are merged by weighting with the exposure time to produce unique spectral products for each source.

Each X-ray spectrum was analyzed with spectral models following the Bayesian methodology of Buchner et al. (2014). We only consider their best model, `torus+pexmon+scattering`, which consists of

1. an intrinsic power-law spectrum modified by photo-electric absorption and Compton scattering from an obscurer with toroidal geometry simulated by the `torus` model of Brightman & Nandra (2011a);
2. a Compton-reflection component approximated by the `pexmon` model (Nandra et al. 2007), and
3. a soft scattering component, which is parameterized by an unabsorbed power law with the same spectral slope as the intrinsic one and normalization that cannot exceed 10% of the intrinsic power-law spectrum.

This model has five free parameters, the slope of the intrinsic power-law spectrum, Γ , the line-of-sight column density of the obscurer, N_{H} , the normalizations of the intrinsic power-law, Compton-reflection, and soft scattering components. For the `torus` model of Brightman & Nandra (2011a), we fix the opening angle parameter to 45° and the viewing angle parameter to the maximum allowed, 87° , i.e., edge on. For `pexmon` component we use the same incident power-law photon index as the intrinsic one spectrum and fix the inclination to 60° . The normalization of the `pexmon` (R parameter) is modeled relative to that of the `torus` component and is allowed to vary between 0.1 and 100. In the analysis, we use flat priors on the luminosity/normalization parameters and the column density, and an informed Gaussian prior of mean 1.95 and standard deviation 0.15 for the photon index Γ (Nandra & Pounds 1994).

In the spectral analysis, redshifts are either fixed, when spectroscopy is available, or are included in the form of priors that follow a probability distribution function (photometric redshifts and missing counterparts; see Section 2.4 above).

Following the methodology of Buchner et al. (2014), we use the 0.5–8 keV spectrum and the redshift information to infer posterior distributions for the X-ray luminosity L_X , redshift z , and column density N_{H} .

2.6. Sample Selection

Based on the initial hard-band detection, normal galaxies (i.e., non-AGNs) should already be largely absent from our sample, except for strongly starbursting galaxies. To exclude these in a conservative manner, we remove intrinsically faint sources that have $L_X < 10^{42}$ erg s $^{-1}$ with 90% probability. The number of objects used in the LF analysis is shown in Table 2. Figure 1 shows how the sample is spread in the luminosity–redshift plane.

3. METHODOLOGY

In this work, we want to study the evolution of AGN obscuration. More specifically, we are interested in the distribution of the population in X-ray luminosity L and column density N_{H} , i.e., the LF. Our sample is a specific draw from this distribution, with a sampling bias in detection and incomplete information on each object. Inference methods for such scenarios have been known for a long time (Marshall et al. 1983).

3.1. Luminosity Function Analysis

Our general approach to establish the evolution of the obscured and unobscured populations of AGNs is to determine the X-ray luminosity function (XLF), accounting for all measurement uncertainties, which can then be analyzed as a function

of obscuration to determine the absolute and relative number densities of obscured and unobscured sources.

Direct visualization of the data is difficult if we want to stay true to the uncertainties. Faint objects, which dominate the sample, are highly uncertain in their properties, namely the intrinsic luminosity L , obscuration N_{H} , but also their redshift z in the case of photometric estimates. For instance, Compton-thick AGNs can have large “probability clouds” for their parameters. This prohibits us from assigning objects to bins for visualization. For direct visualization, three approaches can be considered: (1) assigning each object to a random luminosity bin based on its probability distribution, and then estimating the density in the bin, (2) assigning each object to each luminosity bin with a probability weight, and then estimating the density in the bin, and (3) computing for each luminosity bin the number of objects that have a higher luminosity with, e.g., 90% probability. Method (2) has the difficulty that the “number” in each bin is no longer an integer—requiring interpolation of the Poisson distribution formula. Methods (1) and (2) assume a frequency interpretation of the uncertainty probability distributions—which is not reasonable as every object is different and the sample size is small. Method (3) may be useful for checking whether data and model agree, but does not yield an intuitive visualization. We thus find none of these methods satisfying, and develop a new approach (next section), which visualizes the data and estimates the LF at the same time.

In Appendix A, we review the statistical footing of analyzing population demographics by reviewing and combining the works of Loredó (2004) and Kelly et al. (2008). In this work, we use the usual Poisson likelihood:

$$\mathcal{L} = \frac{\prod_k \int \frac{\phi(\mathcal{C})}{d\mathcal{C}} \cdot p(d_k, D|\mathcal{C}) \cdot \frac{dV}{dz} d\mathcal{C}}{\exp \left\{ \int \frac{\phi(\mathcal{C})}{d\mathcal{C}} \cdot A(\mathcal{C}) \cdot \frac{dV}{dz} d\mathcal{C} \right\}}, \quad (1)$$

where $\mathcal{C} = \{\log L, z, \log N_{\text{H}}\}$, $p(d_k, D|\mathcal{C})$ represents the results of the spectral analysis of data d_k from the detected object k , which is weighted by the LF ϕ . The integral in the denominator computes the expected number of sources by convolving the LF with the sensitivity curve A (shown in Figure 2 as a dotted magenta line). The conversion from L, z, N_{H} to a sensitive area requires a spectral model, for which we use the `torus` model, and we average over Γ in exact correspondence to the Gamma prior used in our analysis. Here, the additional scattering components (`+pexmon+scattering`, see Section 2.5) are not used. Both contribute only minimally to the 2–10 keV flux: Buchner et al. (2014) presents an extreme example in detail, where both components show the strongest contribution in their sample, but the 2–10 keV flux is only affected by 10% (0.04 dex). Future studies may take these components and their normalizations into account as well. Appendix A gives an extensive derivation and caveats for Equation (1).

3.2. Non-parametric Approach

We would like to use the power and safety of a likelihood-based analysis but without the rigidity of a functional form, allowing discovery of the shape of the LF. Our method for analyzing the LF is thus, in a simplistic description, to fit a three-dimensional (L, z, N_{H}) histogram as the LF model.

Using, e.g., $10 \times 10 \times 10 = 1000$ bins already means that the problem is largely underdetermined. We thus need to input additional knowledge. Here we make the reasonable assumption that the function does not vary rapidly between neighboring




Figure 2. Area curve. The bottom plot shows all X-ray sensitivity curves on a logarithmic scale for the sum of all fields (dotted magenta line) and the individual fields. The top plot shows the sensitivity curves for the individual fields on a linear scale and normalized to their respective maximum area (see Table 1), to indicate the flux limit for detection. The *XMM-XXL* curve is limited in flux by 7×10^{-14} erg s $^{-1}$, but we incorporate the uncertainty of measuring a higher flux due to Poisson variance. This introduces a smooth transition.

bins (in particular for the redshift). We encode this smoothness prior in two approaches, both using the normal distribution to penalize large deviations. These two priors, “*constant-value*” and “*constant-slope*” are explained below and illustrated in Figure 3.

1. *Constant-value.* This prior retains the current value unless constraints are imposed by the data. The value of a bin should scatter around its neighbor’s density value $\log \phi_{i+1} = \text{Normal}(\log \phi_i, \sigma)$ with an allowed correlation width σ for each axis (σ_L, σ_z).
2. *Constant-slope.* This prior keeps power-law slopes intact unless constraints are imposed by the data. The log-density slope between bins should scatter around its neighbor’s slope $\log \phi_{i+2} = \log \phi_{i+1} + \text{Normal}(\log \phi_{i+1} - \log \phi_i, \sigma)$, with the deviation from the slope for each axis (σ_L, σ_z).

For the N_H dimension, we always use the constant-value prior, as in our view a power-law dependence is not appropriate here. With this simple prescription, we can recover a smooth field by fitting a model, whose shape is driven by the data.

Whenever the data constrain the result well, the results from either prior prescription (constant-value and constant-slope) will be the same. Where the constraints are poor, the results will differ depending on the adopted prior. We thus take the difference in the reconstructions as an indication of whether the data or the priors dominate, and in the latter case the difference between the two is an indication of the uncertainty in the determination.

It should be stressed that the choice of binning and in particular the correlation strengths σ_L, σ_z , and σ_{N_H} can influence the result. Motivated by the number of data points in each bin, we choose our pixelation as 11 bins of logarithmically spaced luminosity 42–46 (units erg s $^{-1}$), redshift bin edges 0.001, 0.1, 0.3, 0.5, 0.75, 1, 1.25, 1.5, 1.8, 2.1, 2.7, 3.2, 4, and 7, and log N_H

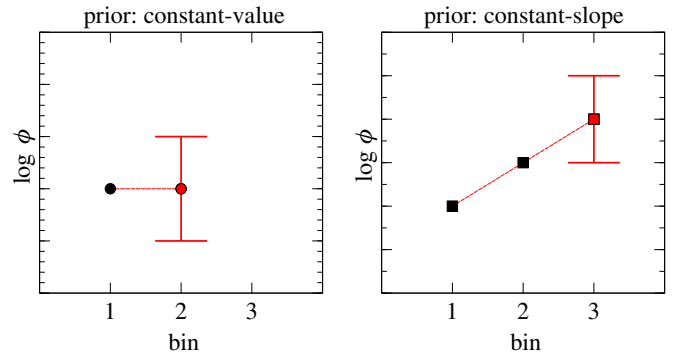


Figure 3. Illustration of the two smoothness priors used. The left panel shows the “constant-value” prior: the extrapolation from well-constrained step-function bins (black) to a neighboring bin (red) is done by assuming the same value with a fixed uncertainty σ , whose value encodes the assumed correlation strength or smoothness. The right panel illustrates the “constant-slope” prior. The value for the neighboring bin (red) is predicted by continuing the same slope from the black points. As the prior is defined in logarithmic units of the density, this behavior corresponds to preferring a power law.

bin edges 20,21,22,23,24,25, and 26. The correlation strength σ is defined for neighboring bins, and this choice is important. Notably, if σ is too small, the model will be flattened out, as the prior dominates, while if σ is too large, no smoothness assumption is used, and uncertainties will be large. Ideally, we would like to recover σ from the data. Unfortunately, our tests show that this computation is not numerically stable. We do find, however, that above some value of σ , the results are stable regardless of the choice of σ . We thus just choose reasonable values for σ , namely $\sigma_L = 0.5$, $\sigma_z = 0.5$, and $\sigma_{N_H} = 0.75$. This encodes, roughly speaking, that neighboring bins have the same order of magnitude in space density. These values have been chosen after a few initial tests, but were not tuned to give optimal results. Rather, we believe, they are one possible, and reasonable, a priori choice.

We are also highly interested in the uncertainties of our smooth field reconstruction method. Bins with data will be tightly constrained, while bins without information will have increasing uncertainty with distance from the data. Obtaining uncertainty estimates with so many parameters is not trivial, especially as the parameters are correlated by definition. We use a Hamiltonian Markov chain Monte Carlo code named “Stan” (Stan Development Team 2014). Stan uses the sophisticated No-U-Turn Sampler (Hoffman & Gelman 2011) technique to ensure rapid mixing of the Markov chain by avoiding cyclic explorations. Our Stan model is shown in Algorithm 1 in Appendix C.

4. RESULTS

A few words on the visualization are warranted. In all relevant figures, we plot the posterior distributions of our three-dimensional step function model from various axis views (L, z , and N_H). We always show the median result for the constant-value prior as a *dashed line*, and the median result for the constant-slope prior as a *solid line*. A feature of our methodology is that uncertainties are realistic and reflect regions of parameter space where data are sparse. Whenever the data constrain the result well, the results from either prior prescription (constant-value or constant-slope) will be the same. Only when constraints from the data are poor do the results differ. We thus take the difference in the reconstructions as an indication of whether the data or the priors dominate the result. We plot the 10%–90%

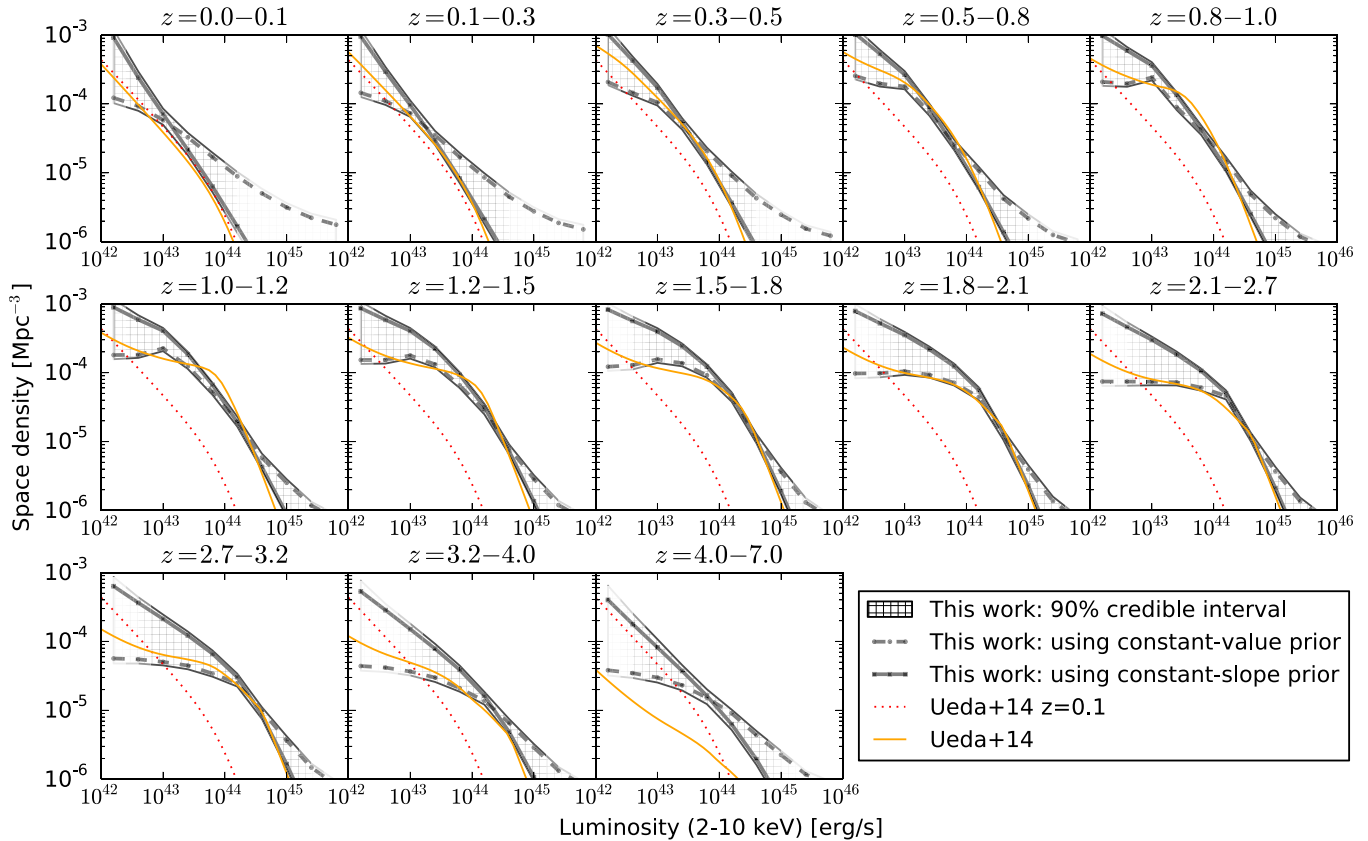


Figure 4. Total X-ray luminosity function in the 2–10 keV spectral band. Each panel corresponds to different redshift intervals. We show our field reconstruction based on a step function in black. The dashed black lines refer to the prior preferring to keep the space density constant, while the solid black line prefers to keep the slope of the density (in z and L) constant. Whenever the data constrain the result well, the space density estimate from either prior prescription (constant-value or constant-slope) is the same. Only when they extrapolate away from any constraints do the results differ. We thus take the difference in the reconstructions as an indication of whether the data or the priors dominate the result. The hatched regions indicate our measure of the uncertainty, using the 10%–90% quantiles of the posterior samples from both priors together. The orange thin solid line shows the reconstruction by Ueda et al. (2014). The dotted red curve is their local ($z = 0.1$) luminosity function kept constant across all panels for comparison. There are important differences between our luminosity function estimates and that of Ueda et al. 2014, discussed in the text.

quantile *hatched regions* as a measure of the uncertainty, by taking together the posterior samples from both priors.

4.1. Total Luminosity Function

Figure 4 shows our non-parametric total XLF, integrated along the N_H axis. Large uncertainties are visible at the bright end at low redshift ($z < 0.3$), where our sample is small due to the limited cosmological volume. At high redshift ($z > 2$), the faint end becomes uncertain as low X-ray fluxes limit the number of detected sources.

Our method captures the general shape, normalization, and evolution with redshift of the XLF as inferred or assumed in previous parametric determinations. In Figure 4, we show in particular the comparison with a recent comprehensive study of the XLF by Ueda et al. (2014). The overall shape of the LF is a double power law with a break or bend at a characteristic luminosity (L_*), the value of which increases with redshift. As found in previous studies, the space density shows a rapid evolution up to around $z \sim 1$ at all luminosities, being most prominent at high luminosities due to the positive evolution of L_* . We find that the positive evolution continues up to $z \sim 3$ (as in Aird et al. 2010), above which the population starts to decline. It is important to emphasize that in our non-parametric approach these features are imposed by the data and not by any assumptions about the functional form of the XLF.

While the general behavior of the population is similar, Figure 4 also shows important differences compared to some previous parametric studies, specifically that of Ueda et al. (2014). At the highest redshift bin ($z = 4-7$), the space density drops sharply toward high redshifts in their XLF reconstruction (see also Kalfountzou et al. 2014; Vito et al. 2014; Civano et al. 2011). However, our reconstruction remains, as suggested by the priors, at space densities comparable to the previous redshift bin $z = 3-4$. In our data, there seems to be no evidence of a steep decline with redshift. The difference may be due to the large uncertainties in the redshift estimates used in this work. At intermediate redshifts, the Ueda et al. (2014) XLF shows a sharp flattening below a luminosity of around 10^{44} erg s $^{-1}$, after which it steepens again at the lowest luminosities probed by the study. This behavior is most apparent at redshifts $z \sim 1-2$ and is only easy to parameterize using the luminosity-dependent density evolution (LDDE) model (e.g., Ueda et al. 2003; Hasinger et al. 2005; Silverman et al. 2008). Our analysis does not support such a behavior. This is most apparent in the redshift range $z = 0.8-2.1$, where our non-parametric field analysis requires a significantly larger space density of AGNs in the critical luminosity range of 10^{43-44} erg s $^{-1}$. Furthermore, we see no strong evidence of a change in the faint-end slope in the individual panels of Figure 4. This may bring into question the need for LDDE to explain the form and evolution of the XLF. As our focus in this paper is on the demographics and evolution

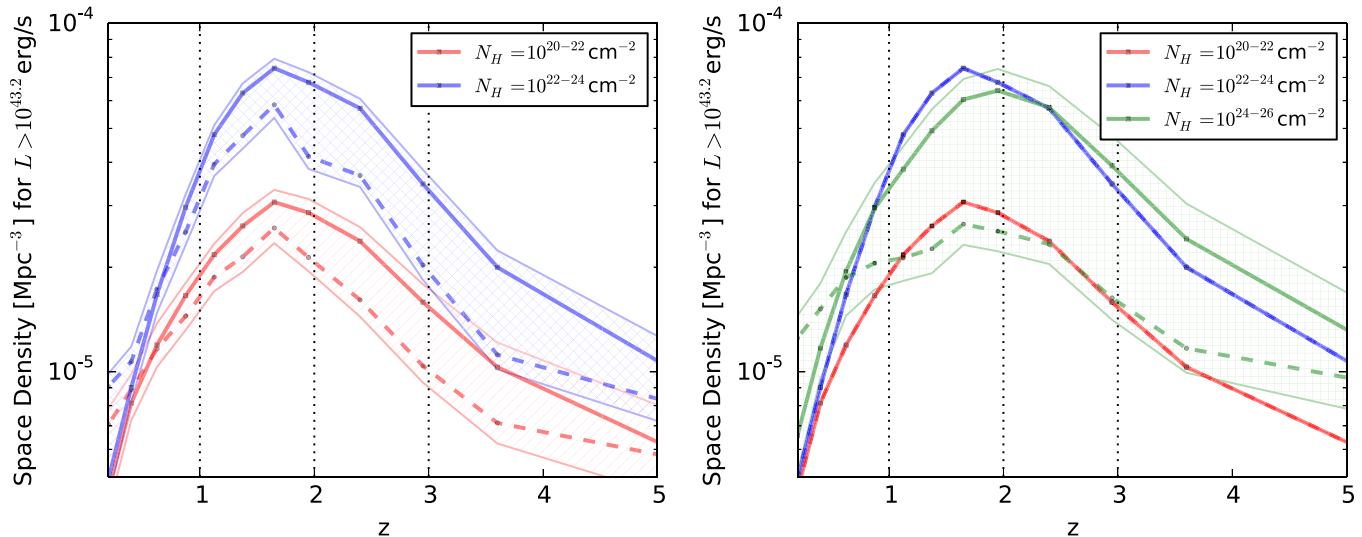


Figure 5. Redshift evolution of the space density of AGNs with $L_X > 10^{43.2} \text{ erg s}^{-1}$, for various column densities. Our step function reconstruction is represented by points at the bin center, which are connected by lines (dashed for the constant-value prior, solid lines for the constant-slope prior). In the left panel, distinct evolutions for the Compton-thin obscured (blue shaded region, top) and unobscured (red shaded region, bottom) can be observed. The right panel plots the evolution of Compton-thick AGNs as a green shaded region. To facilitate the comparison we also plot the evolution of the unobscured and obscured Compton-thin AGN reconstruction in the case of the constant-slope prior (solid lines). All AGN sub-populations split by the level of obscuration experience similar space density evolution, which can be described by a rise from $z = 0.5$ to $z = 1.25$, a broad plateau at $z = 1.25$ – 2.1 , and a decline at higher redshift. There is also evidence that moderately obscured, Compton-thin AGNs ($N_H = 10^{22}$ – 10^{24} cm^{-2}) are evolving faster in the redshift interval 0.5 – 4 in the sense that they reach peak space densities higher than the other AGN sub-populations. The space density of Compton-thick AGNs has the highest uncertainty, due to poor statistics in the low-luminosity range ($L < 10^{44} \text{ erg s}^{-1}$). Nevertheless, there is tentative evidence that the evolution of Compton-thick AGNs is weaker than that of the Compton-thin obscured AGNs (blue), and in fact closer to the evolution of the unobscured AGNs (bottom red solid line).

Table 3
Key Statistics on the Fraction of Obscured and Compton-thick AGNs

	Cosmic Time Average ^a	At $z = 0$ ^{a,b}	Maximum ^{a,c}	z_{Maximum} ^c
Obscured fraction ($>10^{22} \text{ cm}^{-2}$)	$75^{+4}_{-4} \%$	$77^{+4}_{-5} \%$	$83^{+3}_{-3} \%$	>2.25
Compton-thick fraction ($>10^{24} \text{ cm}^{-2}$)	$38^{+8}_{-7} \%$	$39^{+7}_{-6} \%$	$46^{+6}_{-5} \%$	(Unconstrained)

Notes.

^a These fractions relative to the total space density of the population are estimated by integrating the X-ray luminosity function over cosmic time and within X-ray luminosity range $L(2\text{--}10 \text{ keV}) = 10^{43.2}\text{--}10^{46} \text{ erg s}^{-1}$. The uncertainties are computed using the 10% and 90% quantiles of the posterior distribution, i.e., the true value is bracketed with 90% probability.

^b For this estimate we used our lowest redshift bin.

^c The peak is computed by identifying the maximum value (fraction) and location (redshift) in each posterior realization, and considering the distribution of each series. This leads to only an upper limit for the peak location for obscured AGNs. For Compton-thick AGNs, no upper or lower limit could be determined.

of obscuration, rather the LF per se, we defer further discussion of this point to later work.

4.2. Obscured and Compton-thick Fractions

We now report on the fraction of obscured AGNs by comparing the space density above $N_H > 10^{22} \text{ cm}^{-2}$ to the total. Our non-parametric approach allows us to explore this fraction in a model independent way by integrating the LF over luminosity and cosmic time. We computed these fractions using only the luminosity range $L = 10^{43.2}\text{--}46 \text{ erg s}^{-1}$, i.e., without the lowest three luminosity bins. The choice of the luminosity range for the presentation of the results is to minimize uncertainties associated with the typically looser constraints our methodology yields at the faint-end of the XLF. At virtually all redshifts, the AGN space density is better determined at luminosities $L > 10^{43.2} \text{ erg s}^{-1}$. In Table 3, we find that the fraction of obscured objects in the universe is 75%, with narrow uncertainties. The fraction of Compton-thick AGNs ($N_H > 10^{24} \text{ cm}^{-2}$) is approximately 35%.

Additionally, we test whether these fractions are constant over cosmic time. This is done by noting the peak fraction and its redshift in each posterior realization of our reconstruction, and averaging the results. As indicated in Table 3, the obscured fraction tends toward a higher value of almost 85% at $z > 2$. This shows that the fraction of obscured AGNs varies through cosmic time. In the following section, we investigate this evolution of the obscured fraction further. For the Compton-thick fraction, no peak can be identified as the redshift is unconstrained. This indicates that the Compton-thick fraction is constant over cosmic time at approximately 35%.

4.3. Obscuration-dependent Evolution

To explore the evolution of the obscured AGN fraction further, Figure 5 plots the evolution of the space density of unobscured ($N_H = 10^{20}\text{--}22 \text{ cm}^{-2}$), moderately obscured ($N_H = 10^{22}\text{--}24 \text{ cm}^{-2}$) and Compton-thick ($N_H = 10^{24}\text{--}26 \text{ cm}^{-2}$) AGNs with luminosities $L > 10^{43.2} \text{ erg s}^{-1}$. In the left panel, we find that moderately obscured and unobscured AGNs follow similar evolutionary patterns, namely an increase from $z = 0$ to $z = 1.2$




Figure 6. Redshift evolution of space density of AGNs split by the level of obscuration. Different panels correspond to different hydrogen column density intervals as indicated at the top. The plot highlights that the redshift evolution of the AGN space density is strongest for the column density bins $N_{\text{H}} = 10^{22}-10^{23} \text{ cm}^{-2}$ and $N_{\text{H}} = 10^{23}-10^{23.5} \text{ cm}^{-2}$.

where their space density peaks and a decline at higher redshifts. However, there are also differences. Moderately obscured AGNs evolve much faster from $z \sim 0.5$ to $z \sim 1.5$.

The right panel of Figure 5 shows the space density of Compton-thick AGNs. The evolution of the Compton-thick AGN population has larger uncertainties than unobscured and moderately obscured Compton-thin AGNs. Nevertheless we find that their space density shows a broad plateau at $z \approx 1-3$, followed by a decline to both lower and higher redshifts. The evolution of Compton-thick AGNs also appears weaker than that of moderately obscured Compton-thin ones. This behavior is contrary to the N_{H} smoothness prior, which prefers neighboring N_{H} bins to have the same value. Thus it can be concluded that the data drive the result of different evolutions for different obscurations: a strong obscured Compton-thin evolution compared to a weaker evolution in both Compton-thick and unobscured AGNs.

The different evolution of AGNs with different levels of obscuration is further demonstrated in Figure 6, where the AGN sample is split into finer N_{H} bins. All sub-populations experience the same evolutionary pattern, a rise from redshift $z \approx 0.5$, a peak at $z \approx 1.5$, and a decline at higher redshift. The AGNs that undergo the strongest evolution are those with column densities around $N_{\text{H}} = 10^{22}-10^{23.5} \text{ cm}^{-2}$. Both unobscured and Compton-thick AGNs evolve less strongly.

4.4. Luminosity Dependence and Evolution of the Obscured Fraction

Previous studies suggest that the fraction of obscured AGNs is a function of accretion luminosity (e.g., Lawrence 1991; Ueda et al. 2003, 2014; Simpson 2005; Akylas et al. 2006; Silverman et al. 2008; Burlon et al. 2011). The luminosity dependence of the obscured fraction is thought to be related to the nature of the obscurer, e.g., reducing it in physical extent. In this section, we investigate this issue using our model-independent approach. To illuminate the behavior of the obscured population, we analyze the behavior of Compton-thin and Compton-thick AGNs separately. This section only considers the obscured AGNs that are Compton thin ($N_{\text{H}} < 10^{24} \text{ cm}^{-2}$). To this end, we define a new quantity, namely the obscured fraction of the Compton-thin AGNs (Compton-thin obscured fraction (CTNOF))

$$\text{CTNOF} := \frac{\phi[N_{\text{H}} = 10^{22-24} \text{ cm}^{-2}]}{\phi[N_{\text{H}} = 10^{20-24} \text{ cm}^{-2}]} \quad (2)$$

In the top row of Figure 7, we find that the CTNOF is a strong function of luminosity. There is evidence for a peak at a certain luminosity and decline at both brighter and fainter luminosities (the red star symbol provides a constant reference point). Interestingly, the luminosity where the obscured AGN fraction peaks appears to be a function of redshift. With increasing redshift, the drop of the CTNOF at bright luminosities occurs at higher luminosities. These results are in some agreement with the recent analysis in Ueda et al. (2014). One difference, however, is at the brightest luminosities, where our non-parametric method tends toward a higher value (50%). Similar high obscured fractions were suggested by Iwasawa et al. (2012) and Vito et al. (2014), as opposed to 20% in the local universe Burlon et al. (2011). Ueda et al. (2014) have better statistics at bright luminosities compared to our work because they included more wide-area and shallow survey fields in the analysis. It is therefore likely that our sample has poor statistics at the brightest luminosities, which makes the preference of the N_{H} prior toward equipartition apparent. The strength of our data is rather at the faint end of the LF for moderate and high redshifts. There we find a significant turnover at low luminosities (right top panel of Figure 7). This result is independent of the adopted form of the prior, indicating that this behavior is strongly imposed by the data. A similar behavior of a peak luminosity and a turnover have been found in local samples (Burlon et al. 2011; Brightman & Nandra 2011b; top left panel of Figure 7, corrected assuming a constant 20% Compton-thick fraction in their sample). The position of the peak they find is consistent with our results at low redshift (red star symbol, $L = 10^{43} \text{ erg s}^{-1}$).

When considering the obscured fraction at $L = 10^{44} \text{ erg s}^{-1}$, the evolution of the peak is imprinted as a rise with increasing redshift. This is shown in Figure 8, with the same data points from literature. Observers considering mostly AGNs with luminosity $L \geq 10^{44} \text{ erg s}^{-1}$ will see an increase in the fraction of obscured AGNs with redshift as shown here, while observers considering intrinsically faint AGNs ($L \leq 10^{43} \text{ erg s}^{-1}$) would observe the opposite trend (see Figure 7).

4.5. Evolution of Compton-thick AGNs

Compton-thick objects have been hypothesized to play a major role in the accretion phase of AGNs. The identification of such sources in current *XMM-Newton* and *Chandra* is challenging and therefore their space density as a function of redshift has remained controversial. It is therefore

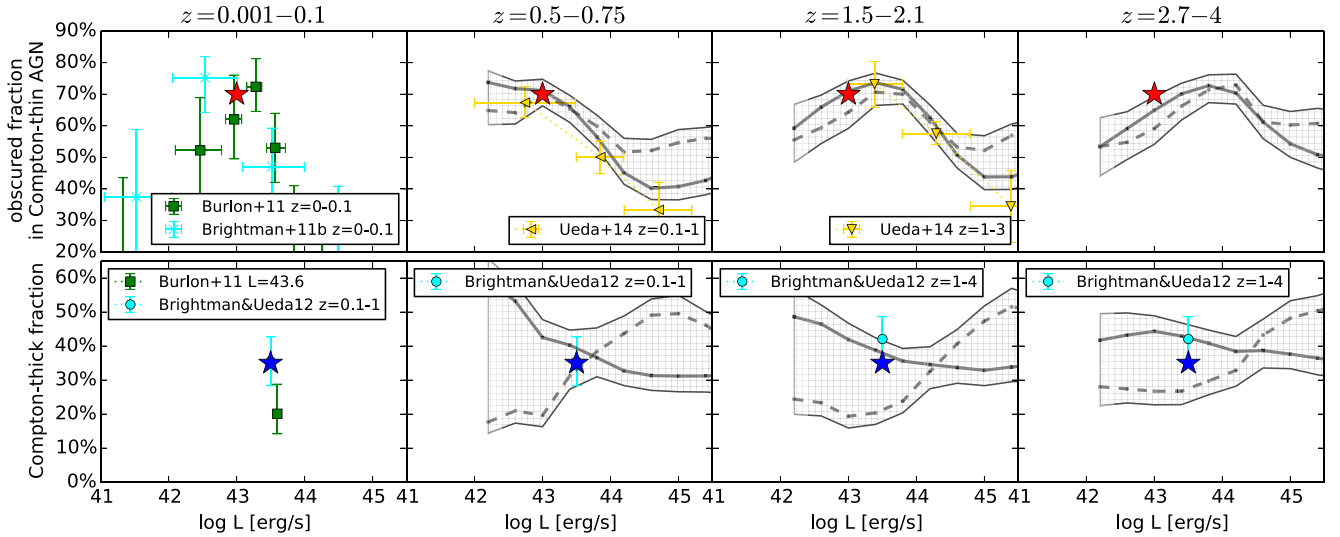


Figure 7. Luminosity dependence of the obscurer. The top row plots the obscured fraction of Compton-thin AGNs (CTNOF, Equation (2)) for various redshift intervals. The shaded gray regions are the constraints from our non-parametric method. Additionally, we compare to the work of Ueda et al. (2014; yellow points). For reference, the red star symbol is placed at 70% and $L = 10^{43} \text{ erg s}^{-1}$ across the panels. In the top left panel, the results from local surveys (Burlon et al. 2011; Brightman & Nandra 2011b) which report a similar shape. The CTNOF shows a distinct peak, which is placed at the red star for local surveys, but appears to move to sequentially higher luminosities at higher redshift. In the bottom row, the luminosity dependence of the Compton-thick fraction (Equation (3)) is plotted. Our results (shaded gray) show that the Compton-thick fraction is compatible with being constant at $\sim 35\%$ (blue star symbol for reference at $L = 10^{43.5} \text{ erg s}^{-1}$). For comparison, previous surveys (Burlon et al. 2011; Brightman & Ueda 2012) are shown (see the text).

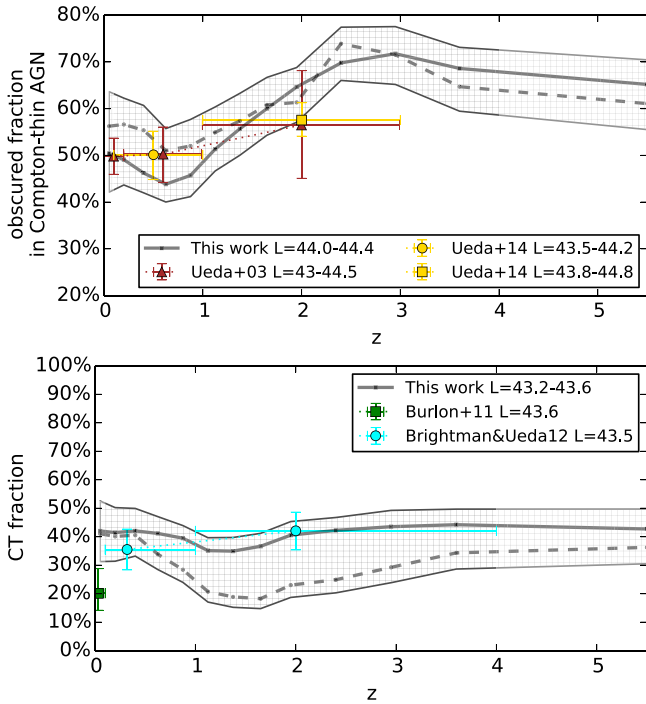


Figure 8. Evolution of the obscured fraction. Top panel: the evolution of the obscured fraction of Compton-thin AGNs (CTNOF, Equation (2)) is plotted as a gray shaded region for $L > 10^{44} \text{ erg s}^{-1}$. At $z \sim 3$, the obscured fraction was higher (75%) than today (50%). For comparison also plotted are the constraints from Ueda et al. (2003; brown triangles) and Ueda et al. (2014; yellow points). Bottom panel: the evolution of the Compton-thick fraction at $L = 43.5$. Our results are compatible with a constant Compton-thick fraction of $\sim 35\%$. We compare our results to Brightman & Ueda (2012; cyan circles). Local hard X-ray *Swift*/BAT observations (green squares, Burlon et al. 2011 for $z < 0.1$) derived a lower Compton-thick fraction of $\sim 20\%$.

important to place constraints on their space density of these sources in the context of previous studies. Figure 9 shows the evolution of the space density evolution of Compton-thick sources inferred by our non-parametric methodology. Three




Figure 9. Evolution of the density of Compton-thick AGNs. Different panels correspond to X-ray luminosities brighter than $10^{43.2}$ (top), $10^{43.6}$ (middle), and 10^{44} (bottom) in units of erg s^{-1} . Comparing our reconstruction (black) to previous results, we find overall good agreement. The result of Fiore et al. (2009) using the $L > 10^{43.5} \text{ erg s}^{-1}$ cut is an exception. Their estimate lies higher compared to our results, potentially due to a starburst galaxy or moderately obscured AGNs contaminating their sample (see the text).

luminosity cuts are shown that allow for a direct comparison with previous studies (Fiore et al. 2008, 2009; Alexander et al. 2011), which select Compton-thick AGNs at infrared wavelengths. Figure 9 shows that the general trend is a decline

of the space density of this population with decreasing redshift. Above $z = 2$, a decline appears toward increasing redshifts. Also, our results are in rough agreement with previous estimates (Fiore et al. 2008; Alexander et al. 2011). The results of Fiore et al. (2009) on moderate luminosity ($L \geq 10^{43.5} \text{ erg s}^{-1}$) Compton-thick AGNs in the COSMOS field is an exception. The Compton-thick AGN space density determined in that study is significantly higher than our result. This may be attributed to contamination of obscured AGN samples selected in the infrared by either dusty starburst or moderately obscured Seyfert galaxies (e.g., Georgakakis et al. 2010; Donley et al. 2010).

To study the contribution of Compton-thick AGNs to the total accretion luminosity output it is necessary to characterize the luminosity dependence of Compton-thick AGNs. The bottom panel of Figure 7 plots the Compton-thick fraction, defined as

$$\text{Compton-thick fraction} := \frac{\phi[N_H = 10^{24-26} \text{ cm}^{-2}]}{\phi[N_H = 10^{20-26} \text{ cm}^{-2}]}, \quad (3)$$

as a function of X-ray luminosity. Within the uncertainties we do not find evidence for a luminosity dependence. If the constant-slope prior is preferred (solid line), a similar behavior as in the obscured fraction is allowed. However, our results are consistent with a constant Compton-thick fraction at $\sim 35\%$. This is the first time the luminosity dependence of the Compton-thick fraction is constrained.

In Figure 8, the bottom panel shows the evolution of the Compton-thick fraction at $L = 10^{43.5} \text{ erg s}^{-1}$. Our results show a minor dip in the $z \sim 1-2$ range. This is due to the strong peak in Compton-thin sources (see above), causing a strong rise in the denominator of the fraction. The blue data point at $z < 0.1$ was taken from the *Swift*/Burst Alert Telescope (BAT) analysis of Burlon et al. (2011), where a considerably lower Compton-thick fraction was found, in disagreement to our findings. This may be because of the small volume of our sample at those redshifts, or differences in the analysis and in particular the methods adopted to correct for incompleteness of the AGN samples in the Compton-thick regime. Figure 8 also compares our results with those of Brightman & Ueda (2012) in the CDFS (cyan). Within the uncertainties, we find broad agreement. However, Brightman & Ueda (2012) argued in favor of a redshift evolution (increase) of the Compton-thick fraction based on their constraints and those of *Swift*/BAT (Burlon et al. 2011). In our result, we find no evidence for such a trend. Our uncertainties are compatible with a constant Compton-thick fraction of $\sim 35\%$.

4.6. The Contribution of Obscured AGNs to the Luminosity Density

Due to the difficulty in capturing the entire population of AGNs (see the previous section), the accretion history of SMBHs has remained uncertain in the literature. Through the relation of mass accretion rate in black holes to luminosity, the total luminosity output of AGNs can be used to study the total accretion history of the universe. We compute the total X-ray luminosity output at any particular redshift, i.e., the luminosity density, by integrating the LF as $\int L_X \cdot \phi \, d \log L_X$. This is shown in the top of Figure 10 (gray). Here, we again use $L > 10^{43.2} \text{ erg s}^{-1}$, which captures essentially all of the luminosity density and focuses on the redshift range $z = 0.5-4$ where our data coverage is best. The luminosity density shows a broad peak in the $z = 1-3$ range, with a decline to both high and low redshifts.




Figure 10. Evolution of the X-ray luminosity density of AGNs with $L_X > 10^{43.2} \text{ erg s}^{-1}$, for various column densities. The luminosity output of AGNs experiences a rise and fall in density in the $z = 1-3.5$ range (total as the top gray shaded region). The strongest contribution to the luminosity density is due to obscured, Compton-thin (blue shaded region) and Compton-thick AGNs (green shaded region), which contribute in equal parts to the luminosity. In contrast, the emission from unobscured AGNs (red shaded region, bottom) is distinctively smaller.

Additionally, in Figure 10 we probe the contribution of AGNs split by their obscuration (unobscured, Compton-thin obscured, and Compton-thick). One of our most striking results is that Compton-thick AGNs and obscured Compton-thin AGNs contribute in equal parts to the luminosity density in the $z = 1-3$ range. These obscured AGNs contribute the majority ($74_{-5}^{+4}\%$), while unobscured AGNs only play a minor role. The shapes in evolution are fairly similar. However, the evolution of unobscured AGNs is weaker compared to the Compton-thin obscured AGNs when considering the $z \sim 1-1.5$ range.

4.7. N_H Distribution

This medium luminosity interval, $L \approx 10^{43.5} \text{ erg s}^{-1}$, where most of the evolution occurs, is also where the obscuration peaks according to Figure 7. To further illustrate the evolution there, we look at how the column-density distribution evolves. Figure 11 plots the intrinsic N_H distribution for three redshift intervals in panels. Here, the boxes indicate the upper and lower 90% quantile on the fraction of sources in the respective N_H bin. The dashed line illustrates a possible N_H density distribution that fits these fractions. When comparing the top panel to the middle panel in Figure 11 in particular, the main effect appears to be that the fractions at $N_H \approx 10^{23} \text{ cm}^{-2}$ increase toward higher redshifts.

5. DISCUSSION

We combined deep and wide-area X-ray surveys conducted by *Chandra* and *XMM-Newton* to initially constrain the space density of X-ray-selected AGNs as a function of accretion luminosity, obscuring column density and redshift. Our methodology accounts for the different sources of errors. For example, we include in the analysis uncertainties, both random and systematic, associated with photometric redshift measurements. We also account for the lack of redshift information for sources




Figure 11. Column-density distributions at various redshifts. Based on our field reconstruction, we compute the fraction of sources in column density bins (red: “unobscured”; blue: Compton-thin obscured; green: Compton-thick). The boxes show the credible regions containing 90% of the posterior probability. The dashed line illustrates a possible N_{H} density distribution that fits these fractions.

without optical or infrared counterparts. For the determination of the obscuration and intrinsic luminosity of individual sources we use the X-ray spectral analysis method of Buchner et al. (2014), which takes into account both the Poisson errors of X-ray spectra and photometric redshift errors. The spectral analysis also uses a physically meaningful, multi-component model that is consistent with recent ideas and observations on dominant AGN emission processes and the structure of matter in the vicinity of active SMBHs. Another important feature of the analysis presented in this paper is the non-parametric method developed to determine the XLF of AGNs. This allows us to explore how the space density of AGNs depends on luminosity, redshift, and column density without imposing any model. This frees us from any assumptions on the dependence of the AGN space density to luminosity, redshift, and column density.

We find that obscured AGNs ($N_{\text{H}} > 10^{22} \text{ cm}^{-2}$) dominate the population of active SMBHs at all redshifts. This is in agreement with previous investigations (e.g., Ueda et al. 2003, 2014; Akylas et al. 2006; Merloni & Heinz 2013), although our analysis highlights this point more robustly and quantitatively (with uncertainties). We show that about 75% of the AGN space density, averaged over redshift, corresponds to sources with column densities $N_{\text{H}} > 10^{22} \text{ cm}^{-2}$ (see Table 3). The contribution of obscured AGNs to the accretion density of the universe is similarly large (75%). The bulk of the black hole growth across cosmic time is therefore taking place behind large column densities of gas and dust clouds.

5.1. The Role of Compton-thick AGNs

We are also able to place constraints on the number fraction of the most heavily obscured, Compton-thick sources to the AGN population, finding it to be $38_{-7}^{+8}\%$ of the total population. Results from previous AGN surveys have been divergent due to the difficulty of identifying Compton-thick AGNs (e.g., 15%–20% in Akylas & Georgantopoulos 2009, 5%–20% in Burlon et al. 2011 for the analyses of *Swift*/BAT surveys), and relatively low compared to the requirements from the XRB (see below). Our constraints on the Compton-thick fraction are in good agreement with the estimates of Brightman & Ueda (2012) of $\sim 35\%$ – 40% . However, in their work they concluded an evolution of the Compton-thick fraction by contrasting their high-redshift data to a local survey, which reported a significantly lower fraction (specifically Burlon et al. 2011). In our analysis, we do not find any evidence for a redshift evolution of the Compton-thick fraction. However, our sample lacks large, shallow fields that can probe the local universe, and thus our results may also benefit from being combined with a local estimate. For instance, the 2–10 keV X-ray-based work of Risaliti et al. (1999) estimated that 50% of all Seyfert 2 galaxies are observed with a Compton-thick line of sight. If this estimation is taken, it is in agreement with our results without the need for any evolution. However, hard X-ray surveys have reported significantly lower values, e.g., $20_{-6}^{+9}\%$ in Burlon et al. (2011), 9%–17% in (Bassani et al. 2006; Malizia et al. 2012; Vasudevan et al. 2013). This shows that estimates for a local Compton-thick fraction have been diverse, making it difficult to make any claim of an evolutionary trend. One possible source of uncertainty is the sensitivity to sources with $N_{\text{H}} = 10^{25-26} \text{ cm}^{-2}$. In this work, we have explicitly assumed that their space density is the same as those of sources with $N_{\text{H}} = 10^{24-25} \text{ cm}^{-2}$. The sparse sampling (only three secure objects in our entire sample) of these heavily buried population prohibits strong inferences.

For discussing the accretion of luminosity onto obscured AGNs, it is noteworthy that the fraction of Compton-thick AGNs does not show a luminosity dependence (Figure 7). When considering currently accreting AGNs, Compton-thick AGNs appear to play an invariant role with luminosity and cosmic time. However, the large uncertainties in the luminosity dependence do not allow firm conclusions. As Figure 5 shows, the evolution of Compton-thick AGNs may follow the unobscured AGNs closely in shape, rather than the Compton-thin obscured AGNs, which evolve strongly. In the former case Compton-thick AGNs can be modeled as $35\%/25\% = 1.4$ times more abundant than unobscured AGNs over cosmic time and luminosity. In the latter case, they would be approximately equal in abundance and follow the luminosity dependence and its evolution (see Section 4.4, discussed below in Section 5.3).

Measuring the luminosity density allows a view of the importance of obscuration with regard to the accretion history of SMBHs. Soltan (1982) originally argued that a major contribution to the accretion has to be accumulated in obscured AGNs, which were missing in surveys at the time. Albeit that obscured AGNs have been detected in great numbers since then, the same argument has been extended to Compton-thick AGNs, whose fraction has remained controversial. Our analysis shows that the accretion growth of SMBHs is indeed dominated by Compton-thick and Compton-thin AGNs ($74_{-5}^{+4}\%$), while unobscured AGNs only play a minor role. However, we also find Compton-thick and Compton-thin AGNs each contributing half of the luminosity output (see Figure 10), with the

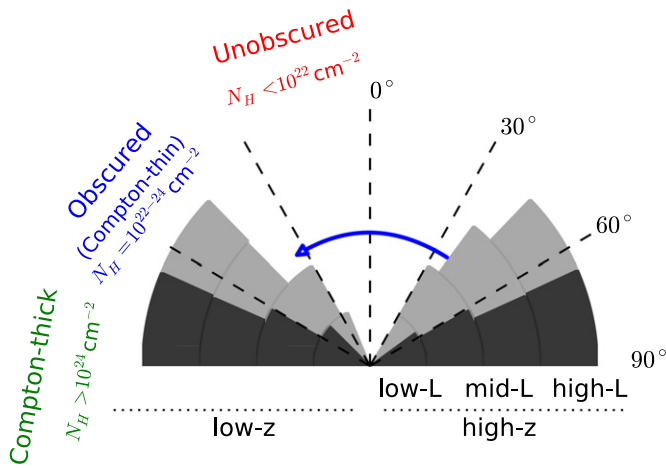


Figure 12. Illustration of the population-averaged obscuring toroid. The obscured and Compton-thick fraction and their luminosity dependence is shown using a corresponding opening angle above which the sky is obscured to the X-ray emitter (center). The luminosity dependence is depicted for low-redshift (left) and high-redshift (right) AGNs using radial shells: as the luminosity increases, more sky is visible. Notice that the difference between a low-redshift and high-redshift source does not affect the lowest and highest luminosities, but only where the transition occurs (blue arrow). The opening angle corresponding to the Compton-thick transition remains roughly constant. For accurate numbers, refer to Figures 7 and 8.

Compton-thick AGNs contributing $40^{+6}_{-6}\%$. This shows that Compton-thick AGNs are an important contributor to the accretion history. Our analysis places, for the first time, tight constraints on the Compton-thick contribution. Results on the Compton-thick fraction from studies of the cosmic XRB have varied between 9% (Treister et al. 2009), 45%–30% (Gilli et al. 2007; luminosity dependent) and 28%–60% (Ueda et al. 2014, not constrained), while Shi et al. (2013) reported 38%. This scatter is also due to the fact that XRB fitting involves other parameters that are degenerate with the Compton-thick fraction (e.g., additional reflection, see Akylas et al. 2012). It is worth emphasizing that our estimate on the Compton-thick fraction is higher than previous survey studies, but in agreement to those values assumed in XRB synthesis analyses.

The overall evolution of the accretion density for the AGN population is presented in Figure 10. There, the peak is in a broad plateau at $z = 1$ –3, with a drop-off toward high redshift. This trend may corroborate claims that the accretion onto SMBHs correlates with the star formation history of galaxies (see Merloni & Heinz 2013). Unlike previous works, our results do not a priori assume any form of the evolution at high redshift.

Having constrained the Compton-thick fraction well, we can speculate on the origin of this obscuration. The fraction of Compton-thick sources is considerably large (35%), necessitating that a large fraction of viewing angles is obscured ($\sim 20^\circ$, see Figure 12). Compton-thick obscuration is associated with the torus, as such column densities are not typically reached by galactic gas. As a Compton-thick, smooth obscuration would be unstable (Krolik & Begelman 1988), the current working hypothesis is that the obscuration comes in discrete clouds. Under this clumpy torus model obscured views are produced when clouds are encountered in the line of sight. If the Compton-thin obscuration is also associated with the torus, Compton-thick AGNs are a simple extension of Compton-thin obscured AGNs. In principle it would also be possible then that Compton-thick AGNs have just a larger number of clouds in a clumpy torus. In this case, their cumulative line-of-sight obscuration would

provide a Compton-thick view. Alternatively, these objects may have denser clouds overall. Another possibility is that the cloud density is distributed such that both regimes are covered. For further research in this regard, the unbiased column-density distributions shown in Figure 11 can provide observational constraints for clumpy torus models. The alternative scenario is that Compton-thin obscuration is associated with a medium of different extent than Compton-thick obscuration (e.g., through galactic and nuclear gas; Matt 2000). We discuss both kinds of models in Section 5.3 when also considering the luminosity dependence.

5.2. Obscuration-dependent Evolution

In Figure 5, we show that the space density of obscured Compton-thin AGNs experience a much stronger rise in the redshift interval $z = 0.5$ –1.2 compared to unobscured AGNs. This demonstrates an observed obscuration-dependent evolution. Our method explicitly contains a smoothness prior, preferring that the space density of AGNs as a function of column densities is the same. In contrast, the results show strong differences in the evolution of unobscured and obscured Compton-thin AGNs. We can thus conclude that this result was driven by the data.

In Figure 6, we find that the fastest evolving population is the one with column densities $N_H = 10^{22}$ – $10^{23.5}$ cm^{-2} . A possible interpretation of these trends is that different levels of obscuration correspond to media with different spatial extents. We investigated this further by focusing on the behavior of moderately obscured AGNs ($N_H = 10^{22}$ – 10^{24} cm^{-2}).

5.3. Luminosity Dependence and Evolution of the Obscured Fraction

An important result from our analysis is that the obscured Compton-thin AGN fraction (CTNOF, Equation (2)) depends on both luminosity and redshift (see Figure 7). The luminosity dependence can be described by a peak of the CTNOF at a given luminosity followed by a decline at both brighter and fainter luminosities. The redshift dependence is manifested by a shift to brighter luminosities of the peak of the obscured AGN fraction with increasing redshift. The shift of this peak causes observers considering mostly AGNs with luminosity $L \geq 10^{44}$ erg s^{-1} to see an increase in the fraction of obscured AGNs with redshift (as shown in Figure 8). Previous studies also find that the fraction of obscured AGNs decreases with increasing luminosity (e.g., Ueda et al. 2003, 2014; Akylas et al. 2006). When we consider a luminosity range where we have good constraints ($L = 10^{43.2}$ – 10^{46} erg s^{-1}), the obscured fraction rises with redshift so that at $z > 2.25$, $83^{+3}_{-3}\%$ of AGNs are obscured ($N_H > 10^{22}$ cm^{-2}), as compared to the local $z = 0$ value $z = 75^{+4}_{-4}\%$.

Our analysis also establishes a decline of the obscured AGN fraction at low luminosities. This trend has only been found in the local universe (Burlon et al. 2011; Brightman & Nandra 2011b). Figure 7 also leads to the conclusion that observers who would consider only intrinsically faint AGNs ($L \leq 10^{43}$ erg s^{-1}) would observe a decrease in the fraction of obscured AGNs with redshift. In other words, the magnitude of the trend is determined by the sample selection.

Obscured models that do not invoke any obscuration dependence are ruled out by our findings. These include the simple torus where the obscuration distribution is produced purely geometrically. Despite these shortcomings, this picture may serve




Figure 13. Four models that address the luminosity dependence of the obscured fraction. This categorization distinguishes by the luminosity-dependent effect, the increase of the obscured fraction (left column), and the decrease of the obscured fraction (right column). The top of Figure 7 is repeated here and simplified. In the four models considered (see the text for details), the relevant obscurer is either the torus (top row) or galactic obscuration (bottom row). In the illustration of the models, the red star represents the X-ray source, while gray clumps illustrate the cold obscurer.

as a visualization of the observed trends. The luminosity dependence and evolution of obscured/Compton-thick AGNs is illustrated in Figure 12. If we assume a toroidal geometry for the obscurer shared by the whole population, we can define an opening angle that reproduces the obscured/Compton-thick fraction. The luminosity-dependent increase of this opening angle from 30° to 40° is depicted using radial shells for low-redshift (left) and high-redshift (right) AGNs. Notice that the difference between a low-redshift and high-redshift source does not affect the lowest and highest luminosities, but only where the transition occurs (blue arrow). Thus, it is not the case that high-redshift sources merely have more obscuration overall. If this were the case, a boost in the obscured fraction would be visible over all luminosities. Instead, the peak of the obscured fraction has moved to lower luminosities, thus reducing the obscurer in moderately bright AGNs (see Figure 7 for accurate numbers).

We now consider a few exemplary models that can qualitatively account for the decrease of the obscured fraction toward low luminosity or the decrease toward high luminosities. An additional uncertainty is whether the relevant obscuration occurs on small or large scales, i.e., in the torus or on galactic scales. In the following, we present for both scenarios a model each for the high-luminosity and the low-luminosity effect. These are illustrated in Figure 13.

The receding torus (RT) model is often invoked to explain the observed decrease of the obscured AGN fraction at bright luminosities. This scenario postulates a toroidal geometry for the obscurer with opening angle that increases with increasing luminosity. This is attributed to photon pressure pushing away the obscuring material, photo-ionization of the gas clouds, or sublimation of the dust by the photon field of the active black hole (Lawrence 1991; Simpson 2005; Akylas & Georgantopoulos 2008). Such a variation of the scale height of the torus has

been claimed via infrared observations (Maiolino & Rieke 1995; Lusso et al. 2013; Toba et al. 2014). In the context of clumpy torus models (Hönig & Beckert 2007; Nenkova et al. 2008), a more intense radiation field leads to the (partial) ionization of individual dust and gas clouds, thereby increasing the effective solid angle of unobscured sightlines. This model (RT) thus represents a direct, causal connection between the X-ray luminosity and the unobscured line of sight. This scenario is illustrated in the top right of Figure 13. In the RT model it is, however, difficult to explain the observed evolution of the luminosity dependence. The main evolutionary effect observed is an increase of the turnover luminosity with redshift, causing the onset of the high-luminosity effect at subsequently higher luminosities (see Figure 7). For a purely nuclear obscurer scenario such as RT, no such evolutionary effect is predicted.

Alternatively, it is also possible to interpret the observed luminosity dependence of the obscured AGN fraction in the context of evolutionary models, in which different levels of obscuration roughly correspond to different stages of the growth of SMBHs. Most black hole formation models include an early stage of gas inflows caused by mergers (e.g., Hopkins et al. 2006a; Somerville et al. 2008) or secular processes (e.g., Fanidakis et al. 2012; Bournaud et al. 2007; Ciotti & Ostriker 1997, 2001). During this period SMBHs grow fast by accreting material close to the Eddington limit and are also typically obscured by the inflowing material. Eventually, however, the energy output of the central engine becomes sufficiently powerful to drive outflows, which can blow away the obscuring clouds of dust and gas. The central engine then shines unobscured for a brief period before its luminosity output declines as a result of the depletion of the available gas reservoirs. It is further proposed (e.g., Hopkins et al. 2005, 2006b) that luminous AGNs brighter than the break of the XLF

are dominated by fast accreting black holes close to the peak of their growth phase. Lower luminosity sources include a large fraction of AGNs during the decline stage of their activity.

In the above scenario the decrease of the obscured AGN fraction with increasing luminosity may be linked to the blow-out stage of AGNs. This scenario (Eddington-limit blowout; EB) is illustrated in the bottom right panel of Figure 13. At brighter accretion luminosities, it is more likely that AGN-related feedback mechanisms become more efficient, thereby shedding the gas and dust cocoons in the vicinity of SMBHs. Lower accretion luminosities, below the break of the XLF, include a large fraction of AGNs past the peak of their activity, when the gas reservoirs that feed and obscure the black hole have already been depleted. Therefore, one might expect a larger fraction of unobscured AGNs at low accretion luminosities. The evolutionary picture outlined above is therefore consistent with the observed luminosity dependence of the obscured AGN fraction: a maximum at a given luminosity and a decline at both brighter and fainter luminosities.

In the EB model, the observed redshift dependence of the turnover luminosity is also difficult to explain. Unless the physical conditions of the obscuring material are a function of cosmic time, it is not obvious how the physical mechanisms that affect the geometry of the obscurer on small scales can produce a luminosity-dependent obscured AGN fraction that changes with redshift.

A rescuing argument for the EB scenario is possible. It is now fairly established that the black hole mass of the average, currently accreting AGNs increases with redshift (Merloni 2004; Merloni & Heinz 2013). This means that at high redshift, the typical black hole undergoing accretion has already accumulated a large mass, while at low redshift, small black holes undergo accretion. Recent simulations have shown that such an anti-hierarchical growth is not contradictory to hierarchical structure formation (e.g., Fanidakis et al. 2012; Hirschmann et al. 2012; Enoki et al. 2014), and can explain the decrease of the peak of the LF. Under the EB scenario, the peak of the LF is linked to the Eddington luminosity, as brighter AGNs remove their obscuration. If the black hole mass increases with redshift, so does the Eddington luminosity, causing the onset of the high-luminosity decrease of the obscured fraction to occur at higher luminosities. The EB scenario is thus compatible with our observations, if the average black hole mass increases with redshift.

This model can affect both the small-scale, nuclear obscurer as well as the gas in the galactic vicinity of the black hole, and our observations may be incapable of distinguishing the effects. For this reason, the illustration of this effect (EB) in Figure 13 might be placed in both categories.

Let us now consider the low-luminosity effect, which is a decrease of the obscured fraction toward low luminosities. The evolutionary scenario might also be capable of explaining the low luminosity increase of the obscuration (Hopkins et al. 2006a). If the observed column density is due to galactic streams, feeding the black hole simultaneously obscures it. One would then expect a simple correlation of the accretion luminosity and the obscuration, at least at luminosities below the Eddington limit. This effect (gas available; GA) is illustrated in the bottom left of Figure 13.

Under the GA scenario (bottom left of Figure 13), the accretion luminosity is simply correlated to the available gas. While observations have shown an overall reduction of the gas and dust content of galaxies over cosmic time (Tacconi et al.

Table 4
Predictions of the Discussed Models

Model ^a	Obscuration ^b	Low- <i>L</i> Effect ^c	High- <i>L</i> Effect ^d	Evolution ^e
RT	Nuclear	✓		BHM ^f
CP	Nuclear		✓	No
EB	Nuclear or galactic		✓	BHM ^f
GA	Nuclear	✓		No
RT+CP	Nuclear	✓	✓	Not for high- <i>L</i>
GA+EB	Galactic	✓	✓	Not for low- <i>L</i>
CP+EB	Nuclear or both	✓	✓	BHM ^f

Notes.

^a Model name as presented in Figure 13.

^b Whether the luminosity-dependent layer of obscuration is associated with galaxy-scale obscuration (“galactic”) or the “torus” (“nuclear”).

^c Whether the model predicts a decrease of the fraction of obscured AGNs toward low luminosities.

^d Whether the model predicts a decrease of the fraction of obscured AGNs toward high luminosities.

^e Whether the model predicts any evolution of the effect with cosmic time.

^f The model scales with black hole mass. Evolution is predicted if the average accreting black hole mass changes over cosmic time.

2013), under the GA model this would increase the total number of AGNs at high luminosities, but the luminosity dependence of the obscured fraction would remain unaffected. In particular, the GA model does not predict the observed increase of the turnover luminosity as suggested by Figure 7.

The low-luminosity increase of the obscured fraction has been observed before in local samples (Burlon et al. 2011; Brightman & Nandra 2011b; Elitzur & Ho 2009). For these sources, it is more probable that the obscurer is nuclear, as many of these sources provide the opportunity to study both the galaxy and the infrared emission of the torus. One model is that the obscurer is caused by clumps in the disk wind (Elitzur & Shlosman 2006). This scenario was originally created to explain the vertical support for a cold dusty torus that would otherwise collapse (Krolik & Begelman 1988). In the disk wind model, the accretion disk is only capable of projecting clouds beyond a certain luminosity, predicting the absence of a clumpy torus in low-luminosity AGNs. This scenario of cloud production (CP) is illustrated in the top left of Figure 13.

While the CP model reproduces the necessary increase of the obscured fraction, it does not predict any evolution over cosmic time when taken at face value. However, under this model, the critical luminosity below which the obscurer cannot be sustained is strictly determined by the black hole mass (Elitzur & Shlosman 2006; Elitzur 2008). This model has received observational support in Elitzur & Ho (2009), who observe a black hole mass dependence of the luminosity dependence from a local survey. This model, taken together with the black hole mass evolution explained above, can thus explain the observed evolution.

To conclude, we have considered four effects (simple models) that are summarized in Table 4. Individually, they only partially explain the observations. However, the CP model, taken together with the EB scenario, can explain the observations if the black hole mass function evolves with redshift. Under this view, the obscurer is vertically extended above a certain luminosity, and the torus disappears when the Eddington luminosity drives away the obscurer, unveiling the bright X-ray source for a (cosmically) brief time. Here it remains uncertain whether the obscuration removed is galactic, nuclear, or both. Both effects engage at

Table 5
Space Density Estimates as a Function of Luminosity
Redshift and Obscuration

$\log L$	z	$\log N_{\text{H}}$	$\log \phi$
44.0–44.4	2.1–2.7	20.0–21.0	–5.583–5.289
44.0–44.4	2.1–2.7	21.0–22.0	–5.439–5.189
44.0–44.4	2.1–2.7	22.0–23.0	–5.157–5.936
44.0–44.4	2.1–2.7	23.0–23.5	–4.917–4.692
44.0–44.4	2.1–2.7	23.5–24.0	–5.128–4.862
44.0–44.4	2.1–2.7	24.0–26.0	–5.251–4.945

(This table is available in its entirety in machine-readable form.)

a luminosity that is dependent on the black hole mass, which provides the cosmic evolution effect.

Future observations may shed light on the dependence of the galactic gas with the AGN X-ray luminosity. For instance, the Atacama Large Millimeter/Submillimeter Array (Beasley et al. 2006) measurements of gas fractions in host galaxies of unobscured and obscured AGNs can be compared. If they do not differ, the GA model is ruled out. A similar approach by comparing luminous and faint AGNs can also probe the predictions of the EB model, and establish whether the obscuration is nuclear. Under the EB model, gas motions are expected in bright AGNs, which may be detectable with future X-ray spectroscopes on the *ATHENA* mission (Nandra et al. 2013). Furthermore, to establish that the evolutionary trend is caused by black hole mass evolution, black hole mass matched samples of luminous AGNs may be considered. If such a sample shows no redshift evolution in the luminosity dependence of the obscured fraction, then the trend is due to black hole mass differences. The CANDELS fields with their high-resolution near-infrared imaging is the best candidate for such a study.

The models as presented here are a rudimentary description of the involved physical processes and require further refinement in their predictions and self-consistency, e.g., via numerical simulations. Among the challenges in this regard is that the nuclear obscuration cannot be resolved in evolutionary models (e.g., Hopkins et al. 2006a; Somerville et al. 2008) and thus the torus is not yet treated self-consistently. Recently, it has become possible to self-consistently model the radiative and hydrodynamic processes that maintain the torus (e.g., Wada 2012). Further research is needed in terms of whether the torus can, by itself, reproduce the luminosity dependence of the obscured fraction. Also, it has to be clarified which processes maintain a cold, Compton-thick torus, with the geometric extent implied by a 35% Compton-thick fraction.

6. CONCLUSIONS

We combined deep and shallow, wide-area X-ray surveys conducted by *Chandra* and *XMM-Newton* to constrain the space density of X-ray-selected AGNs as a function of accretion luminosity, obscuring column density, and redshift (see Table 5). An important feature of our analysis is the non-parametric method developed, which does not require any assumptions on the shape of the LF or its evolution, allowing the data to drive our results. Furthermore, we take into account all sources of uncertainties, allowing us to robustly constrain the evolution of unobscured and obscured AGNs, including the most heavily obscured Compton thick.

We find that obscured AGNs with $N_{\text{H}} > 10^{22} \text{ cm}^{-2}$ account for 77 $_{-5}^{+4}$ % of the number density and 74 $_{-5}^{+4}$ % of the luminosity density of the accretion SMBH population averaged over cosmic

time. Compton-thick objects, with $N_{\text{H}} > 10^{24} \text{ cm}^{-2}$, account for approximately half the number and luminosity density of the obscured population, and 38 $_{-7}^{+8}$ % of the total.

There is evidence that the space density of obscured, Compton-thin AGNs evolves stronger than the unobscured or Compton-thick AGNs. This is connected to the luminosity-dependent fraction of obscured AGNs. At higher luminosities, fewer AGNs are obscured. However, at higher redshift, this effect sets on at significantly higher luminosities. In the luminosity range used in this study, the fraction of obscured AGNs increases from 75 $_{-4}^{+4}$ % to 83 $_{-3}^{+3}$ % at $z > 2.25$. This is due to only a small luminosity range around $L_* \approx 10^{44} \text{ erg s}^{-1}$, which changes its obscured fraction. In particular, the space density evolves fastest around $N_{\text{H}} \approx 10^{23} \text{ cm}^{-2}$. In contrast the fraction of Compton-thick AGNs relative to the total population is consistent with being constant at $\approx 35\%$ independent of redshift and accretion luminosity. The contribution to the luminosity density by Compton-thick AGNs is 40 $_{-6}^{+6}$ %. The robust determination of a large fraction of Compton-thick AGNs consolidates AGN X-ray surveys and studies of the cosmic XRB. It also implies a physically extended nuclear obscurer.

We also find that the fraction of moderately obscured AGNs decreases not only to high but also to low luminosities, in qualitative agreement with findings of local surveys. We have discussed the observed trends of the obscured fraction with L_X and z in the context of models that either assign obscuration to the torus or galaxy–AGN co-evolution effects, with varying luminosity-dependent effects. Both classes of models can qualitatively explain our results but require that SMBHs evolve in a downsizing manner, i.e., larger black holes form earlier in the universe.

J.B. acknowledges financial support through a Max Planck society stipend. We thank the builders and operators of *Chandra X-ray Observatory* and *XMM-Newton*. This research has made use of software provided by the Chandra X-ray Center (CXC) in the application package CIAO and Sherpa. Additionally, the CosmoPy (<http://roban.github.com/CosmoPy/>), BXA (<https://johannesbuchner.github.io/BXA/>), and pymulti-nest (<https://johannesbuchner.github.io/PyMultiNest/>) software packages were used.

A.G.E. acknowledges the THALES project 383549, which is jointly funded by the European Union and the Greek Government in the framework of the program “Education and Lifelong Learning.”

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.

APPENDIX A

LUMINOSITY FUNCTION ANALYSIS

We here go through a derivation of the statistical footing of analyzing population demographics by reviewing and combining the works of Loredó (2004) and Kelly et al. (2008). Their main difference is whether the binomial distribution or its approximation, the Poisson distribution, is used.

We intend to estimate the number density per comoving volume (Mpc^3) of AGNs, as a function of various properties,

specifically X-ray luminosity, redshift, and obscuring column density. This extended LF will describe the evolution of the X-ray population and all sub-populations (e.g., Compton-thick AGNs). The difficulty is that each of these properties influences our ability to detect objects with such properties. Let us assume we can describe the probability to detect a source with properties \mathcal{C} as $p(\mathcal{D}, \mathcal{C})$.

After we analyzed each object in detail, we can bin our sample in such small bins that only one item at most can be in each bin \mathcal{C}_i . Then just applying the binomial/Poisson distribution, assuming that in k bins a source is detected (and no detections in the other $n-k$ bins), yields

$$B(k; n, p) = \underbrace{\binom{n}{k} \times \prod_{i=1}^k p(\mathcal{D}, d_i, \mathcal{C}_i)}_{\square} \times \underbrace{\prod_{i=k}^n p(\bar{\mathcal{D}}, d_i, \mathcal{C}_i)}_{*}, \quad (\text{A1})$$

$$P(k; n, p) = \underbrace{\frac{1}{k!} \times \prod_{i=1}^k n \times p(\mathcal{D}, d_i, \mathcal{C}_i)}_{\square} \times \underbrace{\exp \left\{ - \sum_{i=1}^n p(\mathcal{D}, d_i, \mathcal{C}_i) \right\}}_{*}. \quad (\text{A2})$$

The first equation gives the likelihood based on the Binomial distributions, and here $p(\bar{\mathcal{D}}, d_i, \mathcal{C}_i)$ denotes the probability of not detecting a source with properties \mathcal{C}_i and having obtained the observed data d_i . In the Poisson formula, only $p(\mathcal{D}, d_i, \mathcal{C}_i)$ occurs, which instead denotes the probability of detection.

We notice that the term marked with a star in Equation (A2) remains the same, regardless of the number of detections k in the sample (only dependent on n). The sum in this term therefore has to be always the same, independent of the specific sample d_i . In the Poisson formalism, the exponent of this term has a specific meaning—it refers to the expectation value of the sample (an a priori assumption). We can thus remove the dependency of the data and treat it as an a priori probability. Mathematically, we integrate over all possible data d_i :

$$P(k; n, p) = \exp \left\{ - \sum_{i=1}^n p(\mathcal{D}_i, \mathcal{C}_i) \right\} \times \prod_{i=1}^k (n \times p(\mathcal{D}_i, d_i, \mathcal{C}_i)). \quad (\text{A3})$$

We use $\sum_{i=1}^n p(\mathcal{D}_i, d_i, \mathcal{C}_i) = \sum_{i=1}^n p(\mathcal{D}_i, \mathcal{C}_i)$. By analogy, for the binomial distribution, we use $\int p(\bar{\mathcal{D}}_i, d_i, \mathcal{C}_i) dd_i = p(\bar{\mathcal{D}}_i, \mathcal{C}_i)$. Mathematically, we again integrate over all possible data d_i , but for all non-detected sources (see Kelly et al. 2008):

$$B(k; n, p) = \binom{n}{k} \times \prod_{i=k}^n p(\bar{\mathcal{D}}_i, \mathcal{C}_i) \times \prod_{i=1}^k p(\mathcal{D}_i, d_i, \mathcal{C}_i) \quad (\text{A4})$$

$$= \binom{n}{k} \times \exp \left\{ \sum_{i=k}^n \ln p(\bar{\mathcal{D}}_i, \mathcal{C}_i) + \sum_{i=1}^k \ln p(\mathcal{D}_i, d_i, \mathcal{C}_i) \right\}. \quad (\text{A5})$$

Now we replace our n discrete bins by a continuum. The probability $p(\mathcal{D}_i, d_i, \mathcal{C}_i)$ can then be non-zero over a range of the parameter space. After all, we do not know the true parameter \mathcal{C} . Thus, we replace $\sum_i p(\mathcal{D}_i, d_i, \mathcal{C}_i)$ by $\int_{\theta} p(\mathcal{D}_i, d_i | \mathcal{C}) d\mathcal{C}$:

$$P(k; n, p) = \frac{1}{k!} \times \exp \left\{ - \int p(\mathcal{D}, \mathcal{C}) d\mathcal{C} + \sum_{i=1}^k \ln \int p(\mathcal{D}_i, d_i, \mathcal{C}) d\mathcal{C} \right\}, \quad (\text{A6})$$

$$B(k; n, p) = \binom{n}{k} \times \exp \left\{ (n-k) \cdot \ln \int p(\bar{\mathcal{D}}, \mathcal{C}) d\mathcal{C} + \sum_{i=1}^k \ln \int_{\theta} p(\mathcal{D}_i, d_i, \mathcal{C}) d\mathcal{C} \right\}. \quad (\text{A7})$$

We will now expand $p(\mathcal{D}, \mathcal{C})$, and $p(\mathcal{D}_i, d_i, \mathcal{C})$ using conditional probabilities and discuss the meaning of the occurring terms:

$$p(\mathcal{D}, \mathcal{C}) = p(\mathcal{C}) \cdot p(\mathcal{D} | \mathcal{C}), \quad (\text{A8})$$

$$p(\mathcal{D}, d_i, \mathcal{C}) = p(\mathcal{C}) \cdot p(d_i | \mathcal{C}) \cdot p(\mathcal{D} | d_i, \mathcal{C}). \quad (\text{A9})$$

Here, $p(\mathcal{C})$ is the probability of finding an object with characteristics \mathcal{C} ($\{\log L, z, \log N_{\text{H}}\}$ here). This is essentially the LF we are interested in. The term $p(\mathcal{D} | \mathcal{C})$ in Equation (A8) denotes the sensitivity to such objects.

In Equation (A9), the second term, $p(d_i | \mathcal{C})$, is related to the spectral analysis. It denotes the likelihood that this data was generated from a source with, e.g., luminosity L and column density N_{H} . Finally, we have to consider the probability of detecting this object, given the data and characteristics \mathcal{C} . This is almost surely one, as having data associated with an object implies having detected it. However, there is a subtlety here that has been overlooked so far. Some astronomers use the sensitivity $p(\mathcal{D} | \mathcal{C})$ —specifically the area curve of where the object was detected—here, in place of $p(\mathcal{D} | d_i, \mathcal{C}) = 1$. But Loredo (2004) makes a strong point arguing that the sensitivity function must not be used here. Why is it done in practice anyway?

The X-ray source detection algorithms typically use an area of a certain, relatively small size, which contains 70% of the energy contributed by the source. The background for this region is estimated, and the probability for the background to produce the observed number of counts computed (no-source probability). If a threshold is exceeded, the spectrum is extracted. In this step a much larger region is used, containing more background counts. While in the small detection area it was not possible for the background to produce the observed counts, for faint sources it may be possible for the background to produce the observed counts in the spectrum extraction region. This is because the background contribution grows linearly with

area, but the source contribution almost stays the same. The likelihood of our analysis thus does allow having no flux from a faint source, i.e., luminosity zero. This contradiction to the detection probability stems from the fact that we have not used the information that the counts are *concentrated* in the detection region. In this case, $p(\mathcal{D}|d_i, C) \neq 1$.

Specifically, we could write it as $p(\mathcal{D}|d_i, C) = p(> k|d_i, C)$, the probability of having more than k counts in the detection region out of the extracted spectrum counts d_i , where k denotes the number necessary for detection at the current position. An approximation to this number is $p(> k|C)$, the probability to produce the number of counts required for a detection at the current position. A further approximation is $\int p(> k|C) dA/A = p(\mathcal{D}|C)$, which is the area-average sensitivity curve. Given this, we understand now why some works use the area curve in the data term, even though it should not be done: it is an attempt at fixing a loss of information introduced when the detection and data analysis processes differ.

We now put all the information together and arrive at the relevant likelihoods:

$$B(k; n, p) = \binom{n}{k} \times \left(\int p(\bar{\mathcal{D}}|C) \cdot p(C) dC \right)^{n-k} \times \prod_{i=1}^k \int p(C) \cdot p(d_i|C) \cdot p(\mathcal{D}|d_i, C) dC, \quad (\text{A10})$$

$$P(k; n, p) = \frac{1}{k!} \times \exp \left\{ - \int p(C) \cdot p(\mathcal{D}|C) dC \right\} \times \prod_{i=1}^k \int p(C) \cdot p(d_i|C) \cdot p(\mathcal{D}|d_i, C) dC. \quad (\text{A11})$$

Comparing Equation (9) in Loredo (2004) with Equation (A11) above, we have obtained the same result with the notation key $p(d_i|\theta) = l_i$, $p(\mathcal{D}|d_i, C) = 1$, $C = m$, $n = \omega\delta m$, and $p(\mathcal{D}|C) = \eta(m)$. The derivation is similar; however, we started from first principles and continued in a sound and complete way for both the Poisson and the binomial distribution.

Comparing Equation (5) in Kelly et al. (2008) with Equation (A10) above, we have obtained the same result with the notation key $C = (L_j, z_j)$, $p(\bar{\mathcal{D}}|C) = p(I_j = 0|L_j, z_j)$, $p(C) = p(L_j, z_j|\theta)$, and $p(\mathcal{D}|d_i, C) = 1$ as well as $p(d_i|C) = 1$, since Kelly et al. (2008) neglects any measurement uncertainties in the derivation, assuming that luminosity and redshift can be determined perfectly.

Finally, we summarize the logarithms of the likelihoods, neglecting constants:

$$\ln \mathcal{L}_B = (n - k) \times \ln \int p(\bar{\mathcal{D}}|C) \cdot p(C) dC + \sum_{i=1}^k \ln \int p(C) \cdot p(d_i|C) \cdot p(\mathcal{D}|d_i, C) dC, \quad (\text{A12})$$

$$\ln \mathcal{L}_P = - \int p(C) \cdot p(\mathcal{D}|C) dC + \sum_{i=1}^k \ln \int p(C) \cdot p(d_i|C) \cdot p(\mathcal{D}|d_i, C) dC. \quad (\text{A13})$$

The general-purpose framework presented here is advocated by Kelly et al. (2008). It is a general approach for the inference of population demographics based on specific samples. The importance of this step—rather than analyzing trends based on sample statistics—and incorporating selection biases, cannot be overstated.

In this work, we restrict ourselves to the Poisson likelihood. This allows us to use the space density directly (rather than separating sampled volume and probability). For $p(C)$, we insert the LF model $\phi(\log L, z, \log N_H) \times (dV/dz)$. For $p(\mathcal{D}|C)$, we insert the area curve. For $p(\mathcal{D}|d_i, C)$, we use the area curve of the specific field as an approximation (see above). For $p(d_i|C)$, we should use the likelihood of the spectral analysis. However, we did use intermediate priors in the data analysis. We thus extend $p(d_i|C) = (p(d_i|C) \times p(C))/p(C)$ where the top term is the posterior distribution computed in the spectral analysis, and the prior has to be divided away again. As we used flat priors in $\log L$, $\log N_H$, and z (subject to additional information from other wavelengths), which are the units of the integral over $C = \{\log L, z, \log N_H\}$, this division only contributes to the likelihood as a fixed offset, and is not relevant for further analysis. If different intermediate priors had been used, here they would have to be divided away. In fact, we employed a Gaussian prior on the photon index Γ . The LFs ϕ we consider, however, are in fact $\phi(\log L, z, \log N_H) \times f(\Gamma)$ where f is the Gaussian prior used, and thus we also do not require a correction here.

Regarding the practical evaluation of the likelihood function, we use posterior samples from our spectral analysis as just described. For the first integral, which describes the expected number of detected sources, we employ a fixed grid, as these are fast to compute but also ensure a consistent estimate between likelihood evaluations. We use 40 grid points in each dimension. By pre-computation of all weights except for the LF ϕ , the likelihood function reduces to (fast) addition and multiplication operations.

APPENDIX B

ROBUST PHOTOMETRIC REDSHIFT PROBABILITY DISTRIBUTIONS

In this section we elaborate on incorporating systematic uncertainties into photometric probability distributions from SED fitting. The photometric fitting method of Salvato et al. (2011, 2009), using software from Ilbert et al. (2006, 2009) is capable of not just computing a most likely redshift point estimate, but producing redshift probability distributions. Previous studies, such as Buchner et al. (2014), used the photo- z probability distributions directly to incorporate the uncertainty of the redshift estimate. However, additional to the uncertainty due to measurement errors, the method of photometric redshift has systematic errors due to incomplete template libraries, failures in automatically extracting correct fluxes at the edge of images, blending of sources causing the flux to be convoluted, etc. The most important systematic contributions are probably incorrect associations.

On a high level, the quality of redshift point estimates for a specific sample can be described by two quantities, which are estimated from a subsample where spectroscopic redshifts are available: the outlier fraction η and the scatter σ (Salvato et al. 2011). The outlier fraction describes the failure rate or accuracy of the redshifts, i.e., how frequent the redshift estimate is off

by far

$$\eta : \text{fraction where } \left| \frac{z_{\text{phot}} - z_{\text{spec}}}{1 + z_{\text{spec}}} \right| > 0.15.$$

The scatter width σ describes how precise the redshifts are scattered around the true value. The sample can be characterized using the normalized median absolute deviation (NMAD; Hoaglin et al. 1983) as a robust estimator:

$$\sigma_{\text{NMAD}} = 1.48 \times \text{median} \left| \frac{z_{\text{phot}} - z_{\text{spec}}}{1 + z_{\text{spec}}} \right|.$$

These two quantities are interpreted that the redshift estimates z_{phot} are distributed along a normal distribution centered at the true value z_{spec} with standard deviation $\sigma_{\text{NMAD}} \cdot (1 + z_{\text{spec}})$. Additionally, a fraction η of the redshift estimates are drawn from a different distribution, which can be described as flat over a redshift range $(0, z_{\text{max}})$. Here we considered $z_{\text{max}} = 5$ for outliers. The probability distribution of the point estimator incorporating systematic errors (SYSPE) can be written as

$$\text{SYSPE}(z) := \eta \times U(0, z_{\text{max}}) + (1 - \eta) \times N(z_{\text{phot}}, \sigma).$$

This modeling in fact matches the empirical distribution very well when considering the cumulative distributions of $|z_{\text{phot}} - z_{\text{spec}}| / (1 + z_{\text{spec}})$.

In this work, we would like to similarly incorporate catastrophic outliers. These are not contained in the photometric redshift distribution (PDZ) from SED fitting. Thus, we follow an analogous modeling of the uncertainty contribution. We allow a broadening of the PDZ by convolution with a Gaussian kernel of width $\bar{\sigma}$ and add a flat probability plateau with weight $\bar{\eta}$:

$$\text{SYSPDZ}(z) := \bar{\eta} \times U(0, z_{\text{max}}) + (1 - \bar{\eta}) \times (\text{PDZ} * N(0, \bar{\sigma})).$$

Note that for the special case $\bar{\sigma} = 0$ and $\bar{\eta} = 0$, SYSPDZ becomes PDZ. The two parameters $\bar{\sigma}$ and $\bar{\eta}$ have a slightly different definition than σ and η above, but take the same roles for characterize systematic uncertainty. We fit for the parameters using the spectroscopic subsample by maximizing the likelihood

$$\mathcal{L}(\bar{\sigma}, \bar{\eta}) = \prod_i \text{SYSPDZ}_i(z_{\text{spec}, i})$$

defined as the product of the value of the modified PDZ at the true redshifts z_{spec} . PDZs that are far off the true value will demand an increase in $\bar{\sigma}$ and $\bar{\eta}$; however, raising both values diminishes the SYSPDZ value overall due to normalization of probability distributions.

We find the best-fit parameters $\bar{\sigma} = 0.024$, 0.048 , and 0.029 and $\bar{\eta} = 1.8\%$, 0.6% , and 2.3% for the CDFS, AEGIS-XD, and C-COSMOS samples, respectively. These characterizations are not directly equivalent to the point estimate based definition of σ_{NMAD} and η (see Salvato et al. 2011 and Hsu et al. 2014), but the numbers are comparable. We thus use the smoothed SYSPDZ distributions instead of the PDZs.

A simpler method for incorporating the systematic uncertainty would be to use the modeling based on the point estimator (SYSPE). However, when comparing the likelihood value, we find that SYSPE always has lower likelihood values than the kernel smoothing SYSPDZ. This shows that the PDZ contains valuable information due to, e.g., secondary solutions, which is missing in the point estimator z_{phot} , and that SYSPDZ is the most faithful representation of the state of redshift information.

APPENDIX C

```

data {
  int indices[27344, 3];
  real weights[27344];
  vector[11-1] widths0;
  vector[14-1] widths1;
  vector[7-1] widths2;
  int lengths[2046];
  int chain_indices[780, 3];
  real chain_weights[780];
}
parameters {
  real<lower=-40, upper=0> y[11, 14, 7];
}
model {
  real sigmaL;
  real sigmaZ;
  real sigmaNH;
  real dataterms[2046];
  real detectionterms[780];
  real loglike;

  sigmaL <- 0.5;
  sigmaZ <- 0.5;
  sigmaNH <- 0.75;

  /* L smoothness prior:
  2nd derivative is small */
  for (i in 3:11) {
    for (j in 1:14) {
      for (k in 1:7) {
        y[i, j, k] ~ normal((y[i-1, j, k] -
y[i-2, j, k]) * widths0[i-1] / widths0[i-2] +
y[i-1, j, k], sigmaL);
      } } }

  /* z smoothness prior:
  2nd derivative is small */
  for (i in 1:11) {
    for (j in 3:14) {
      for (k in 1:7) {
        y[i, j, k] ~ normal((y[i, j-1, k] -
y[i, j-2, k]) * widths1[j-1] / widths1[j-2] + y[i,
j-1, k], sigmaZ);
      } } }

  /* NH smoothness prior:
  1st derivative is small */
  for (i in 1:11) {
    for (j in 1:14) {
      for (k in 2:7) {
        y[i, j, k] ~ normal(y[i, j, k-1],
sigmaNH);
      } } }

  /* individual objects */
  int m;
  m <- 0;
  for (i in 1:2046) {
    int a; int b; int c;
    real v[lengths[i]];

    for (j in 1:lengths[i]) {

```



```

int l;
l <- m + j;
/* look up interpolation point and
compute weights */
a <- indices[l, 1] + 1;
b <- indices[l, 2] + 1;
c <- indices[l, 3] + 1;
v[j] <- exp(y[a, b, c]) * weights[l];
}
m <- m + lengths[i];
/* add value to likelihood */
dataterms[i] <- log(sum(v) / lengths[i]);
}
}

/* detection integral */
for (k in 1:780) {
int a; int b; int c;
real v;
/* look up interpolation point and
compute weights */
a <- chain_indices[k, 1] + 1;
b <- chain_indices[k, 2] + 1;
c <- chain_indices[k, 3] + 1;
v <- chain_weights[k] * exp(y[a, b, c]);
detectionterms[k] <- v;
}
loglike <- sum(dataterms) +
sum(detectionterms);
increment_log_prob(loglike);
}

```

Algorithm 1 Stan code for estimating the field. The constant-slope prior implementation is shown. We prepared the integrals to be computed using importance sampling chain points evaluated the respective field bin value (indices) and multiplied by the weights.

REFERENCES

- Aihara, H., Allende Prieto, C., An, D., et al. 2011, *ApJS*, 193, 29
- Aird, J., Nandra, K., Laird, E. S., et al. 2010, *MNRAS*, 401, 2531
- Akylas, A., Georgakakis, A., Georgantopoulos, I., Brightman, M., & Nandra, K. 2012, *A&A*, 546, A98
- Akylas, A., & Georgantopoulos, I. 2008, *A&A*, 479, 735
- Akylas, A., & Georgantopoulos, I. 2009, *A&A*, 500, 999
- Akylas, A., Georgantopoulos, I., Georgakakis, A., Kitsionas, S., & Hatziminaoglou, E. 2006, *A&A*, 459, 693
- Alexander, D. M., Bauer, F. E., Brandt, W. N., et al. 2011, *ApJ*, 738, 44
- Alexander, D. M., Stern, D., Del Moro, A., et al. 2013, *ApJ*, 773, 125
- Antonucci, R. 1993, *ARA&A*, 31, 473
- Antonucci, R. R. J., & Miller, J. S. 1985, *ApJ*, 297, 621
- Barro, G., Pérez-González, P. G., Gallego, J., et al. 2011a, *ApJS*, 193, 13
- Barro, G., Pérez-González, P. G., Gallego, J., et al. 2011b, *ApJS*, 193, 30
- Bassani, L., Molina, M., Malizia, A., et al. 2006, *ApJL*, 636, L65
- Beasley, A. J., Murowinski, R., & Tarengi, M. 2006, *Proc. SPIE*, 6267, 2
- Bolton, A. S., Schlegel, D. J., Aubourg, É., et al. 2012, *AJ*, 144, 144
- Bournaud, F., Elmegreen, B. G., & Elmegreen, D. M. 2007, *ApJ*, 670, 237
- Brightman, M., & Nandra, K. 2011a, *MNRAS*, 413, 1206
- Brightman, M., & Nandra, K. 2011b, *MNRAS*, 414, 3084
- Brightman, M., Nandra, K., Salvato, M., et al. 2014, *MNRAS*, 443, 1999
- Brightman, M., & Ueda, Y. 2012, *MNRAS*, 423, 702
- Broos, P. S., Townsley, L. K., Feigelson, E. D., et al. 2010, *ApJ*, 714, 1582
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, 564, A125
- Budavári, T., & Szalay, A. S. 2008, *ApJ*, 679, 301
- Burlon, D., Ajello, M., Greiner, J., et al. 2011, *ApJ*, 728, 58
- Capak, P., Aussel, H., Ajiki, M., et al. 2007, *ApJS*, 172, 99
- Cappi, M., Panessa, F., Bassani, L., et al. 2006, *A&A*, 446, 459
- Cardamone, C. N., van Dokkum, P. G., Urry, C. M., et al. 2010, *ApJS*, 189, 270
- Ciotti, L., & Ostriker, J. P. 1997, *ApJL*, 487, L105
- Ciotti, L., & Ostriker, J. P. 2001, *ApJ*, 551, 131
- Civano, F., Brusa, M., Comastri, A., et al. 2011, *ApJ*, 741, 91
- Civano, F., Elvis, M., Brusa, M., et al. 2012, *ApJS*, 201, 30
- Coil, A. L., Georgakakis, A., Newman, J. A., et al. 2009, *ApJ*, 701, 1484
- Comastri, A., Ranalli, P., Iwasawa, K., et al. 2011, *A&A*, 526, L9
- Cooper, M. C., Aird, J. A., Coil, A. L., et al. 2011, *ApJS*, 193, 14
- Cooper, M. C., Griffith, R. L., Newman, J. A., et al. 2012, *MNRAS*, 419, 3018
- Davis, M., Guhathakurta, P., Konidaris, N. P., et al. 2007, *ApJL*, 660, L1
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, 145, 10
- Dickey, J. M., & Lockman, F. J. 1990, *ARA&A*, 28, 215
- Donley, J. L., Rieke, G. H., Alexander, D. M., Egami, E., & Pérez-González, P. G. 2010, *ApJ*, 719, 1393
- Ebrero, J., Carrera, F. J., Page, M. J., et al. 2009, *A&A*, 493, 55
- Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, *AJ*, 142, 72
- Elitzur, M. 2008, *MmSAI*, 79, 1124
- Elitzur, M., & Ho, L. C. 2009, *ApJL*, 701, L91
- Elitzur, M., & Shlosman, I. 2006, *ApJL*, 648, L101
- Elvis, M., Civano, F., Vignali, C., et al. 2009, *ApJS*, 184, 158
- Enoki, M., Ishiyama, T., Kobayashi, M. A. R., & Nagashima, M. 2014, *ApJ*, 794, 69
- Fabian, A. C. 1999, *MNRAS*, 308, L39
- Fabian, A. C., & Iwasawa, K. 1999, *MNRAS*, 303, L34
- Fanidakis, N., Baugh, C. M., Benson, A. J., et al. 2012, *MNRAS*, 419, 2797
- Ferrarese, L., & Ford, H. 2005, *SSRv*, 116, 523
- Ferrarese, L., & Merritt, D. 2000, *ApJL*, 539, L9
- Fiore, F., Grazian, A., Santini, P., et al. 2008, *ApJ*, 672, 94
- Fiore, F., Puccetti, S., Brusa, M., et al. 2009, *ApJ*, 693, 447
- Gebhardt, K., Bender, R., Bower, G., et al. 2000, *ApJL*, 539, L13
- Georgakakis, A., & Nandra, K. 2011, *MNRAS*, 414, 992
- Georgakakis, A., Nandra, K., Laird, E. S., Aird, J., & Trichas, M. 2008, *MNRAS*, 388, 1205
- Georgakakis, A., Rowan-Robinson, M., Nandra, K., et al. 2010, *MNRAS*, 406, 420
- Georgantopoulos, I., Comastri, A., Vignali, C., et al. 2013, arXiv:1303.5556
- Gilli, R., Comastri, A., & Hasinger, G. 2007, *A&A*, 463, 79
- Gilli, R., Vignali, C., Mignoli, M., et al. 2010, *A&A*, 519, A92
- Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, *AJ*, 131, 2332
- Guo, Y., Ferguson, H. C., Giavalisco, M., et al. 2013, *ApJS*, 207, 24
- Hasinger, G. 2008, *A&A*, 490, 905
- Hasinger, G., Miyaji, T., & Schmidt, M. 2005, *A&A*, 441, 417
- Hirschmann, M., Naab, T., Somerville, R. S., Burkert, A., & Oser, L. 2012, *MNRAS*, 419, 3200
- Hoaglin, D. C., Mosteller, F., & Tukey, J. W. 1983, *Understanding Robust and Exploratory Data Analysis* (New York: Wiley)
- Hoffman, M. D., & Gelman, A. 2011, arXiv:1111.4246
- Hönig, S. F., & Beckert, T. 2007, *MNRAS*, 380, 1172
- Hopkins, P. F., Hayward, C. C., Narayanan, D., & Hernquist, L. 2012, *MNRAS*, 420, 320
- Hopkins, P. F., Hernquist, L., Cox, T. J., et al. 2005, *ApJ*, 630, 716
- Hopkins, P. F., Hernquist, L., Cox, T. J., et al. 2006a, *ApJS*, 163, 1
- Hopkins, P. F., Somerville, R. S., Hernquist, L., et al. 2006b, *ApJ*, 652, 864
- Hsieh, B.-C., Wang, W.-H., Hsieh, C.-C., et al. 2012, *ApJS*, 203, 23
- Hsu, L.-T., Salvato, M., Nandra, K., et al. 2014, *ApJ*, 796, 60
- Ilbert, O., Arnouts, S., McCracken, H. J., et al. 2006, *A&A*, 457, 841
- Ilbert, O., Capak, P., Salvato, M., et al. 2009, *ApJ*, 690, 1236
- Iwasawa, K., Gilli, R., Vignali, C., et al. 2012, *A&A*, 546, A84
- Jia, J., Ptak, A., Heckman, T., & Zakamska, N. L. 2013, *ApJ*, 777, 27
- Kalberla, P. M. W., Burton, W. B., Hartmann, D., et al. 2005, *A&A*, 440, 775
- Kalfountzou, E., Civano, F., Elvis, M., Trichas, M., & Green, P. 2014, *MNRAS*, 445, 1430
- Kelly, B. C., Fan, X., & Vestergaard, M. 2008, *ApJ*, 682, 874
- King, A. 2003, *ApJL*, 596, L27
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, 51, 511
- Krolik, J. H., & Begelman, M. C. 1988, *ApJ*, 329, 702
- La Franca, F., Fiore, F., Comastri, A., et al. 2005, *ApJ*, 635, 864
- Laird, E. S., Nandra, K., Georgakakis, A., et al. 2009, *ApJS*, 180, 102
- Lanzuisi, G., Ranalli, P., Georgantopoulos, I., et al. 2015, *A&A*, 573, 137
- Lawrence, A. 1991, *MNRAS*, 252, 586
- Lilly, S. J., Le Brun, V., Maier, C., et al. 2009, *ApJS*, 184, 218
- Loredo, T. J. 2004, in *AIP Conf. Proc. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. R. Fischer, R. Preuss, & U. V. Toussaint (Melville, NY: AIP), 195
- Lusso, E., Hennawi, J. F., Comastri, A., et al. 2013, *ApJ*, 777, 86
- Magorrian, J., Tremaine, S., Richstone, D., et al. 1998, *AJ*, 115, 2285
- Maiolino, R., & Rieke, G. H. 1995, *ApJ*, 454, 95
- Malizia, A., Bassani, L., Bazzano, A., et al. 2012, *MNRAS*, 426, 1750

- Marshall, H. L., Tananbaum, H., Avni, Y., & Zamorani, G. 1983, *ApJ*, **269**, 35
- Matt, G. 2000, *A&A*, **355**, L31
- McCracken, H. J., Capak, P., Salvato, M., et al. 2010, *ApJ*, **708**, 202
- Merloni, A. 2004, *MNRAS*, **353**, 1035
- Merloni, A., & Heinz, S. 2008, *MNRAS*, **388**, 1011
- Merloni, A., & Heinz, S. 2013, in *Planets, Stars and Stellar Systems, Vol. 6, Evolution of Active Galactic Nuclei*, ed. T. D. Oswalt & W. C. Keel (Dordrecht: Springer), 503
- Mignoli, M., Vignali, C., Gilli, R., et al. 2013, *A&A*, **556**, A29
- Nandra, K., Barret, D., Barcons, X., et al. 2013, arXiv:1306.2307
- Nandra, K., Laird, E. S., Adelberger, K., et al. 2005, *MNRAS*, **356**, 568
- Nandra, K., Laird, E. S., Aird, J. A., et al. 2015, *ApJ*, submitted
- Nandra, K., O'Neill, P. M., George, I. M., & Reeves, J. N. 2007, *MNRAS*, **382**, 194
- Nandra, K., & Pounds, K. A. 1994, *MNRAS*, **268**, 405
- Nenkova, M., Sirocky, M. M., Nikutta, R., Ivezić, Ž., & Elitzur, M. 2008, *ApJ*, **685**, 160
- Newman, A. B., Ellis, R. S., Bundy, K., & Treu, T. 2012, *ApJ*, **746**, 162
- Norman, C., Hasinger, G., Giacconi, R., et al. 2002, *ApJ*, **571**, 218
- Panessa, F., Bassani, L., Cappi, M., et al. 2006, *A&A*, **455**, 173
- Pérez-González, P. G., Rieke, G. H., Villar, V., et al. 2008, *ApJ*, **675**, 234
- Rangel, C., Nandra, K., Laird, E. S., & Orange, P. 2013, *MNRAS*, **428**, 3089
- Richstone, D., Ajhar, E. A., Bender, R., et al. 1998, *Natur*, **395**, A14
- Risaliti, G., Maiolino, R., & Salvati, M. 1999, *ApJ*, **522**, 157
- Salvato, M., Hasinger, G., Ilbert, O., et al. 2009, *ApJ*, **690**, 1250
- Salvato, M., Ilbert, O., Hasinger, G., et al. 2011, *ApJ*, **742**, 61
- Sanders, D. B., Salvato, M., Aussel, H., et al. 2007, *ApJS*, **172**, 86
- Santini, P., Rosario, D. J., Shao, L., et al. 2012, *A&A*, **540**, A109
- Sazonov, S., Revnivtsev, M., Burenin, R., et al. 2008, *A&A*, **487**, 509
- Scoville, N., Aussel, H., Brusa, M., et al. 2007, *ApJS*, **172**, 1
- Shi, Y., Helou, G., & Armus, L. 2013, *ApJ*, **777**, 6
- Silk, J., & Rees, M. J. 1998, *A&A*, **331**, L1
- Silverman, J. D., Green, P. J., Barkhouse, W. A., et al. 2008, *ApJ*, **679**, 118
- Simpson, C. 2005, *MNRAS*, **360**, 565
- Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *AJ*, **146**, 32
- Soltan, A. 1982, *MNRAS*, **200**, 115
- Somerville, R. S., Hopkins, P. F., Cox, T. J., Robertson, B. E., & Hernquist, L. 2008, *MNRAS*, **391**, 481
- Stalin, C. S., Petitjean, P., Srianand, R., et al. 2010, *MNRAS*, **401**, 294
- Stan Development Team. 2014, *Stan: A C++ Library for Probability and Sampling, Version 2.2*
- Stark, A. A., Gammie, C. F., Wilson, R. W., et al. 1992, *ApJS*, **79**, 77
- Strüder, L., Briel, U., Dennerl, K., et al. 2001, *A&A*, **365**, L18
- Sutherland, W., & Saunders, W. 1992, *MNRAS*, **259**, 413
- Tacconi, L. J., Neri, R., Genzel, R., et al. 2013, *ApJ*, **768**, 74
- Toba, Y., Oyabu, S., Matsuhara, H., et al. 2014, *ApJ*, **788**, 45
- Tozzi, P., Gilli, R., Mainieri, V., et al. 2006, *A&A*, **451**, 457
- Treister, E., Urry, C. M., Chatzichristou, E., et al. 2004, *ApJ*, **616**, 123
- Treister, E., Urry, C. M., & Virani, S. 2009, *ApJ*, **696**, 110
- Tremaine, S., Gebhardt, K., Bender, R., et al. 2002, *ApJ*, **574**, 740
- Trump, J. R., Impey, C. D., Elvis, M., et al. 2009, *ApJ*, **696**, 1195
- Turner, M. J. L., Abbey, A., Arnaud, M., et al. 2001, *A&A*, **365**, L27
- 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
- Vasudevan, R. V., Brandt, W. N., Mushotzky, R. F., et al. 2013, *ApJ*, **763**, 111
- Vignali, C., Alexander, D. M., Gilli, R., & Pozzi, F. 2010, *MNRAS*, **404**, 48
- Vignali, C., Mignoli, M., Gilli, R., et al. 2014, *A&A*, **571**, A34
- Vito, F., Gilli, R., Vignali, C., et al. 2014, *MNRAS*, **445**, 3557
- Wada, K. 2012, *ApJ*, **758**, 66
- Xue, Y. Q., Luo, B., Brandt, W. N., et al. 2011, *ApJS*, **195**, 10
- York, D. G., Adelman, J., Anderson, J. E., Jr., et al. 2000, *AJ*, **120**, 1579