



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

X-rays across the galaxy population I

tracing the main sequence of star formation

Citation for published version:

Aird, J, Coil, AL & Georgakakis, A 2017, 'X-rays across the galaxy population I: tracing the main sequence of star formation', *Monthly Notices of the Royal Astronomical Society*, vol. 465, no. 3, pp. 3390-3415.
<https://doi.org/10.1093/mnras/stw2932>

Digital Object Identifier (DOI):

[10.1093/mnras/stw2932](https://doi.org/10.1093/mnras/stw2932)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Monthly Notices of the Royal Astronomical Society

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



X-rays across the galaxy population I: tracing the main sequence of star formation

J. Aird^{1*}, A. L. Coil² and A. Georgakakis^{3,4}

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

²*Center for Astrophysics and Space Sciences (CASS), Department of Physics, University of California, San Diego, CA 92093, USA*

³*Max Planck Institute für Extraterrestrische Physik, Giessenbachstrasse, 85748 Garching, Germany*

⁴*IAASARS, National Observatory of Athens, GR-15236 Penteli, Greece*

Accepted 2016 Nov 10. Received 2016 Oct 19; in original form 2016 July 01.

ABSTRACT

We use deep *Chandra* imaging to measure the distribution of X-ray luminosities (L_X) for samples of star-forming galaxies as a function of stellar mass and redshift, using a Bayesian method to push below the nominal X-ray detection limits. Our luminosity distributions all show narrow peaks at $L_X \lesssim 10^{42}$ erg s⁻¹ that we associate with star formation, as opposed to AGN that are traced by a broad tail to higher L_X . Tracking the luminosity of these peaks as a function of stellar mass reveals an “X-ray main sequence” with a constant slope $\approx 0.63 \pm 0.03$ over $8.5 \lesssim \log \mathcal{M}_*/\mathcal{M}_\odot \lesssim 11.5$ and $0.1 \lesssim z \lesssim 4$, with a normalization that increases with redshift as $(1+z)^{3.79 \pm 0.12}$. We also compare the peak X-ray luminosities with UV-to-IR tracers of star formation rates (SFRs) to calibrate the scaling between L_X and SFR. We find that $L_X \propto \text{SFR}^{0.83} \times (1+z)^{1.3}$, where the redshift evolution and non-linearity likely reflect changes in high-mass X-ray binary populations of star-forming galaxies. Using galaxies with a broader range of SFR, we also constrain a stellar-mass-dependent contribution to L_X , likely related to low-mass X-ray binaries. Using this calibration, we convert our X-ray main sequence to SFRs and measure a star-forming main sequence with a constant slope $\approx 0.76 \pm 0.06$ and a normalization that evolves with redshift as $(1+z)^{2.95 \pm 0.33}$. Based on the X-ray emission, there is no evidence for a break in the main sequence at high stellar masses, although we cannot rule out a turnover given the uncertainties in the scaling of L_X to SFR.

Key words: galaxies: evolution – galaxies: star formation – X-rays: galaxies

1 INTRODUCTION

It is now well established that most galaxies can be divided into two, fairly distinct populations: *quiescent* galaxies, which appear red due to their passively evolving stellar populations and generally have elliptical morphologies; and *star-forming* galaxies, which have blue colours due to the ongoing formation of new stars and tend to have “late-type” morphologies with significant disk components (e.g. Bell et al. 2004; Faber et al. 2007; Blanton & Moustakas 2009). Over cosmic time, the total combined mass of quiescent galaxies is found to increase (e.g. Brown et al. 2007; Moustakas et al. 2013), indicating that some process is shutting down star formation and transforming massive star-forming galaxies into quiescent systems.

Star-forming galaxies, on the other hand, are found to follow a relatively tight correlation between the current star-formation rate (SFR) and stellar mass (\mathcal{M}_*), referred to as the “main sequence of star formation” (Noeske et al. 2007). The average SFR increases with \mathcal{M}_* , following approximately a power law of the

form $\text{SFR} \propto \mathcal{M}_*^m$, where m denotes the logarithmic “slope” of the main sequence and is generally found in the range $m \sim 0.6 - 1$ (e.g. Elbaz et al. 2007; Karim et al. 2011; Speagle et al. 2014). The tightness of this correlation, with an intrinsic scatter of $\sim 0.3 - 0.5$ dex (e.g. Rodighiero et al. 2011; Ilbert et al. 2015; Schreiber et al. 2015), indicates that the majority of star-forming galaxies build up their stellar populations in a smooth and continuous manner. The main sequence evolves strongly with redshift, shifting to lower SFRs (at a fixed \mathcal{M}_*) as redshift decreases (e.g. Noeske et al. 2007; Daddi et al. 2007; Elbaz et al. 2011), indicating a uniform reduction in the rate at which star formation is fuelled as cosmic time progresses. This drop in the normalization of the main sequence drives the rapid decline in the overall density of star formation in the Universe between $z \sim 2$ and the present day.

Given the fundamental constraints that the main sequence places on our understanding of galaxy evolution, there have been substantial efforts to measure the slope and normalization over a wide range of stellar masses and redshifts (e.g. Karim et al. 2011; Whitaker et al. 2012; Schreiber et al. 2015, see Speagle et al. 2014 for an overview). A roughly linear relation ($m \approx 1$) is found for low stellar masses ($\mathcal{M}_* \lesssim 10^{10} \mathcal{M}_\odot$) at all redshifts (e.g.

* jaird@ast.cam.ac.uk

Whitaker et al. 2014; Schreiber et al. 2015; Tomczak et al. 2016). At higher stellar masses ($\mathcal{M}_* \gtrsim 10^{10} \mathcal{M}_\odot$), however, the slope appears to flatten as redshift decreases (e.g. Whitaker et al. 2014; Schreiber et al. 2015; Tomczak et al. 2016). The slope of the main sequence reflects the efficiency at which new stars are formed relative to the sum of past star formation (i.e. the total stellar mass of the galaxy). A turnover at high stellar masses and late cosmic times could indicate a reduction in the star formation efficiency at the highest stellar masses or the onset of any overall quenching process.

One of the major challenges in studies of the main sequence is to measure the SFRs of galaxies with sufficient accuracy over a wide dynamic range in stellar mass and redshift (see Kennicutt & Evans 2012 and Madau & Dickinson 2014 for recent reviews of SFR estimators). The SFR can be inferred from the rest-frame UV emission, which directly traces the emission from massive stars formed within the last ~ 10 – 100 Myr (e.g. Kennicutt 1998). However, the UV emission is extremely sensitive to the effects of dust; thus large (and in many cases uncertain) corrections are required to recover the total SFR (e.g. Gordon et al. 2000; Salim et al. 2007; Hao et al. 2011). Alternatively, obscured star formation (occurring within the last ~ 100 Myr) can be traced more directly via the re-radiated emission from dust at infrared (IR) wavelengths of ~ 8 – $1000 \mu\text{m}$. With the advent of *Herschel* it has become possible to trace the peak of the far-IR emission (at rest-frame ~ 60 – $100 \mu\text{m}$) out to $z \sim 4$ (e.g. Elbaz et al. 2011), although the limited sensitivity means only highly star-forming galaxies are detected and a stacking analysis is usually required to probe the typical SFRs of main sequence galaxies (e.g. Schreiber et al. 2015). Fainter galaxies can be detected in $24 \mu\text{m}$ imaging from *Spitzer* and used to infer the dust emission at longer wavelengths, although stacking analyses are still required to reach the very lowest SFRs (e.g. Whitaker et al. 2014). While such stacking analyses are powerful, they can be affected by selection effects (when defining galaxy samples for stacking), are sensitive to background subtraction and blending effects, and provide an average that may not accurately represent the underlying distribution (see e.g. Viero et al. 2013; Schreiber et al. 2015, and references therein).

Alternatively, dust-corrected SFRs can be estimated by modelling the broadband photometry of galaxies (spanning UV, optical and IR wavelengths) with stellar population synthesis models (e.g. Wuyts et al. 2011; Conroy 2013), although this method is sensitive to the assumed models and the corrections for dust extinction. Furthermore, the same data are often used to estimate the galaxy stellar mass, which can introduce correlated uncertainties between parameters that could bias measurements of the main sequence (e.g. Reddy et al. 2012). The SFR can also be probed via nebular line emission such as $\text{H}\alpha$ (providing a prompt tracer of star formation occurring within the last ~ 3 – 10 Myr), although dust corrections must be applied and large spectroscopically observed galaxy samples are required (e.g. Moustakas et al. 2006; Shivaei et al. 2015). Radio emission also provides a probe of the SFR, although again a stacking analysis is required to probe the bulk of the star-forming galaxy population out to high redshifts (e.g. Karim et al. 2011).

X-ray emission also provides a tracer of the SFR. The overall X-ray emission from a galaxy is due to the combined emission from high-mass X-ray binaries (HMXBs) and low-mass X-ray binaries (LMXBs) within the galaxy, hot gas throughout the galaxy, and accretion onto the central supermassive black hole that powers an Active Galactic Nucleus (AGN), if present (e.g. Fabbiano 1989). HMXBs are expected to form promptly after a burst of star formation (within ~ 5 Myr) and have relatively short lives (~ 100 –

300 Myr e.g. Fragos et al. 2013). Thus, HMXBs should be a relatively direct tracer of recent star-formation and provide the dominant contribution to the X-ray luminosity in normal star-forming galaxies (i.e. with negligible current AGN activity). LMXBs, on the other hand, tend to form later (~ 100 – 300 Myr after a burst of star formation) and then gradually fade over $\sim \text{Gyr}$ time-scales. Thus, the X-ray luminosity from LMXBs will provide a delayed tracer of the SFR (e.g. Ghosh & White 2001) and may scale with the overall stellar mass of the galaxy, rather than recent star formation (e.g. Lehmer et al. 2010). Diffuse, hot, ionized gas, thought to be heated by supernovae and winds from massive stars, is also observed in star-forming galaxies, with the total luminosity scaling with SFR and contributing up to a third of the overall emission at softer X-ray energies (e.g. Mineo et al. 2012). The X-ray luminosity from an AGN, if present, will typically dominate over any other X-ray emission from a galaxy (e.g. Brandt & Alexander 2015).

A number of studies have explored the use of the X-ray luminosity as a tracer of the SFR. In the local universe, it is possible to directly detect moderate-to-high SFR galaxies, which have been used to calibrate the X-ray luminosity as a SFR tracer (e.g. David et al. 1992; Ranalli et al. 2003; Rovilos et al. 2009; Lehmer et al. 2010). The deepest *Chandra* surveys are able to detect the highest SFR galaxies (e.g. ultraluminous infrared galaxies) out to $z \sim 1$ and thus extend such studies of the X-ray properties to higher redshifts (e.g. Symeonidis et al. 2011, 2014; Mineo et al. 2014). Stacking analyses have been used to push to fainter limits and probe more typical, high-redshift galaxy populations (e.g. Nandra et al. 2002; Laird et al. 2006; Basu-Zych et al. 2007; Lehmer et al. 2016), providing crucial constraints on the relation between the physical properties of galaxies and their X-ray binary populations. However, extremely deep X-ray data are needed to directly detect star-forming galaxies, as well as reliably identify and exclude AGNs from any stacking analyses. Such studies are further complicated by the variation in the depths of X-ray survey data across different fields and within a single field. Thus, the adoption of X-ray tracers of SFRs has been somewhat limited, with most studies restricted to the small areas covered by the *Chandra* Deep Fields (e.g. Lehmer et al. 2008; Symeonidis et al. 2014; Lehmer et al. 2016).

In this paper we develop a new approach to probe the X-ray emission from large samples of star-forming galaxies, combining data from a number of *Chandra* surveys of varying depths. We determine the *distribution* of X-ray luminosities for samples of star-forming galaxies as a function of redshift and stellar mass, spanning $z \sim 0$ – 4 and $\mathcal{M}_* \sim 10^{8.5-11.5} \mathcal{M}_\odot$, using a Bayesian technique that allows us to push substantially below the nominal detection limits in a given field. Our measured distributions allow us to identify and separate two different origins for the X-ray emission: 1) galactic emission (predominantly from high-mass X-ray binaries), which produces a peak in the distribution at low luminosities that traces the average star formation rate for galaxies of a given stellar mass and redshift and is the focus of this paper; and 2) the emission from active galactic nuclei (AGN), tracing the distribution of supermassive black hole accretion activity within galaxies of a given stellar mass and redshift, which is explored in a companion paper (Aird et al. 2016, in preparation, hereafter Paper II).

In Section 2 we describe our data and define our samples. We construct large, stellar-mass limited samples of star-forming galaxies from near-infrared selected catalogues from the CANDELS/3D-HST and UltraVISTA surveys. We extract X-ray information for all galaxies in our samples from the available *Chandra* imaging. Section 3 presents our measurements of the intrinsic distributions of X-

ray luminosities within star-forming galaxies as a function of stellar mass and redshift. We identify peaks in these distributions at low luminosities that we relate to star formation processes. In Section 4 we measure the position of these peaks as a function of stellar mass and redshift, revealing the “X-ray main sequence”. In Section 5 we compare with UV-to-IR SFR estimates to constrain the relation between X-ray luminosity and SFR, allowing us to relate the X-ray main sequence to the main sequence of star formation. Section 6 discusses our findings and compares to previous work. We summarize our findings in Section 7. We adopt a flat cosmology with $\Omega_{\Lambda} = 0.7$ and $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and assume a Chabrier (2003) stellar initial mass function (IMF) throughout this paper.

2 DATA AND SAMPLE SELECTION

2.1 Near-infrared selected photometric catalogues

We construct a large sample of star-forming galaxies by combining data from four of the CANDELS survey fields (Grogin et al. 2011; Koekemoer et al. 2011) and the UltraVISTA survey of the wider COSMOS field (McCracken et al. 2012). We adopt catalogues of objects selected at near-infrared (NIR) wavelengths, which allows us to define a sample of galaxies in an homogenous manner over a wide range of redshifts and down to a well-defined stellar mass completeness limit (see Section 2.4 below for details).

The CANDELS survey is an ultra-deep NIR imaging campaign with the Wide Field Camera 3 (WFC3) on the *Hubble Space Telescope* (*HST*), covering $\sim 800 \text{ arcmin}^2$ across five premier multiwavelength survey fields, four of which are used in this study (GOODS-N, GOODS-S, AEGIS, and COSMOS). The 3D-HST survey (Brammer et al. 2012) complements the deep CANDELS imaging with low-resolution near-infrared spectroscopy, primarily using the the WFC3 G141 grism on *HST*. In this paper, we use the v4.1.5 photometric and redshift catalogues provided by the 3D-HST collaboration (Skelton et al. 2014; Momcheva et al. 2016). The photometric catalogues are based on detection in $F125W + F140W + F160W$ combined images. PSF-matched photometry is then obtained from a wide range of ground- and space-based imaging campaigns in the various fields, covering the UV ($\sim 0.4 \mu\text{m}$) to mid-infrared (MIR: $\sim 8 \mu\text{m}$) with up to 44 broad or medium band filters (full details are given by Skelton et al. 2014, hereafter S14). Deblended photometry at $24 \mu\text{m}$ from *Spitzer*/MIPS is also provided by Whitaker et al. (2014). The combined catalogue consists of $\sim 150,000$ objects over the four CANDELS/3DHST fields.

The UltraVISTA survey is a deep, wide-area near-infrared imaging survey, covering $\sim 1.6 \text{ deg}^2$ in the COSMOS field (Scoville et al. 2007) and thus complements the deeper CANDELS imaging, providing a better sampling of galaxies with lower redshifts ($z \lesssim 1$) and higher stellar masses ($\mathcal{M}_* \gtrsim 10^{10} \mathcal{M}_{\odot}$). Near-infrared imaging in the Y , J , H and K_S bands was obtained as part of the deepest component of the ESO public survey programme with the Visible and Infrared Survey Telescope for Astronomy (VISTA). In this paper, we adopt the K_S -selected catalogues provided by Muzzin et al. (2013b, hereafter M13). This catalogue provides PSF-matched photometry in 30 photometric bands, covering $0.15 - 24 \mu\text{m}$, and includes the data from *GALEX* (Martin et al. 2005), CFHT/Subaru (Capak et al. 2007), UltraVISTA (McCracken et al. 2012), and S-COSMOS (Sanders et al. 2007). We mask out any areas contaminated by bright stars, as defined by M13. The CANDELS WFC3 imaging of the COSMOS field is fully contained within the larger UltraVISTA survey area;

in this paper we adopt the deeper CANDELS/3D-HST photometric catalogues when available and thus also mask out this area from the UltraVISTA catalogues. Our UltraVISTA catalogue thus contains $\sim 230,000$ K_S -selected sources.

2.2 Chandra X-ray data

Chandra X-ray imaging of varying depths is available in our four CANDELS fields, with exposure times of ~ 160 ks in COSMOS (Elvis et al. 2009), ~ 800 ks in AEGIS (Nandra et al. 2015), ~ 2 Ms in GOODS-N (Alexander et al. 2003) and ~ 4 Ms in GOODS-S (Xue et al. 2011).¹ *Chandra* imaging of $\sim 50 - 160$ ks depth of the entire UltraVISTA field has also recently become available as part of the COSMOS-Legacy programme (Civano et al. 2016; Marchesi et al. 2016), which we exploit in this paper.

All of the *Chandra* observations were analysed using our own data reduction and source detection procedure (see Laird et al. 2009; Nandra et al. 2015, for details), providing homogenous and well-characterised data products that are required for our statistical analysis (see also Georgakakis et al. 2008, 2014, 2015; Rangel et al. 2013; Aird et al. 2008, 2010, 2015). We apply the point source detection procedure described by Laird et al. (2009) with a Poisson false probability threshold of $< 4 \times 10^{-6}$ in the soft (0.5–2 keV), hard (2–7 keV), full (0.5–7 keV) and ultrahard (4–7 keV) energy bands, combining the source lists to create a merged point source catalogue. This provides samples of 914 significant X-ray detections in the area of our four CANDELS fields and 2851 significant X-ray detections in our UltraVISTA area.

We cross-match our sample of significant X-ray detections with the NIR-selected photometric catalogues using the likelihood ratio technique (see Ciliegi et al. 2003; Brusa et al. 2007), applying a threshold that maximizes the sum of the completeness and reliability (as described by Luo et al. 2010, see also Aird et al. 2015). Using this technique, we are able to robustly associate an NIR-selected object with 831 (91 per cent) of the detected X-ray sources in the CANDELS fields and 2663 (93 per cent) of the X-ray sources in the UltraVISTA area.

We extract X-ray information for the remaining NIR-selected objects within our fields that are not already matched with a significant X-ray detection. We first remove any NIR objects from our catalogues that lie close to a significant X-ray source (but are not associated with the X-ray source according to our likelihood ratio matching), removing objects within 1.5 times the radius corresponding to the 90 per cent enclosed energy fraction (EEF) of the exposure-weighted *Chandra* PSF. This cut removes between 5 and 10 per cent of objects from our NIR catalogues in an unbiased manner (assuming the NIR catalogues are not strongly clustered at $\lesssim 5$ arcsecond scales). We then extract the total X-ray counts in the full, soft and hard energy bands from within a circular aperture at the position of every remaining NIR object. The radius of the aperture is based on the exposure-weighted PSF and corresponds to an EEF of 70 per cent. We also estimate the background counts within the same aperture from our background maps² and the effective exposure at the position of each object. Our measurements

¹ Recent *Chandra* imaging of the GOODS-S field (that will provide a total depth ~ 7 Ms, P.I. Brandt) and the ~ 600 ks *Chandra* imaging of the fifth CANDELS field (UDS, P.I. Hasinger) are not considered in this work.

² Background maps are generated by replacing counts in the region of detected sources with values from a local background region and smoothing to create an overall background map. See Georgakakis et al. (2008) for full details.

of the total and background counts are identical in our source detection procedure but are applied here to the larger catalogue of NIR-selected objects. The extracted total counts, background counts and exposures contain information for sources that fall below our direct detection threshold but may correspond to a lower significance detection. The overall sensitivity of our X-ray observations is also described by these data; i.e. information on the fraction of sources that are *not* directly detected in the X-ray images. All of this information is used by our Bayesian analysis method, described in Appendix B and implemented in Section 3 below.

2.3 Redshifts, rest-frame colours, stellar masses and star-formation rates

The 3DHST and UltraVISTA catalogues provide spectroscopic redshifts, compiled from a number of ground-based spectroscopic follow-up campaigns (see S14 and M13 and references therein for details). We also include additional spectroscopic redshifts from the compilations described by Aird et al. (2015) and Marchesi et al. (2016) and references therein, which include follow-up campaigns focused on the X-ray source populations. In addition, we include recently obtained spectroscopic redshifts from the first data releases of the MOSDEF survey (Kriek et al. 2015) and VIMOS Ultra Deep Survey (VUDS: Le Fèvre et al. 2015, Tasca et al. in preparation). If a higher resolution spectroscopic redshift is unavailable, we adopt redshifts based on the combination of the low-resolution WFC3 grism spectroscopy and the photometry in the CANDELS fields (Momcheva et al. 2016) or the low-resolution spectroscopic redshifts of large numbers of galaxies obtained with PRIMUS in the UltraVISTA field (Coil et al. 2011; Cool et al. 2013). We match the spectroscopic redshift catalogues to our NIR-selected catalogues using a $2''$ search radius, correct for any global offset, then re-match using a stricter $0.5''$ search radius. In total, we are able to assign high- or low-resolution spectroscopic redshifts to 17,692 sources in our CANDELS catalogue and 13,860 sources in the UltraVISTA catalogue, including 680 (82 per cent) and 1648 (59 per cent) of the X-ray detected sources in CANDELS and UltraVISTA, respectively.

For the sources without spectroscopic redshifts, we must resort to photometric redshift estimates (*photo-z*). For the bulk of our sources, we adopt the *photo-z* provided by S14 and M13. These *photo-z* are calculated using the *EaZY* code (Brammer et al. 2008) with a standard set of galaxy templates and achieve an accuracy of $\sigma \lesssim 0.02^3$ and outlier rates of $\sim 2 - 5$ per cent when compared with high-quality spectroscopic galaxy redshift samples. For X-ray detected sources, we instead adopt *photo-z* that are optimised for sources which may have significant AGN contributions to the optical and IR light. We adopt high-quality AGN *photo-z* from Hsu et al. (2014), Nandra et al. (2015), or Marchesi et al. (2016) for the GOODS-S, AEGIS and COSMOS/UltraVISTA fields, respectively, all of which were calculated using techniques developed by Salvato et al. (2009, 2011) and achieve accuracies of $\sigma \lesssim 0.03$ with outlier rates of ~ 5 per cent. If a high-quality AGN *photo-z* is not available from these catalogues (e.g. in the GOODS-N field or due to slight differences in the X-ray source lists), we recalculate the *photo-z* based on the 3DHST or UltraVISTA photometry using the method described by Aird et al. (2015). These *photo-z* have a lower accuracy ($\sigma \sim 0.05$) and higher outlier rate (~ 10 -15

per cent) than the work by Salvato and collaborators but are only adopted for a small fraction (~ 7 per cent) of our X-ray sources. Given our high spectroscopic completeness for X-ray sources and the high quality of the bulk of our galaxy and AGN *photo-z*, we neglect uncertainties in the *photo-z* for this work.

We calculate rest-frame colours in the standard *U*, *V* and *J* band passes for all sources in our catalogues. We use *EaZY* with the standard set of galaxy templates to interpolate from the observed photometry (see Brammer et al. 2011), fixing the redshift at the spectroscopic value (if available) or our best *photo-z* estimate. We do not make any corrections for an AGN contribution to the spectral energy distributions (SEDs); thus our estimates correspond to the *observed* rest-frame colours and will be contaminated by bright AGNs in some cases (see further discussion in Section 2.4 below).

We use the FAST code (Kriek et al. 2009) to fit the SEDs of our sources with stellar population synthesis models and estimate both stellar masses and SFRs. We apply a number of important modifications in our implementation of the FAST code to ensure we recover accurate and physically meaningful stellar population parameters across our wide redshift range. We adopt delayed- τ models of the star-formation histories but apply a redshift-dependent *minimum* age, ensuring our code adopts realistic galaxy templates at both low and high redshifts. We also use a χ^2 -weighted average over the grid of stellar population parameters assumed in FAST when estimating stellar masses and SFRs, which reduces the effects of the relatively coarse grid and provides a fairer weighting of the range of viable models, removing some clear outliers (e.g. star-forming galaxies with blue optical colours that can be assigned steeply declining star-formation histories and erroneously low SFRs). Further details are given in Appendix A below. We also account for an AGN contribution to the SED, as described in the appendix to Paper II, although at the low X-ray luminosities probed in this paper the AGN contribution to the optical-to-MIR SED is generally negligible and thus does not have an impact on the results presented here. Our final SED-based SFR estimates use an ‘‘SFR ladder’’ (e.g. Wuyts et al. 2011): when a $24\mu\text{m}$ detection is available we use UV+IR SFRs that sum the unobscured and obscured star formation seen in the UV and far-IR respectively⁴; for sources that are not detected at $24\mu\text{m}$ we instead adopt the dust-corrected estimates of SFRs from the fits to the UV-to-MIR SEDs. Changing the details of the SED modelling (such as adopting different star formation histories or neglecting our χ^2 weighting scheme) has a negligible impact on our estimates of the total galaxy stellar mass. Thus, we are able to estimate accurate stellar masses for the bulk of our sample (including the majority of our X-ray detected sources). Our final SED-based SFR estimates are more sensitive to the assumptions in the SED modelling and are thus less certain for individual sources. Nonetheless, our consistent calculations of SFRs ensure that when combining large samples of galaxies we are able to accurately recover overall trends between SFR and X-ray luminosity. Furthermore, the main results of this paper (see Section 4) do not rely on these multiwavelength measurements of the SFR.

2.4 Star-forming galaxy sample definition

Our study requires well-defined, unbiased samples of the star-forming galaxy population down to a given stellar mass limit. To

³ Photo-*z* accuracy, σ , is given in terms of the median absolute deviation, normalized relative to $(1 + z_{\text{spec}})$.

⁴ For X-ray detected sources, we apply a correction to allow for an AGN contribution in the UV and IR; see the appendix of Paper II for full details.

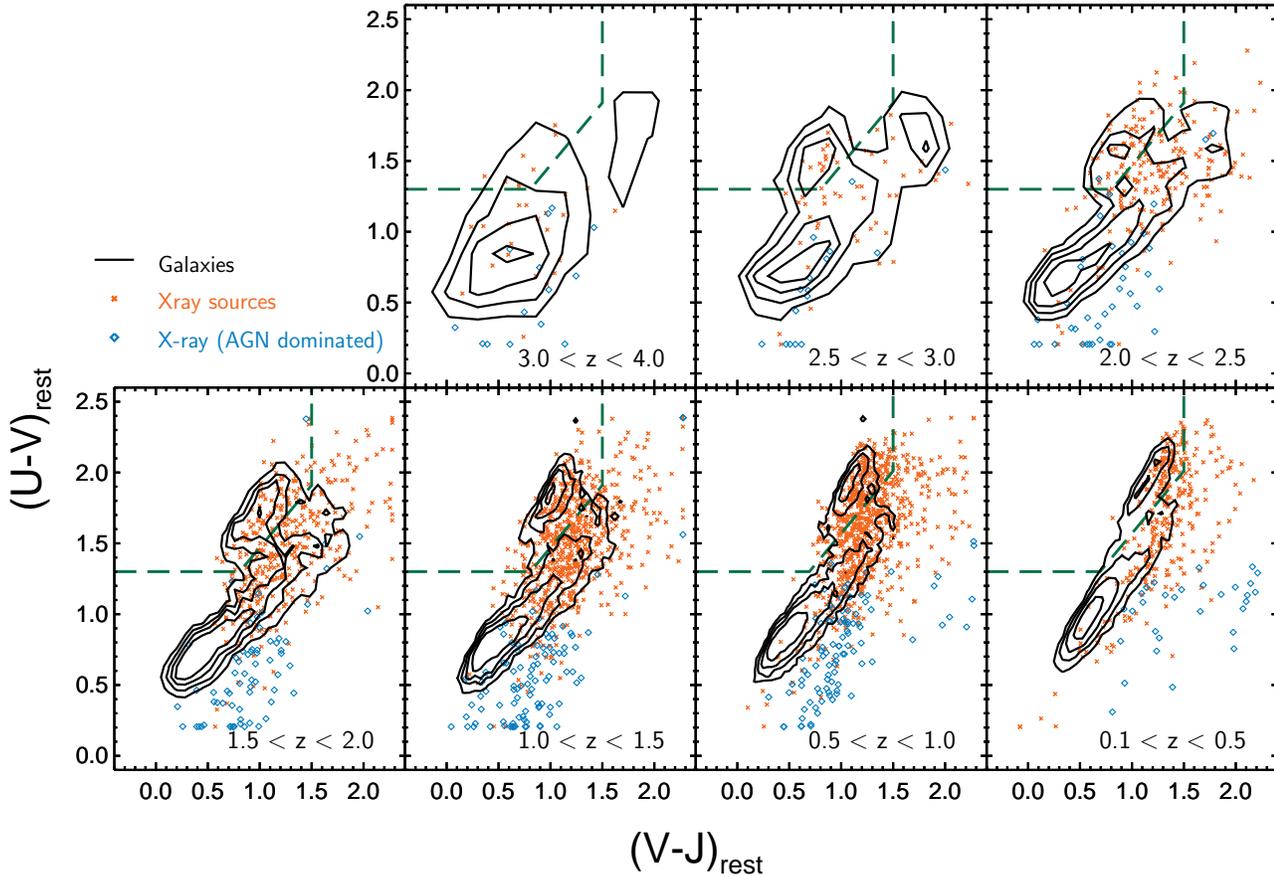


Figure 1. Rest-frame $U - V$ versus $V - J$ colours of galaxies (black contours) in our combined stellar-mass-limited CANDELS and UltraVISTA samples. A clear bimodality—seen out to at least $z \sim 2.5$ —allows us to separate quiescent and star-forming galaxies (including dust-reddened systems) using the colour-colour cuts defined by Equation 2 (green dashed lines). X-ray detected objects (orange crosses) span the galaxy population, although they are biased toward the dusty star-forming region due to the selection bias toward high stellar mass host galaxies in X-ray selected samples (Aird et al. 2012). Blue diamonds indicate X-ray sources where our two-component SED fitting indicates that the optical light is dominated by the AGN resulting in blue, contaminated colours; these sources are excluded from the sample of star-forming galaxies used in this paper.

construct such samples, we first generate magnitude-limited catalogues by applying limits of $H_{F160W} < 25.1$ and $K_S < 23.4$ to our CANDELS and UltraVISTA catalogues, respectively. These limits correspond to the 90 per cent completeness limits for point sources given by S14 and M13.

Next, we remove any sources that are spectroscopically identified as stars. We also remove objects that are photometrically flagged as stars based on the positions in the size-magnitude plane (in the S14 catalogues of the CANDELS fields) or a cut in the $u^* - J$ vs. $J - K_S$ colour-colour space (in the M13 catalogues of the UltraVISTA field).

We then apply stellar mass completeness limits. For the UltraVISTA field we adopt the empirically determined 95 per cent stellar mass completeness limits (as a function of redshift) that were calculated by Muzzin et al. (2013a) by comparing the distribution of stellar masses in the UltraVISTA catalogue to deeper NIR-selected catalogues. We shift the mass limits down by 0.04 dex as we adopt a Chabrier (2003) initial mass function in our SED fitting. We restrict our UltraVISTA sample to galaxies lying above this stellar mass completeness limit at a given redshift. For the CANDELS fields we shift the UltraVISTA completeness limits down by an additional order of magnitude (1 dex), allowing for the significantly

deeper NIR imaging. This provides a completeness limit as a function of redshift that is in agreement with the empirically determined 95 per cent completeness limits calculated by Tal et al. (2014) for the CANDELS fields at $z \lesssim 1.5$ and is more conservative at higher redshifts.

We restrict our samples to include only star-forming galaxies by applying a cut based on the rest-frame $U - V$ versus $V - J$ colour-colour diagram (hereafter the “UVJ diagram”). This diagnostic divides the galaxy population according to the well-established bimodality between actively star-forming galaxies and quiescent galaxies, with the added advantage that dust-reddened star-forming systems can be separated from the quiescent galaxy population and classified correctly (Labbé et al. 2005; Wuyts et al. 2007; Williams et al. 2009; Whitaker et al. 2011). Figure 1 shows the UVJ diagram for the combination of our stellar-mass-limited samples from CANDELS and UltraVISTA (grey contours). We observe a clear bimodality in our galaxy sample up to at least $z \sim 2.5$ that allows us to separate quiescent galaxies from the star-forming population. We adopt colour-colour cuts

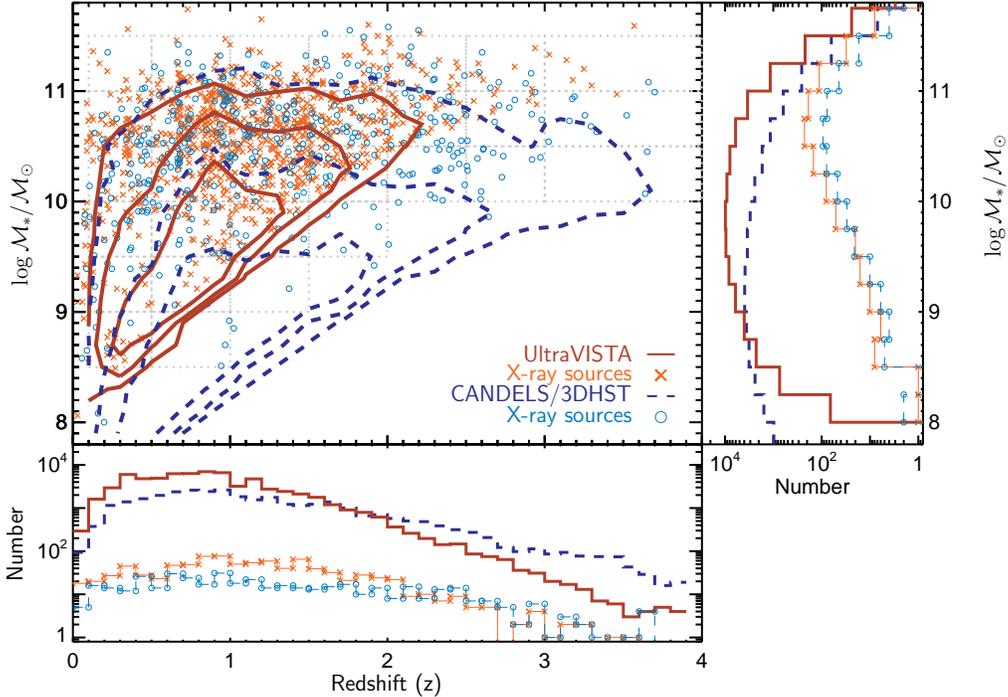


Figure 2. Distributions of stellar mass and redshift for the CANDELS/3DHST (dashed blue contours) and UltraVISTA (solid red contours) star-forming galaxy samples. Contours enclose 50, 80 and 95 per cent of the star-forming galaxies in each survey. Points (light blue circles and orange crosses for CANDELS and UltraVISTA, respectively) indicate sources with soft band (0.5–2 keV) X-ray detections. The grey dotted lines indicate our bins of stellar mass and redshift, which are required to lie wholly above the stellar mass completeness limits of the CANDELS fields.

from Muzzin et al. (2013a), where quiescent galaxies satisfy

$$\begin{aligned}
 (U - V) > 1.3, \quad (V - J) < 1.5 & \quad (\text{all redshifts}) \\
 (U - V) > 0.88 \times (V - J) + 0.69 & \quad (0.0 < z < 1.0) \\
 (U - V) > 0.88 \times (V - J) + 0.59 & \quad (1.0 < z < 4.0)
 \end{aligned} \quad (1)$$

which were adapted from the colour cuts originally proposed by Williams et al. (2009).

The orange crosses in Figure 1 indicate the colours of sources with significant X-ray detections. As previously seen by Georgakakis et al. (2014), the X-ray detections are distributed across the quiescent and star-forming galaxy populations, although there is a strong trend toward the dusty star-forming region of the UVJ diagram (redder $U - V$ and $V - J$ colours but below the quiescent galaxy selection box). This trend is due to the selection bias toward high stellar mass host galaxies in X-ray selected AGN samples (Aird et al. 2012) and the fact that higher stellar mass galaxies are known to contain higher levels of dust (e.g. Santini et al. 2014, see also Georgakakis et al. 2014, Cowley et al. 2016, Azadi et al. in preparation). The blue diamonds indicate X-ray detections where our two-component SED fitting finds that more than 50 per cent of the light at rest-frame 5000 Å is coming from the AGN component; thus, the rest-frame UVJ colours are contaminated by the AGN light, leading to the very blue colours in the UVJ diagram and making it impossible to classify the host galaxy as star-forming or quiescent. These sources could constitute at most 0.5 per cent of our total star-forming galaxy sample (assuming they all had star-forming hosts) and their X-ray emission is dominated by the AGN (rather than tracking the SFR, which is the aim of this paper). Thus, we exclude these AGN-dominated sources from the sample of star-forming galaxies used in this paper.

Our final sample consists of 32,927 star-forming galaxies in the CANDELS fields and 68,718 star-forming galaxies in the UltraVISTA field, of which 422 and 881 are directly detected in the soft (0.5–2 keV) X-ray band in CANDELS and UltraVISTA respectively. Figure 2 shows the distributions of stellar masses and redshifts of the galaxy samples and the X-ray detected sources from the two surveys. Grey dotted lines indicate the bins of stellar mass and redshift that are used in Section 3 below and are required to lie wholly above the stellar mass completeness limit of CANDELS. The higher stellar mass completeness limit applied to the UltraVISTA catalogue, however, cuts through some of our stellar mass bins, and thus galaxies with lower stellar masses and higher redshifts may be under-represented in some bins. Removing sources from one part of a bin due to the mass completeness limits should not have an impact on our results under the reasonable assumption (implicit in our analysis) that the true distribution of luminosities does not change significantly over a stellar mass–redshift bin. The numbers of galaxies in each stellar mass–redshift bin (and the number with significant soft X-ray detections) are provided in each of the panels of Figure 3.

3 THE INTRINSIC DISTRIBUTIONS OF X-RAY LUMINOSITIES WITHIN SAMPLES OF STAR-FORMING GALAXIES

In this section we measure the distributions of X-ray luminosities within our star-forming galaxy sample, divided into bins of stellar mass and redshift. We use a non-parametric Bayesian mixture modelling approach—described fully in Appendix B below—to re-

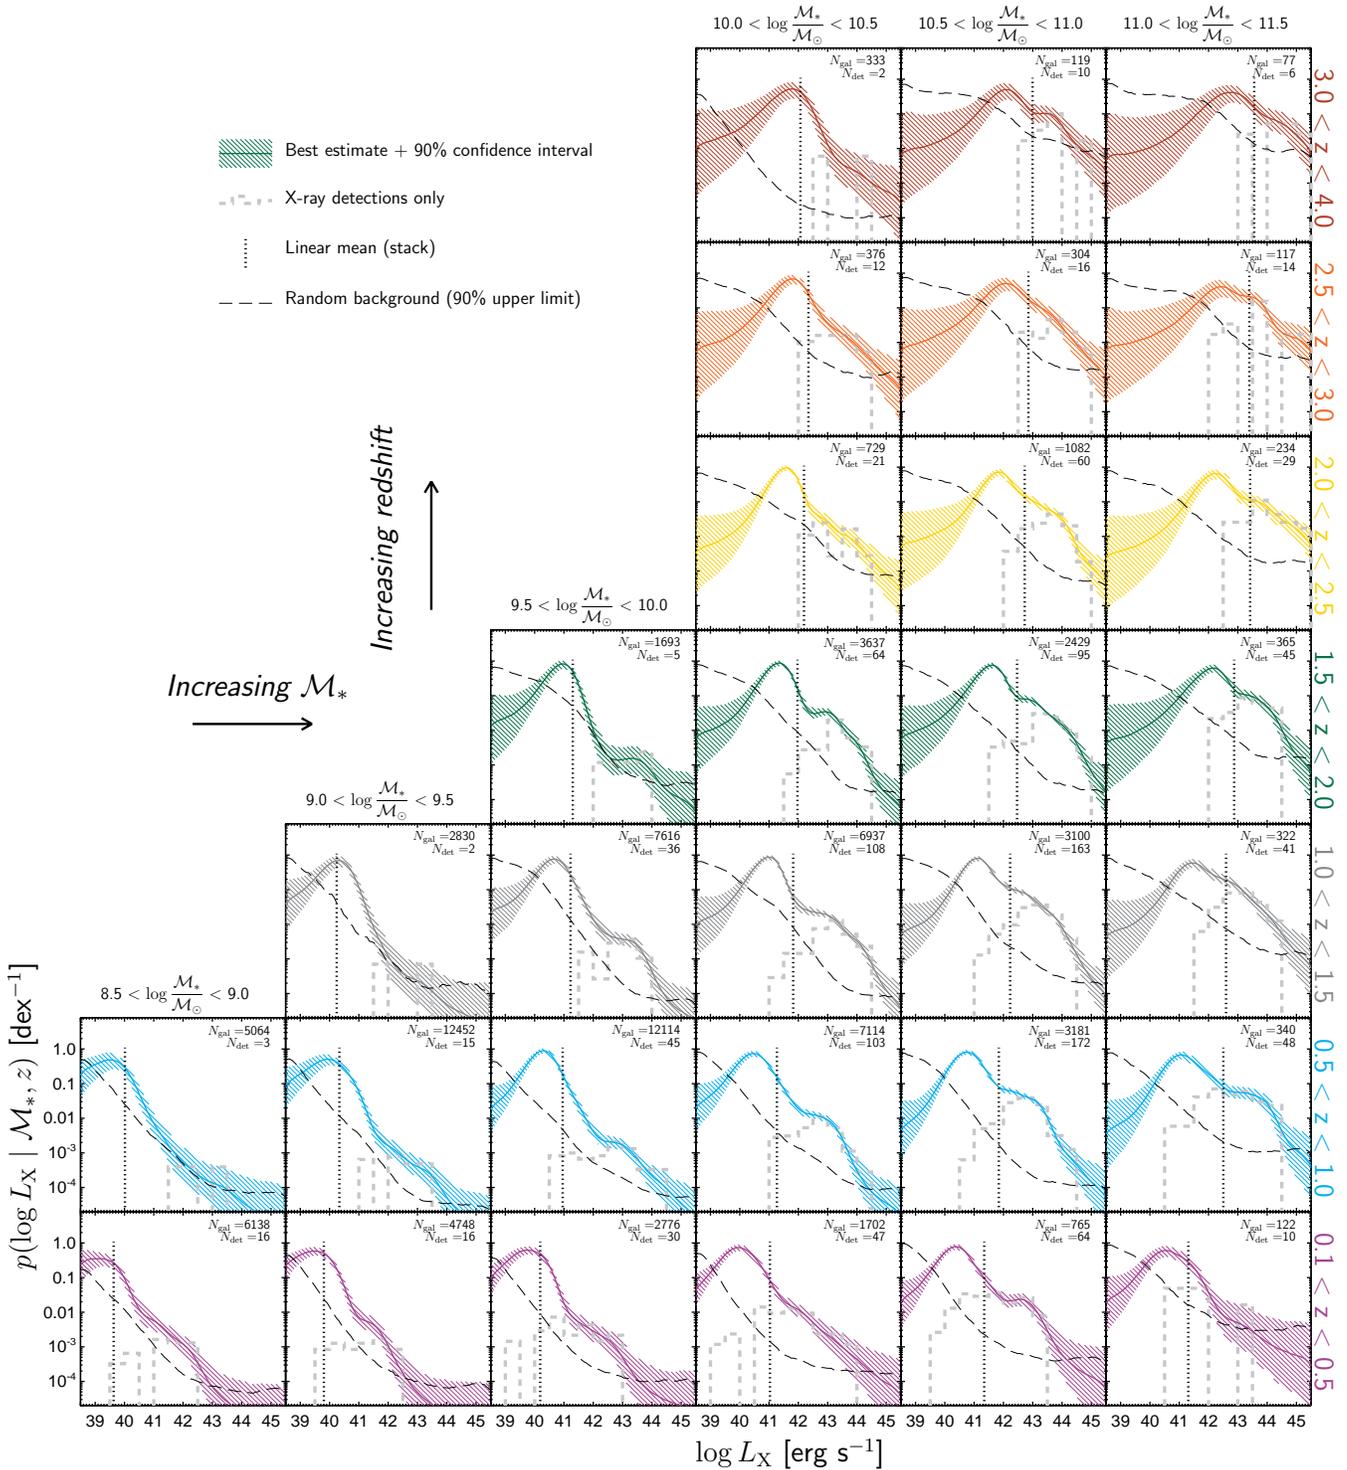


Figure 3. Measurements of the intrinsic probability distributions of X-ray luminosities for samples of star-forming galaxies, binned by stellar mass and redshift. The thick coloured lines show the best estimate of $p(\log L_X | M_*, z)$ using our flexible Bayesian mixture modelling approach (shaded regions indicate the 90 per cent confidence intervals). The grey dashed histograms show the observed distributions of X-ray luminosities for galaxies that satisfy the nominal X-ray detection criterion in the 0.5–2 keV band. Our best estimates of $p(\log L_X | M_*, z)$ use data from all galaxies in a sample (including non-detections) and account for the varying depths of the X-ray imaging to recover the intrinsic distributions, which differ significantly from these observed distributions. The total number of galaxies and the total number of X-ray detections in each bin is given in the legend of each panel. The vertical dotted lines indicate the mean X-ray luminosity from a simple stacking analysis (including both detected X-ray sources and non-detections). The thin black dashed curves are calculated by randomly shifting the positions of all galaxies in a given stellar mass–redshift bin and repeating our Bayesian analysis for this random sampling of the *Chandra* background (the curve indicates the 90 per cent upper limit on the recovered distribution). In all panels, we observe a significant peak at low luminosities ($L_X \lesssim 10^{42} \text{ erg s}^{-1}$) in our best estimates of $p(\log L_X | M_*, z)$, which traces the distribution of SFRs in each galaxy sub-sample, and a tail to higher luminosities that traces the distribution of AGN accretion activity.

cover $p(\log L_X | \mathcal{M}_*, z)$, the underlying probability distribution function of X-ray luminosities for galaxies of a given stellar mass and redshift, with limited assumptions on the shape of the distribution (other than a prior that prefers smoothly varying distributions, see also Buchner et al. 2015). We derive rest-frame 2–10 keV luminosities (hereafter, L_X) based on the observed count rates in the soft (0.5–2 keV) *Chandra* energy band, adopting a fixed conversion factor that assumes an intrinsic X-ray spectral shape with photon index $\Gamma = 1.9$, subject to Galactic absorption only⁵. We use data from the 0.5–2 keV *Chandra* band as previous studies have shown that this energy range is most sensitive to the star-forming galaxy population (e.g. Lehmer et al. 2012; Aird et al. 2015). Our method accounts for the Poisson nature of the X-ray data, thus allowing for uncertainties in the X-ray flux measurements and naturally correcting for the effects of the Eddington bias (see also Laird et al. 2009; Aird et al. 2010, 2015).

Figure 3 presents our estimates of the intrinsic distributions of L_X within sub-samples of the star-forming galaxy population, divided into bins of stellar mass (increasing to the right) and redshift (increasing towards higher panels). The solid coloured lines show our best estimates of the intrinsic distributions; the shaded regions indicate the 90 per cent confidence intervals. Our estimates of the intrinsic distributions differ significantly from the observed distributions for X-ray detected sources, shown by the grey histograms in each panel. Our method accounts for the varying sensitivity of the X-ray imaging over the CANDELS and UltraVISTA fields (due to both the overall differences in exposure times between fields and the variation in sensitivity *within* a field) and thus applies a completeness correction to the X-ray detected populations.⁶

Additionally, our method uses the X-ray data for every galaxy in our sample. We are thus able to use information from sources that fall just below the nominal X-ray detection threshold and the combined X-ray information from the large numbers of galaxies in our samples, which may not individually produce detectable X-ray emission but when combined provide important constraints. Fully utilising all of this information allows us, in many cases, to probe substantially below the luminosity of the faintest detected source and constrain the shape of $p(\log L_X | \mathcal{M}_*, z)$.

In each panel of Figure 3, our recovered distributions show a relatively narrow peak at lower luminosities ($L_X \lesssim 10^{42}$ erg s⁻¹) with a broad tail to higher luminosities and possibly a second higher-luminosity peak. We propose that this observed structure is related to the two different processes that produce the X-ray emission in star-forming galaxies: 1) the integrated emission from X-ray binaries, hot gas and supernovae, which roughly trace the SFR and results in the low-luminosity peak; and 2) AGN accretion activity, which produces the high luminosity tail. Our interpretation is consistent with the expected X-ray luminosities due to star for-

mation in normal galaxies and the previous identification of star-forming galaxies at faint X-ray fluxes in deep survey fields (e.g. Georgakakis et al. 2006; Lehmer et al. 2012; Aird et al. 2015).

We note that in some cases the position of the low-luminosity peak falls within the range of luminosities probed by the directly detected X-ray populations (shown by the grey dashed histograms). To further assess whether these peaks are real features, rather than an artefact of our analysis, we randomly shift the positions of all galaxies in a given stellar mass–redshift bin by 30–60'' and repeat our full analysis: extracting the total counts (from the *Chandra* images), background estimates (from the background maps) and exposures at the new position of each galaxy; then applying our Bayesian analysis to produce an estimate of $p(\log L_X | \mathcal{M}_*, z)$ for this random sampling of the *Chandra* background. The thin black dashed curves in Figure 3 show the upper 90 per cent confidence interval from the random background samples.⁷ In some panels, the 90 per cent upper limits from the background samples exceed the upper limits from the real data at the very lowest luminosities (below the peaks observed in the real data). These differences reflect our requirement that $p(\log L_X | \mathcal{M}_*, z)$ must integrate to unity; the existence of sources at higher luminosities in the real data thus requires a smaller fraction of galaxies with the lowest luminosities. While our nominal uncertainties are larger at these very low luminosities, the exact shape and normalization of our estimates of $p(\log L_X | \mathcal{M}_*, z)$ should be treated with caution in this regime. Nonetheless, the low-luminosity peaks in the observed distributions all lie significantly above the upper limits from the background samples in all panels. Furthermore, we do *not* observe a low-luminosity peak in our random background samples, indicating that this feature is not due to uncertainties in our modelling of the *Chandra* background. We also note that the existence of a distinct, narrow peak in the distribution of L_X goes against the underlying prior in our Bayesian method which prefers a smoothly varying distribution. We are therefore confident that the observed low-luminosity peaks are real features in our data and that we are able to track their positions accurately over a wide range of redshifts, luminosities and stellar masses. The validity of our Bayesian method and our ability to recover the position of L_X^{mode} in the different stellar mass and redshift bins is explored further using extensive simulations described in Appendix C below.

The vertical black dotted lines in Figure 3 show the results of a simple X-ray stacking analysis for the galaxies in each stellar mass–redshift bin (including X-ray detections). Stacking recovers the linear mean of the distributions but loses the information on the shape of the underlying distributions that is recovered by our Bayesian method. To trace the low-luminosity peak via a stacking analysis would require the bulk of higher luminosity AGN to be excluded to ensure the signal is dominated by the galactic emission. Our results show that appropriate threshold depends on stellar mass and redshift. Furthermore, it is difficult with a stacking analysis to remove the signal from AGNs that fall below the nominal X-ray detection thresholds, which are traced correctly and separated in our analysis.

The positions of the low-luminosity peaks provide estimates of the *mode* of the distribution of X-ray luminosities from star formation. By tracing the position of the peak as a function of stellar mass and redshift, we are able to track the mode of the main sequence of star formation via the X-ray emission. We note that a dis-

⁵ We do not apply any corrections for *intrinsic* absorption, local to the X-ray source when estimating luminosities and thus our distribution functions track the *observed* luminosity. The X-ray emission that traces the overall SFR of a galaxy, coming from the integrated emission from X-ray binaries throughout the host galaxy, is not expected to be subject to significant line-of-sight absorption (even in relatively dusty galaxies) and thus can be tracked by the observed X-ray luminosity.

⁶ A additional concern is the difference in the stellar mass completeness limits of the CANDELS and UltraVISTA fields. We have verified that our results are not significantly altered if we exclude galaxies from the UltraVISTA field in stellar mass–redshift bins that are intersected by the completeness limits of UltraVISTA (see Figure 2) and thus only use galaxies from the deeper CANDELS/3DHST fields for these bins.

⁷ An *a priori* assumption that there is an underlying distribution of L_X is implicit in our Bayesian modelling. See also panel e) of Figure C1.

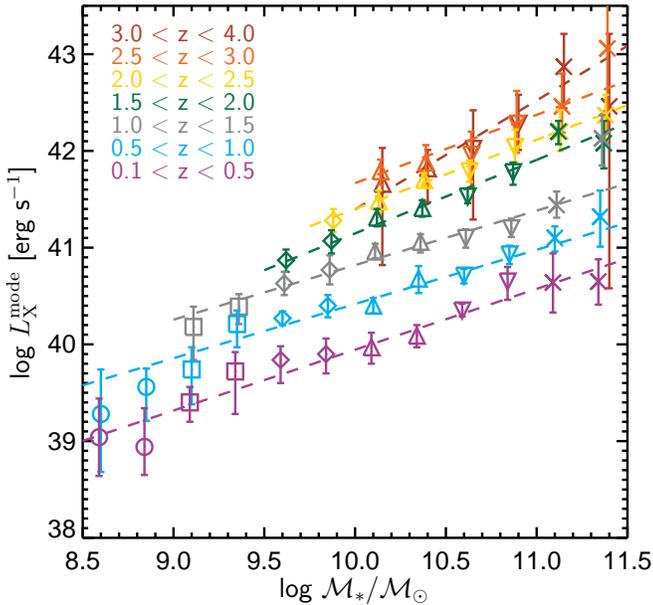


Figure 4. The peak of the distribution of X-ray luminosities from star-forming galaxies, L_X^{mode} , as a function of stellar mass and redshift, revealing the “X-ray main sequence.” Errors on L_X^{mode} are 1σ equivalent uncertainties, based on the combination of our jack-knife resampling and the statistical uncertainty from the posterior distributions of L_X^{mode} in our Bayesian analysis. In each redshift bin we find a linear relationship between $\log L_X^{\text{mode}}$ and $\log M_*$, which we fit using Equation 2 (dashed lines are the fits in each redshift bin). We find that the X-ray main sequence has an approximately constant slope of ≈ 0.6 and a normalization that increases with redshift as $\sim (1+z)^{3.79 \pm 0.12}$ (see Figure 5 and Section 4 for further details).

tinct low-luminosity peak is seen up to high redshifts ($z \sim 3-4$), even in lower mass galaxies ($M_* \lesssim 10^{10.5} M_\odot$), indicating that with our advanced method we are able to probe star formation via the X-ray emission even for the faintest star-forming galaxies in the CANDELS fields.

4 THE X-RAY MAIN SEQUENCE IN STAR-FORMING GALAXIES

In this section we use the peaks in the X-ray luminosity distributions identified in Section 3 above (and shown in Figure 3) to constrain the “X-ray main sequence” as a function of stellar mass and redshift. Figure 4 shows our measurements of the position of the low-luminosity peak, tracking the mode of the distribution of X-ray luminosity due to star formation processes (hereafter L_X^{mode}), as a function of stellar mass for our various redshift bins. We have further sub-divided our star-forming galaxy sample into 0.25 dex wide bins of stellar mass, measuring L_X^{mode} for each bin, to accurately track the stellar mass dependence. We note that our flexible Bayesian modelling allows us to isolate the low luminosity peak from the broader distribution of luminosities related to AGNs and thus AGN contamination should not have a significant impact on our measurements of L_X^{mode} , in contrast to simpler stacking analyses (see Section 6.2 for further discussion).

To estimate the uncertainty in L_X^{mode} we adopt a jack-knife resampling approach (e.g. Lupton 1993), whereby we sequentially

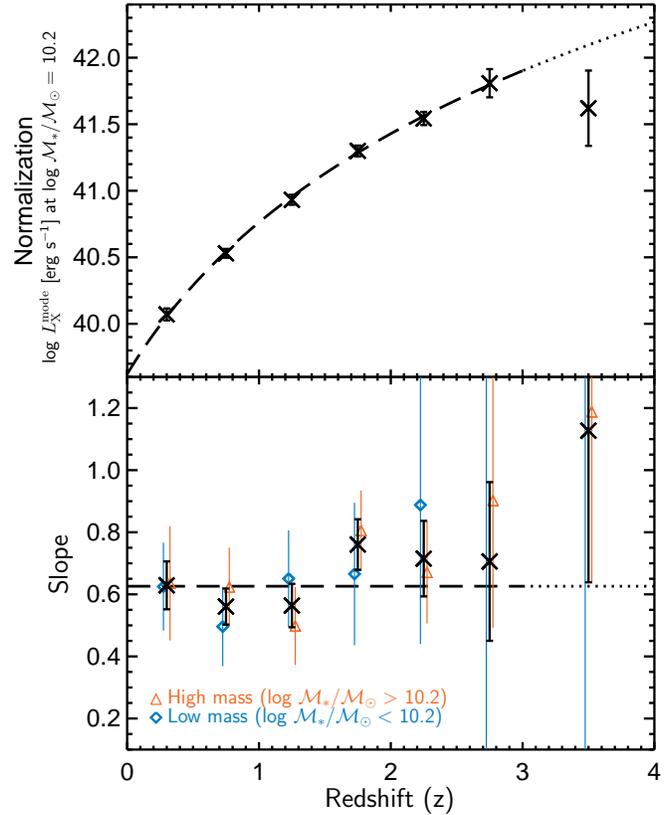


Figure 5. Normalization and slope of the X-ray main sequence as a function of redshift. Black crosses indicate the normalization and slope from fits to Equation 2 in each individual redshift bin. The black dashed line is the result of the overall fit, allowing for evolution in the normalization and assuming a constant slope at all redshifts (see Equation 3). The dotted black line shows the fit at $z > 3$ where our parametrization does not provide a good fit to the evolution. In the bottom panel we also show individual constraints for the high- and low-mass slopes when fitting with a broken power-law function in an individual redshift bin. There is no evidence for a difference between the high- and low-mass slopes at any redshift, although the low-mass slope is very poorly constrained at $z \gtrsim 2$.

remove an individual field and repeat our full analysis on the remaining fields. For the jack-knife re-sampling, we split the large UltraVISTA field into eight sub-fields of approximately equal size such that each sub-field contains a similar number of star-forming galaxies as an individual CANDELS field (to within ~ 20 per cent). The jack-knife re-sampling approach allows us to assess any impact from “cosmic variance” (the sample variance due to the existence of large-scale cosmological structures in our different fields) and to check for any systematic variations due to the differing data quality over our fields. We are able to identify a low-luminosity peak in $p(\log L_X | M_*, z)$ in every bin for each of our jack-knife samples. Most reassuringly, the existence and position of the recovered peak is not significantly altered by the exclusion of the GOODS-S field, which has the deepest X-ray data, indicating that our results are not dominated by this field. We find that the dispersion in L_X^{mode} between jack-knife samples is generally smaller than the statistical error (based on the posterior distributions from our Bayesian analyses) for an individual jack-knife sample. We therefore combine the posterior distributions from each jack-knife sample (giving each equal weight) and determine the 68 per cent central confidence in-

terval (1σ equivalent) for L_X^{mode} from the combined posterior distribution. Our errors on L_X^{mode} in Figure 4 are therefore conservative, allowing for both the statistical uncertainty and field-to-field variations. Our best estimates of L_X^{mode} and the uncertainties are given in Table 1.

Figure 4 reveals a clear trend whereby L_X^{mode} is higher for higher stellar mass galaxies—revealing a main sequence in X-rays—and the overall normalization of this relation increases with redshift.⁸ We fit the main sequence in an individual redshift bin with a log-linear relationship,

$$\log L_X^{\text{mode}}(\mathcal{M}_*) [\text{erg s}^{-1}] = a + b \left(\log \frac{\mathcal{M}_*}{\mathcal{M}_\odot} - 10.2 \right). \quad (2)$$

We adopt a standard χ^2 fitting approach, using our estimates of the uncertainties in L_X^{mode} described above. A fully Bayesian analysis of the slope and normalization of the main sequence is beyond the scope of this paper and is deferred to future studies. The dashed lines in Figure 4 indicate our log-linear fits in each redshift bin.

Figure 5 shows the normalization (top panel) and slope (bottom panel) of our fits to the X-ray main sequence as a function of redshift. Our results are consistent with a constant slope of $b \approx 0.6$ at all redshifts. We also attempt to fit a broken power-law relation with a different slope above and below $\log \mathcal{M}_*/\mathcal{M}_\odot = 10.2$ (following Whitaker et al. 2014), shown by the red and blue points in the lower panel of Figure 5 for the high- and low-mass slopes, respectively. We find no significant evidence for a difference between the high- and low-mass slopes at any redshift for our X-ray main sequence.

We observe a strong evolution in the overall normalization of the X-ray main sequence as a function of redshift out to at least $z \sim 3$. We attempt to fit this overall evolution as

$$\log L_X^{\text{mode}}(\mathcal{M}_*, z) [\text{erg s}^{-1}] = a + b \left(\log \frac{\mathcal{M}_*}{\mathcal{M}_\odot} - 10.2 \right) + c \log \frac{1+z}{1+z_0} \quad (3)$$

where we fix $z_0 = 1.0$ and assume the same slope of the X-ray main sequence, b , at all redshifts. Our overall fit to the X-ray main sequence gives $a = 40.76 \pm 0.02$ with a constant slope of $b = 0.63 \pm 0.03$ and redshift evolution described by $c = 3.79 \pm 0.12$. The black dashed lines in Figure 5 show our best overall fit for the evolving normalization and constant slope. We note that our parametrization may not be a good fit to the evolution at $z \gtrsim 3$ (indicated by the dotted line in Figure 5), where there is weak evidence ($\sim 2\sigma$) that the normalization does not increase further toward higher redshifts (see also Speagle et al. 2014).

We discuss these findings and the implications for the evolution of the galaxy population in Section 6.3 below.

5 THE RELATIONSHIP BETWEEN X-RAY LUMINOSITY AND STAR FORMATION RATE

Our results reveal an ‘‘X-ray main sequence’’ in star-forming galaxies over a wide range of redshifts and stellar masses. To interpret this relation, we need to understand how the observed X-ray luminosities are related to the SFRs of galaxies. As discussed in Section 1, the overall X-ray emission from normal star-forming galaxies (i.e. with negligible current AGN activity) is predominantly due

Table 1. Measurements of the X-ray main sequence (see Figures 4 and 8).

Redshift (z)	$\log \mathcal{M}_*$ [\mathcal{M}_\odot]	$\log L_X^{\text{mode}}$ [erg s^{-1}]	$\log \text{SFR}_X^{(a)}$ [$\mathcal{M}_\odot \text{ yr}^{-1}$]	$\log \text{SFR}_X^{\text{lim}(b)}$ [$\mathcal{M}_\odot \text{ yr}^{-1}$]
0.10–0.50	8.50–8.75	$39.04^{+0.40}_{-0.40}$	$-0.71^{+0.48}_{-0.48}$	—
0.10–0.50	8.75–9.00	$38.94^{+0.40}_{-0.29}$	$-0.83^{+0.48}_{-0.35}$	—
0.10–0.50	9.00–9.25	$39.40^{+0.16}_{-0.20}$	$-0.28^{+0.19}_{-0.24}$	–0.55
0.10–0.50	9.25–9.50	$39.72^{+0.20}_{-0.44}$	$0.11^{+0.24}_{-0.53}$	–0.50
0.10–0.50	9.50–9.75	$39.84^{+0.14}_{-0.24}$	$0.25^{+0.17}_{-0.29}$	–0.08
0.10–0.50	9.75–10.00	$39.90^{+0.16}_{-0.20}$	$0.32^{+0.19}_{-0.24}$	–0.02
0.10–0.50	10.00–10.25	$39.97^{+0.15}_{-0.17}$	$0.41^{+0.18}_{-0.20}$	0.01
0.10–0.50	10.25–10.50	$40.09^{+0.12}_{-0.12}$	$0.55^{+0.13}_{-0.14}$	0.14
0.10–0.50	10.50–10.75	$40.35^{+0.07}_{-0.06}$	$0.86^{+0.08}_{-0.07}$	0.61
0.10–0.50	10.75–11.00	$40.65^{+0.19}_{-0.15}$	$1.23^{+0.18}_{-0.23}$	0.73
0.10–0.50	11.00–11.25	$40.64^{+0.30}_{-0.31}$	$1.21^{+0.36}_{-0.37}$	–1.00
0.10–0.50	11.25–11.50	$40.65^{+0.23}_{-0.23}$	$1.23^{+0.28}_{-0.29}$	–1.00
0.50–1.00	8.50–8.75	$39.28^{+0.46}_{-0.60}$	$-0.63^{+0.55}_{-0.72}$	—
0.50–1.00	8.75–9.00	$39.56^{+0.19}_{-0.35}$	$-0.30^{+0.23}_{-0.42}$	–0.77
0.50–1.00	9.00–9.25	$39.74^{+0.23}_{-0.34}$	$-0.08^{+0.28}_{-0.43}$	–0.60
0.50–1.00	9.25–9.50	$40.21^{+0.34}_{-0.24}$	$0.49^{+0.43}_{-0.29}$	—
0.50–1.00	9.50–9.75	$40.27^{+0.07}_{-0.07}$	$0.56^{+0.08}_{-0.08}$	—
0.50–1.00	9.75–10.00	$40.40^{+0.11}_{-0.12}$	$0.72^{+0.13}_{-0.14}$	—
0.50–1.00	10.00–10.25	$40.40^{+0.08}_{-0.08}$	$0.72^{+0.10}_{-0.10}$	0.53
0.50–1.00	10.25–10.50	$40.68^{+0.13}_{-0.15}$	$1.05^{+0.16}_{-0.18}$	0.76
0.50–1.00	10.50–10.75	$40.71^{+0.07}_{-0.08}$	$1.09^{+0.08}_{-0.10}$	0.72
0.50–1.00	10.75–11.00	$40.93^{+0.09}_{-0.10}$	$1.35^{+0.11}_{-0.12}$	0.88
0.50–1.00	11.00–11.25	$41.10^{+0.12}_{-0.14}$	$1.56^{+0.14}_{-0.17}$	0.67
0.50–1.00	11.25–11.50	$41.32^{+0.27}_{-0.31}$	$1.82^{+0.33}_{-0.37}$	–1.00
1.00–1.50	9.00–9.25	$40.18^{+0.21}_{-0.34}$	$0.27^{+0.25}_{-0.41}$	—
1.00–1.50	9.25–9.50	$40.39^{+0.13}_{-0.18}$	$0.53^{+0.16}_{-0.22}$	—
1.00–1.50	9.50–9.75	$40.63^{+0.12}_{-0.12}$	$0.82^{+0.14}_{-0.14}$	—
1.00–1.50	9.75–10.00	$40.77^{+0.15}_{-0.15}$	$0.99^{+0.18}_{-0.18}$	—
1.00–1.50	10.00–10.25	$40.96^{+0.08}_{-0.07}$	$1.21^{+0.10}_{-0.08}$	—
1.00–1.50	10.25–10.50	$41.06^{+0.05}_{-0.11}$	$1.34^{+0.10}_{-0.13}$	1.16
1.00–1.50	10.50–10.75	$41.10^{+0.09}_{-0.10}$	$1.38^{+0.11}_{-0.12}$	1.00
1.00–1.50	10.75–11.00	$41.20^{+0.10}_{-0.09}$	$1.50^{+0.12}_{-0.11}$	0.75
1.00–1.50	11.00–11.25	$41.44^{+0.14}_{-0.13}$	$1.79^{+0.17}_{-0.16}$	0.66
1.00–1.50	11.25–11.50	$42.13^{+0.24}_{-0.28}$	$2.62^{+0.29}_{-0.34}$	2.18
1.50–2.00	9.50–9.75	$40.87^{+0.11}_{-0.11}$	$0.97^{+0.13}_{-0.14}$	—
1.50–2.00	9.75–10.00	$41.07^{+0.13}_{-0.13}$	$1.21^{+0.14}_{-0.16}$	—
1.50–2.00	10.00–10.25	$41.31^{+0.09}_{-0.09}$	$1.50^{+0.11}_{-0.11}$	—
1.50–2.00	10.25–10.50	$41.41^{+0.09}_{-0.09}$	$1.62^{+0.11}_{-0.11}$	—
1.50–2.00	10.50–10.75	$41.53^{+0.08}_{-0.07}$	$1.76^{+0.10}_{-0.08}$	1.63
1.50–2.00	10.75–11.00	$41.78^{+0.13}_{-0.13}$	$2.06^{+0.16}_{-0.16}$	1.81
1.50–2.00	11.00–11.25	$42.20^{+0.11}_{-0.13}$	$2.57^{+0.13}_{-0.16}$	—
1.50–2.00	11.25–11.50	$42.08^{+0.23}_{-0.26}$	$2.42^{+0.28}_{-0.31}$	–1.00
2.00–2.50	9.75–10.00	$41.28^{+0.12}_{-0.12}$	$1.34^{+0.14}_{-0.14}$	—
2.00–2.50	10.00–10.25	$41.48^{+0.11}_{-0.11}$	$1.58^{+0.13}_{-0.13}$	—
2.00–2.50	10.25–10.50	$41.70^{+0.07}_{-0.06}$	$1.85^{+0.08}_{-0.07}$	—
2.00–2.50	10.50–10.75	$41.79^{+0.11}_{-0.11}$	$1.96^{+0.13}_{-0.13}$	1.78
2.00–2.50	10.75–11.00	$42.04^{+0.18}_{-0.16}$	$2.26^{+0.22}_{-0.19}$	1.97
2.00–2.50	11.00–11.25	$42.20^{+0.22}_{-0.20}$	$2.45^{+0.27}_{-0.24}$	1.88
2.00–2.50	11.25–11.50	$42.37^{+0.21}_{-0.25}$	$2.66^{+0.25}_{-0.30}$	1.30
2.50–3.00	10.00–10.25	$41.79^{+0.12}_{-0.14}$	$1.86^{+0.14}_{-0.17}$	—
2.50–3.00	10.25–10.50	$41.88^{+0.18}_{-0.23}$	$1.96^{+0.22}_{-0.28}$	1.66
2.50–3.00	10.50–10.75	$41.96^{+0.24}_{-0.34}$	$2.06^{+0.29}_{-0.41}$	0.88
2.50–3.00	10.75–11.00	$42.28^{+0.34}_{-0.30}$	$2.45^{+0.41}_{-0.36}$	1.79
2.50–3.00	11.00–11.25	$42.46^{+0.34}_{-0.27}$	$2.66^{+0.41}_{-0.33}$	1.94
2.50–3.00	11.25–11.50	$43.06^{+0.46}_{-0.86}$	$3.39^{+0.55}_{-1.04}$	–1.00
3.00–4.00	10.00–10.25	$41.66^{+0.37}_{-0.84}$	$1.57^{+0.45}_{-1.01}$	–1.00
3.00–4.00	10.25–10.50	$41.82^{+0.19}_{-0.35}$	$1.76^{+0.23}_{-0.42}$	–1.00
3.00–4.00	10.50–10.75	$42.02^{+0.40}_{-0.73}$	$2.01^{+0.48}_{-0.88}$	–1.00
3.00–4.00	10.75–11.00	$42.28^{+0.30}_{-0.30}$	$2.32^{+0.36}_{-0.36}$	–1.00
3.00–4.00	11.00–11.25	$42.87^{+0.34}_{-0.43}$	$3.03^{+0.41}_{-0.52}$	2.00
3.00–4.00	11.25–11.50	$42.46^{+0.75}_{-1.88}$	$2.54^{+0.90}_{-2.27}$	–1.00

(a) Estimate of the SFR corresponding to the X-ray main sequence, converted from the observed L_X^{mode} using Model 4 in Table 2.

(b) Lower limit on the SFR converted from the lower limit on the observed L_X^{mode} allowing for a stellar-mass-dependent contribution from LMXBs to the X-ray luminosity (i.e. using Model 6 in Table 2).

⁸ Hints of such a relation were previously identified by Lehmer et al. (2008) but over a much more limited range in \mathcal{M}_* and z .

to the combined emission from HMXBs, LMXBs and diffuse hot gas. All of these components will, to some extent, be related to the star formation history of the galaxy. The hot gas and HMXB components are expected to trace relatively recent star formation (and thus track the overall SFR), whereas the LMXB population provides a delayed tracer of star formation that persists over much longer time-scales (thus the total LMXB luminosity may instead track the total stellar mass of a galaxy rather than the current SFR).

To assess how well our measurements of L_X^{mode} are tracing the average SFRs in our star-forming galaxy samples, in Figure 6 we compare our measurements of L_X^{mode} to the median SFR (estimated using our multiwavelength ‘‘SFR ladder’’, see Appendix A) for galaxies in each of the stellar mass–redshift bins used in Figure 4. We observe a roughly linear correlation between L_X^{mode} and the median SFR, indicating that the X-ray luminosity is tracing, on average, the SFRs of the star-forming galaxy population. Performing a simple least-squares fit (with errors based on our jack-knife resampling, as described in Section 4 above), we find a linear relation of the form

$$\log L_X [\text{erg s}^{-1}] = 39.72 \pm 0.02 + \log \text{SFR} [\mathcal{M}_\odot \text{ yr}^{-1}] \quad (4)$$

or equivalently $L_X \approx 5.2 \times 10^{39} \times \text{SFR}$ (with units given in Equation 4), which is shown by the thin black line in the large left panel of Figure 6. We note that this overall correlation will be partly driven by changes in luminosity distance over the wide redshift range covered by our sample. Furthermore, the fit has a relatively large reduced χ^2 ($\chi^2/\nu = 1.96$) and thus does not provide a satisfactory description of the data. We therefore attempt to find a more general functional form to describe the relationship between SFR and L_X . The different functional forms, along with the best-fitting parameters and the resulting χ^2 values, are given in Table 2.

First, we allowed for a non-linear slope in Equation 4 (denoted as Model 2 in Table 2), which does not significantly improve the fit. Next, we allowed for a redshift dependence of the form $L_X \propto (1+z)^C$ with both a linear and non-linear slope (Models 3 and 4, respectively). Allowing for a redshift dependence (Model 3) provides a significant improvement in χ^2 (at the $> 3\sigma$ level, based on an F test) over the basic relation (Model 1). Allowing for a non-linear SFR-dependence (Model 4) subsequently improves the fit (at $> 3\sigma$) compared to Model 3 and provides a statistically acceptable fit ($\chi^2/\nu \approx 1$). The evolving relationship described by Model 4 is shown by the coloured lines in the right panels of Figure 6. The existence of a non-linear, redshift-dependent correlation between L_X^{mode} and $\text{SFR}^{\text{median}}$ is robust and not driven by the luminosity distance as it is observed *at a fixed redshift* and the median redshift of galaxies in each data point does not change within a given redshift bin (i.e. within a single small panel of Figure 6).

These results indicate that a redshift-dependent, linear scaling between $\log L_X$ and $\log \text{SFR}$ holds, *on average*, for star-forming galaxies over the range of stellar masses and redshifts probed by our study. We note that our SFR– L_X relation is for the properties of star-forming galaxies averaged over large numbers of sources at a given stellar mass and redshift. A significantly larger intrinsic scatter and potential systematic variations may exist for *individual* galaxies but would not be revealed by our work. In addition, while our SFR estimates appear robust, the derived L_X –SFR scaling relation could be affected by any remaining systematic biases (e.g. that depend on stellar mass, redshift, or a related parameter such as metallicity) in either the UV+IR or SED-based estimates that form our SFR ladder.

In Appendix D we also find evidence for X-ray emission from quiescent galaxies with comparable luminosities to our ‘‘X-

ray main sequence’’ (at $\log \mathcal{M}_*/\mathcal{M}_\odot \gtrsim 10.5$), despite the much lower SFRs of these quiescent galaxies. These results indicate the presence of LMXB populations and thus the overall X-ray emission may have a component that roughly scales with the total stellar mass of a galaxy, as well as a component due to HMXBs that directly scales with the current SFR (see also Lehmer et al. 2010; Fragos et al. 2013; Lehmer et al. 2016). To investigate this relation further, we divide our *entire* galaxy sample (combining those classified as either star-forming or quiescent in the UVJ diagram) into bins of stellar mass, redshift, and SFR. We measure $p(\log L_X | \mathcal{M}_*, z, \text{SFR})$ for each of these sub-samples and determine the position of the low-luminosity peak, L_X^{mode} . The data points in Figure 7 show our results for $z < 2$ (where our galaxy samples are large enough to provide meaningful constraints), with colours indicating different stellar mass bins. Our measurements show a clear plateau at low SFRs (for fixed \mathcal{M}_*), related to the LMXB contribution to the overall X-ray emission. The level of this plateau evolves with redshift (indicating a higher LMXB luminosity per unit galaxy stellar mass at higher redshifts). We fit these data with the following function (Model 5 in Table 2):

$$L_X [\text{erg s}^{-1}] = \alpha(1+z)^\gamma \mathcal{M}_* + \beta(1+z)^\delta \text{SFR}^\theta \quad (5)$$

where the total X-ray luminosity is assumed to be the sum of the contribution from LMXBs and HMXBs that scale with the total stellar mass and the SFR, respectively. We allow for a redshift dependence in each of the scale factors and a non-linear scaling with the SFR with exponent θ , motivated by our Model 4. We fit for α , β , γ , δ and θ using the data from all galaxies at $0.1 < z < 2$ (binned by SFR, \mathcal{M}_* and z , as shown in Figure 7) and report the best fit values in Table 2. We also apply Model 5 to our original data for star-forming galaxies only (shown in Figure 6), fixing the parameters to the best-fitting values for all galaxies. The resulting $\chi^2/\nu = 1.05$ indicates that this model also provides a good description of the data for star-forming galaxies. In addition, the parameters describing the SFR-dependence in Model 5 ($\log \beta$, δ , and θ) are statistically consistent with the equivalent parameters in Model 4 (to within 2σ). Fixing $\theta = 1.0$ gives a worse fit to the SFR-binned data from all galaxies (although the other parameters do not change significantly compared with their uncertainties) and a significantly higher χ^2 than Model 5 when applied to the star-forming galaxy data only, indicating that a non-linear scaling with SFR is required.

These results indicate a more complicated relationship between L_X and SFR than initially revealed in Figure 6. For lower stellar mass galaxies ($\log \mathcal{M}_*/\mathcal{M}_\odot \lesssim 10$) with moderate SFRs ($\gtrsim 1 \mathcal{M}_\odot \text{ yr}^{-1}$) the X-ray emission appears to be dominated by the HMXB population and thus the observed X-ray luminosity traces the SFR. However, in more massive galaxies there can be a significant contribution from LMXBs and the observed X-ray luminosity may be relatively high, even if the current SFR of the galaxy (seen at other wavelengths) is low. Nevertheless, when considering just the star-forming galaxy samples, a simple scaling between L_X and SFR (Model 4) can be applied across all stellar masses and redshifts. In such star-forming galaxies—pre-selected on the basis of their blue rest-frame optical colours indicating young stellar populations—the observed X-ray luminosities *are* correlated with the SFR estimated at other wavelengths. We discuss these findings further and compare with previous studies of the X-ray emission from normal galaxies in Section 6.1 below.

In Figure 8 we use the scaling relation between L_X and SFR for star-forming galaxies (Model 4) to convert our measurements of the ‘‘X-ray main sequence’’ to SFRs (black crosses and error bars).

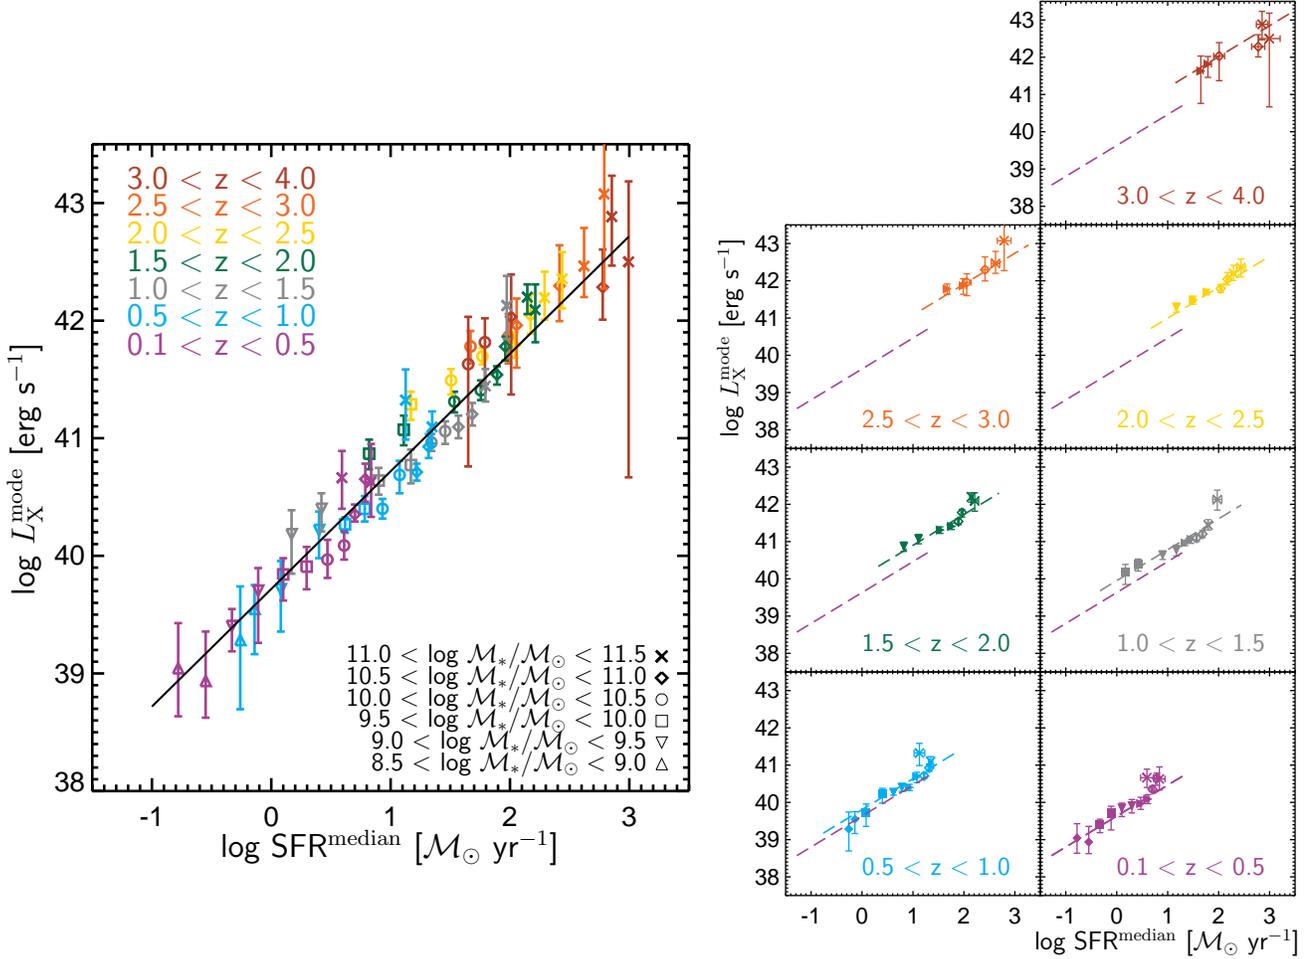


Figure 6. The peak of the distribution of X-ray luminosities from star-forming galaxies, L_X^{mode} , compared to the median SFR of each sample (based on either the optical SED or UV+IR estimates, see Appendix A). The large left panel shows the data from all redshift bins (indicated by the different colours) and stellar mass bins (indicated by the symbol types), compared to a simple linear scaling (black line, Model 1 in Table 2). The smaller right panels show the same data in each redshift bin, compared to our best fitting relation with a redshift dependence and a non-linear slope (Model 4 in Table 2). The model evaluated at $z = 0.3$ is reproduced in every panel (purple dashed line) to illustrate the evolution. Errors on L_X^{mode} are 1σ equivalent, based on the jack-knife resampling and Bayesian posteriors as in Figure 4.

These estimates of the SFR, based on the observed L_X^{mode} , are also listed in Table 1. In addition, we set a *conservative* lower limit on the SFR (black upward pointing arrows in Figure 8), based on the lower limit on L_X^{mode} and allowing for a stellar-mass dependent contribution to the X-ray luminosity (i.e. using Model 5). We believe these lower limits are conservative as we have pre-selected star-forming galaxies based on their optical colours; we know there must be significant star formation in such galaxies and a substantial population of HMXBs that contribute to the observed L_X . We also note that the conversion of our L_X measurements to SFRs applied here is reliant on the calibration relative to our multiwavelength SFR estimates and as such are *not* completely independent of the UV-to-IR data. Nevertheless, converting our measured X-ray luminosities (and the associated uncertainties) to SFRs provides a new, distinct measurement of the main sequence of star formation and allows us to compare our results with previous studies that use a wide variety of multiwavelength datasets and different methodologies (see Section 6.3 below).

6 DISCUSSION

In this paper, we measure an “X-ray main sequence” in star-forming galaxies (Section 4) and determine the scaling between L_X and SFR (Section 5). This section discusses our results and compares with previous studies. In Section 6.1 we discuss our measurements of the scaling between X-ray luminosity, SFR, and M_* , compare with previous studies, and discuss the constraints our results place on the evolution of X-ray binary populations within galaxies. Section 6.2 discusses whether emission from AGN could be contaminating our measurements and influencing our results. In Section 6.3 we convert our X-ray main sequence to SFRs, compare with previous studies of the main sequence of star formation, and discuss the implications of our X-ray-based measurements.

6.1 X-ray binary populations and the X-ray emission of normal galaxies

In Section 5 above we used multiwavelength (UV+IR or SED-based) estimates of the SFRs to determine the relationship between L_X and the SFR for our star-forming galaxy sample. We also considered all galaxies (both star-forming and quiescent) and probed

Table 2. Summary of fits for the relationship between the average X-ray luminosity (L_X^{mode}) and SED-based estimates of the SFR and stellar mass (\mathcal{M}_*). Units of L_X in erg s^{-1} , SFR in $\mathcal{M}_\odot \text{ yr}^{-1}$, and \mathcal{M}_* in \mathcal{M}_\odot are assumed for all fits.

Model	Description	Parameter	Value	Star-forming galaxies			All galaxies, binned by SFR		
				χ^2	ν	χ^2/ν	χ^2	ν	χ^2/ν
1	$\log L_X = A + \log \text{SFR}$	A	39.72 ± 0.02	117.36	60	1.96	—	—	—
2	$\log L_X = A + B \log \text{SFR}$	A	39.62 ± 0.04	111.03	59	1.88	—	—	—
		B	1.07 ± 0.03						
3	$\log L_X = A + \log \text{SFR} + C \log(1+z)$	A	39.46 ± 0.04	79.83	59	1.35	—	—	—
		C	0.76 ± 0.12						
4	$\log L_X = A + B \log \text{SFR} + C \log(1+z)$	A	39.48 ± 0.05	66.14	58	1.14	—	—	—
		B	0.83 ± 0.05						
		C	1.34 ± 0.20						
5	$L_X = \alpha(1+z)^\gamma \mathcal{M}_* + \beta(1+z)^\delta \text{SFR}^\theta$	$\log \alpha$	28.81 ± 0.08	(58.97) ^a	(56)	(1.05)	92.86	88	1.06
		γ	3.90 ± 0.36						
		$\log \beta$	39.50 ± 0.06						
		δ	0.67 ± 0.31						
		θ	0.86 ± 0.05						

^a For Model 5 the fit is performed using data from all galaxies at $0.1 < z < 2$, binned by SFR, z and \mathcal{M}_* ; here we report the resulting χ^2 when applying this function to the original star-forming galaxy samples (binned by z and \mathcal{M}_*).

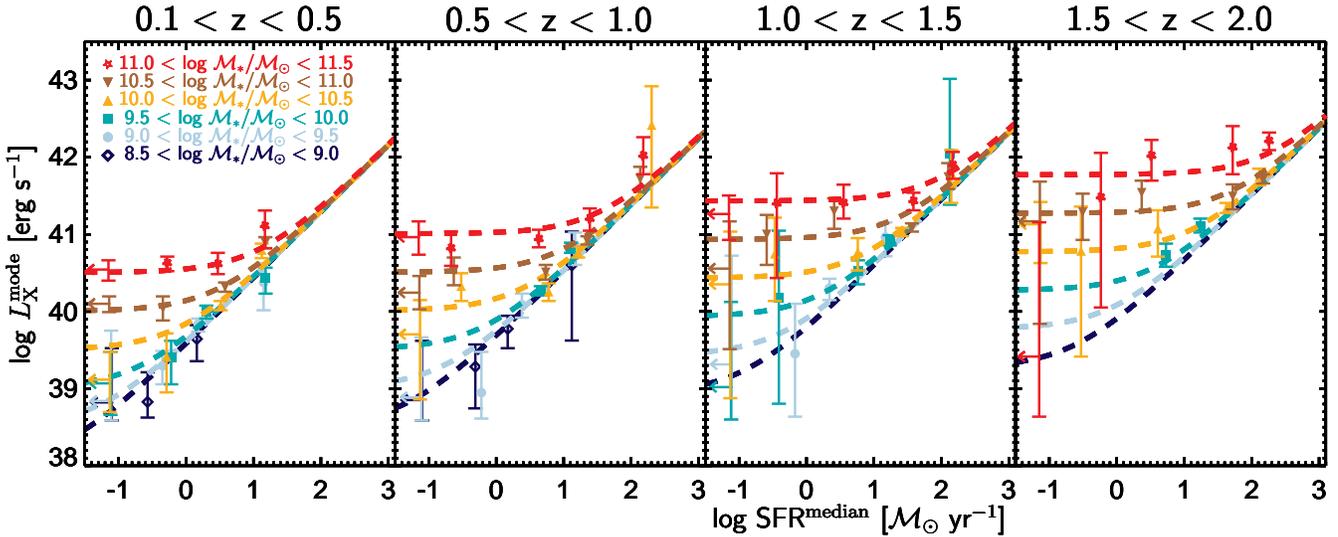


Figure 7. The peak of the distribution of X-ray luminosities, L_X^{mode} , for all galaxies—divided into bins of redshift, stellar mass, and SFR—versus the median SFR of galaxies in a given bin (data points). Colours and symbol types correspond to the stellar mass bin, as indicated. A single SFR bin is used for galaxies with $\text{SFR} < 10^{-1} \mathcal{M}_\odot \text{ yr}^{-1}$ and is indicated by the arrow at $\log \text{SFR} [\mathcal{M}_\odot \text{ yr}^{-1}] \approx -1$. The dashed lines indicate the best overall fit allowing for a stellar-mass-dependent contribution to the X-ray luminosity, as described by Model 5 in Table 2, leading to a plateau in the L_X -SFR relation at low SFR for high stellar mass galaxies.

the X-ray emission as a function of SFR, \mathcal{M}_* , and z , finding evidence for an additional contribution to L_X that scales with stellar mass. Here we compare with previous studies of the X-ray luminosities of normal galaxies and discuss the origin of the X-ray emission. We first consider our L_X -SFR relation for star-forming galaxies, which primarily traces HMXBs, and compare with previous measurements (predominantly at low redshifts). Next, we discuss the redshift evolution of the L_X -SFR relation, found in both our work and previous studies, as well as the non-linear scaling between L_X and SFR ($L_X \propto \text{SFR}^{0.83}$), that may reflect changes

in the overall stellar populations (and thus the HMXBs) across our stellar-mass-selected samples of star-forming galaxies. Finally, we discuss the additional contribution to the X-ray luminosity that scales with the stellar mass and is attributed to the LMXB population.

A number of previous studies have investigated the scaling between L_X and SFR using samples of individual galaxies with X-ray detections (e.g. Ranalli et al. 2003; Grimm et al. 2003; Persic & Rephaeli 2007; Symeonidis et al. 2011). Lehmer et al. (2010) determined the L_X -SFR relation based on a sample of

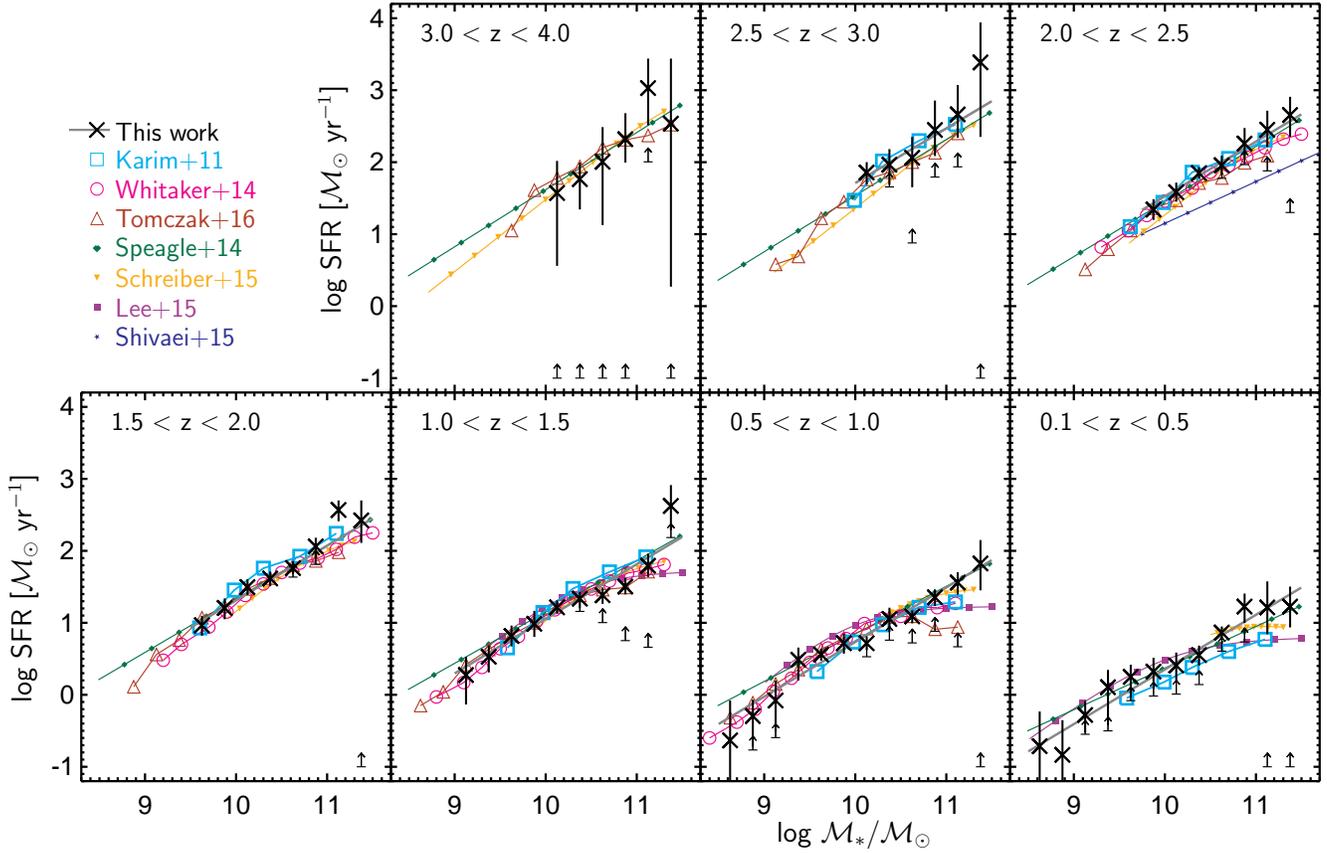


Figure 8. Our estimates of the main sequence of star formation from our X-ray analysis compared to previous measurements in the literature at other wavelengths. Black crosses indicate the X-ray based measurements from this work, converting L_X^{mode} to SFR using the scaling given by Model 4 in Table 2. The grey lines indicate our best fit for the evolution of the X-ray main sequence to $z \approx 3$ with a constant slope, converted to SFRs (see Equation 6). We do not show this fit in the highest redshift panel where there is poor agreement with our measurements. Black upward pointing arrows indicate a conservative lower limit from the X-ray data, allowing for a stellar-mass-dependent contribution from LMXBs to the observed X-ray luminosity (converting the lower limit on L_X^{mode} to an SFR using Model 5 from Table 2); the arrow is placed at $\log \text{SFR} [\mathcal{M}_\odot \text{yr}^{-1}] = -1$ when the lower limit on L_X^{mode} is consistent with the stellar-mass-dependent contribution only. Previous measurements of the main sequence of star formation are shown by the coloured symbols (we omit uncertainties for clarity). Large, open symbols indicate direct measurements that are tabulated in the following papers: Karim et al. (2011, cyan squares) based on stacking of the 1.4 GHz radio emission in the 2 deg² COSMOS field; Whitaker et al. (2014, pink circles) based on stacking of the UV+IR(24 μ m) emission in the CANDELS-3DHST fields; and Tomczak et al. (2016, dark red triangles) based on stacking of UV+IR(24 μ m) in ZFOURGE. Smaller, solid symbols instead indicate parametrizations of the main sequence from Speagle et al. (2014), Schreiber et al. (2015), Lee et al. (2015) and Shivaee et al. (2015). Previous results are only shown over the range of stellar masses and redshifts covered by a given study; the exception is Speagle et al. (2014), who compiled multiple earlier results from the literature, covering a wide stellar mass and redshift range, and thus we show their best fit across all stellar masses in all redshift panels. For Karim et al. (2011), we show their results at $z = 0.2\text{--}0.4, 0.6\text{--}0.8, 1.2\text{--}1.6$ and $1.6\text{--}2.0$ in our lowest four redshift panels (the redshift bins are identical for $z = 2.0\text{--}2.5$ and $2.5\text{--}3.0$). For Tomczak et al. (2016), we take the logarithmic mean of their narrower redshift bins for $z = 0.5\text{--}1.0$ and $1.0\text{--}1.5$. The Whitaker et al. (2014) work matches our redshift bins exactly (over the range considered). For Lee et al. (2015), we show the best fits in their $z = 0.25\text{--}0.46, 0.63\text{--}0.78$, and $1.11\text{--}1.30$ bins in our three lowest redshift panels. For Shivaee et al. (2015) we show the best fit to a sample at $z = 2.09\text{--}2.61$. For Speagle et al. (2014) and Schreiber et al. (2015) we show the best fit including redshift evolution, evaluated at the centre of our redshift bins.

66 local star-forming galaxies (including luminous and ultraluminous infrared galaxies with high SFRs). Their best-fitting linear relation ($L_X \approx 1.8 \times 10^{39} \times \text{SFR}$) gives a factor ~ 3 lower L_X per unit SFR than our best linear fit for star-forming galaxies ($L_X \approx 5.2 \times 10^{39} \times \text{SFR}$; Model 1, black line in Figure 6). However, more recent work using a sample of star-forming galaxies spanning $z \sim 0 - 1.4$ by Mineo et al. (2014) found a higher normalisation: $L_X \approx 4.0 \times 10^{39} \times \text{SFR}$ (converted to our X-ray band and assumed IMF), which is closer to our result. These discrepancies could be due to the redshift evolution of L_X -SFR relation. Indeed, extrapolating our Model 3 (with a simple redshift-dependence only) gives $L_X \approx 2.9 \times 10^{39} \times \text{SFR}$ at $z = 0$, which

is in closer agreement with the Lehmer et al. (2010) measurement for local galaxies.

Evidence for redshift evolution in the L_X -SFR relation has been found in a number of prior studies that is consistent with our measured evolution of $\sim (1+z)^{1.3}$ (e.g. Basu-Zych et al. 2013a; Lehmer et al. 2016, but see also Symeonidis et al. 2014 who find no evidence for an intrinsic evolution in the ratio of X-ray to far-IR luminosities for star-forming galaxies). The evolution most likely reflects changes in the properties of the stellar populations of galaxies (e.g. stellar ages, star formation histories, metallicities) which in turn affect the number and luminosities of HMXBs within a galaxy (e.g. Belczynski et al. 2008; Fragos et al. 2013; Tzanavaris et al.

2013). Indeed, Fragos et al. (2013) predict a higher L_X per unit SFR at higher redshifts, primarily due to the lower metallicities of high redshift galaxies that results in more numerous and more luminous HMXBs (see also Kaaret et al. 2011; Basu-Zych et al. 2013b; Douna et al. 2015). However, the extent of the predicted evolution is somewhat weaker than our observed trend indicating that younger galaxy ages at high redshift or differing star formation histories may also play a role.

The sub-linear scaling between L_X and SFR that we observe for our stellar-mass-selected samples of star-forming galaxies (i.e. $L_X \propto \text{SFR}^{0.83}$, based on Model 4) may also be related to changes in the underlying stellar populations that alter the properties of HMXBs. Our highest SFRs at a given redshift correspond to higher stellar mass samples of galaxies (see Figure 6). Thus, the expected changes in the properties of the stellar population across our mass range (e.g. higher metallicities in higher mass galaxies; Tremonti et al. 2004; Erb et al. 2006; Kewley & Ellison 2008) may alter the X-ray luminosity per unit SFR, leading to our observed non-linear dependence. Nonetheless, our results constrain the relationship between the *average* X-ray luminosity (traced by the mode of the overall distribution of L_X) and the *average* SFRs for different (complete) stellar mass limits. To reveal the underlying driver of the non-linear SFR-dependence we would need to identify and select galaxies at fixed redshift, stellar mass, SFR, age, metallicity etc. and then study their X-ray properties, which is not possible with the current dataset. Ultimately, combining X-ray luminosity measurements with data at other wavelengths may allow constraints to be placed on a variety of galaxy properties (e.g. both SFR and metallicity, see Brorby et al. 2016).

We also find evidence for a contribution to the X-ray luminosity that appears to scale more directly with the galaxy stellar mass (independent of the non-linear SFR-dependent contribution), when considering *both* star-forming and quiescent galaxies and thus probing a broader range of SFR for a given stellar mass and redshift (see Figure 7). This contribution to the X-ray luminosity is most likely due to LMXBs, which are a long-lived population that trace the result of star formation within the last ~ 100 Myr – 3 Gyr, and thus will scale more closely with the total stellar mass of a galaxy than the current SFR. Lehmer et al. (2010) also found that a scaling of L_X with both stellar mass and SFR (corresponding to the LMXB and HMXB populations, respectively) reduced the scatter compared to a single scaling between L_X and SFR. In addition, we find that the stellar-mass-dependent contribution to L_X evolves strongly with redshift, $\sim (1+z)^{3.9}$ (Model 5), indicating that the overall X-ray luminosity from LMXBs per unit stellar mass is greater at higher redshifts. Lehmer et al. (2016) stacked samples of galaxies in the CDFS (using recently acquired data bringing the total exposure to ~ 6 Ms) and also found evidence for a stellar-mass and redshift dependent contribution to the observed X-ray luminosities of normal galaxies, given by $\alpha(z) \approx 1.6 \times 10^{29} (1+z)^{3.8}$ (based on the 0.5–2 keV band, converted here to our X-ray band and IMF). This relation has a similar redshift dependence and a slightly higher normalisation than our relation, but is generally consistent with our work considering the differing methodologies and potential systematic issues in \mathcal{M}_* and SFR estimates. Lehmer et al. (2016) and Fragos et al. (2013) both attribute the strong evolution of the total X-ray luminosity from LMXBs to the younger ages of galaxies at higher redshifts (resulting in higher mass donor stars in LMXBs, increasing the X-ray luminosity per unit stellar mass, but see also Zhang et al. 2012). Ultimately, the total observed X-ray luminosity from LMXBs in a galaxy will reflect the star formation history of a galaxy and thus trace the relic of star formation over

longer time-scales than the HMXB emission or SFR tracers at other wavelengths. The strong evolution of the X-ray luminosity for our quiescent galaxy sample, $\sim (1+z)^{3.8}$, indicates that at high redshift the quiescent population may be dominated by galaxies where the star formation has been quenched relatively recently.

Our results place important constraints on the formation and evolution of X-ray binary populations out to $z \sim 4$ and indicate that large populations of HMXBs are forming at high redshifts. Some fraction of these short-lived HMXB systems will form neutron star–neutron star, black hole–neutron star or black hole–black hole binaries that will eventually merge and could be detected by Advanced LIGO and future gravitational wave experiments. Indeed, the initial Advanced LIGO discoveries indicate that black hole binary systems may be fairly common (e.g. Abbott et al. 2016; The LIGO Scientific Collaboration et al. 2016), although we defer a direct comparison of merger rates with our constraints on the high-redshift HMXB population to future studies (see also e.g. Belczynski et al. 2008).

In conclusion, the overall X-ray luminosity (from both HMXBs and LMXBs) provides a unique and independent probe of star formation in galaxies that will trace the history of star formation on different time-scales to estimates at other wavelengths (e.g. UV emission that traces recent, unobscured star formation; or far-IR emission that traces dusty, obscured star formation on slightly different time-scales). Our estimates of L_X^{mode} appear to provide a robust and relatively direct tracer of the SFRs in *star-forming* galaxies over the wide range of stellar mass and redshift probed by our study; however, X-ray based estimates of SFR remain subject to uncertainties related to the underlying stellar populations, star-formation histories, and metallicities of the galaxy population. SFR estimators at other wavelengths are also affected by these issues (see e.g. Madau & Dickinson 2014 for an overview and discussion). An advantage of the X-ray luminosity as an SFR tracer is that it should be relatively immune to the effects of dust, which severely affects UV-based SFR estimates and must be modelled correctly to determine IR-based SFRs. The sensitivity of current X-ray observatories limits the use of X-rays as an SFR tracer for individual galaxies, although our advanced analysis provides estimates of the underlying distributions and—most crucially—recovers the overall mode of the population, reducing biases compared to a more basic stacking analysis and allowing X-ray data of varying depths to be combined robustly. The existence of an “X-ray main sequence” with a single slope over a wide range of stellar masses and redshifts thus places important constraints on the evolution of the galaxy population (see Section 6.3 below).

6.2 AGN contamination

The majority of the star-forming galaxies in our sample are not individually detected in our *Chandra* X-ray imaging (at our nominal detection threshold). For most of the galaxies that *are* detected, the X-ray emission tends to be associated with an AGN with $L_X \gtrsim 10^{42}$ erg s $^{-1}$, not the galactic processes that we are tracking in this paper. A relatively weak AGN can easily overwhelm the X-ray emission associated with star formation in a galaxy (e.g. Brandt & Alexander 2015). Thus, it is worth considering the potential impact of AGNs on our estimates of L_X^{mode} for the star-forming galaxy samples and our ability to trace the main sequence of star formation via the X-ray emission.

In our analysis, we do not separate or exclude AGNs (other than optically luminous QSOs, that are also very X-ray bright), in-

stead estimating the overall distributions of X-ray luminosities.⁹ By modelling the full distribution we are able to account for the broad, roughly power-law distributions of AGN luminosities (Aird et al. 2012; Bongiorno et al. 2012; Georgakakis et al. 2014; Bongiorno et al. 2016; Jones et al. 2016, see also Paper II) and subsequently separate this component from the distinct low-luminosity peaks that we associate with galactic processes. Our technique, using the available X-ray information for all galaxies, allows for AGN emission from galaxies that are not detected in the X-ray imaging (at the nominal detection threshold). Our method should therefore be significantly more robust than previous studies, particularly stacking analyses that may not reliably exclude all AGN and could give a biased indication of the “average” luminosity from star formation (e.g. Symeonidis et al. 2014; Lehmer et al. 2016). For the AGN emission to have a significant impact on our results would require a very large population of AGNs that all have luminosities within a very narrow range, close to our observed peaks. Furthermore, a very high AGN duty cycle ($\gtrsim 60$ per cent, based on our measured luminosity distributions) would need to be inferred to make a substantial contribution and explain the peak. Thus, we expect any systematic effects due to AGN contamination are much smaller than our (relatively conservative) estimates of the uncertainties in L_X^{mode} and our ability to trace the main sequence of star formation via the X-ray emission is not significantly affected by the presence of AGNs.

6.3 The evolution of the main sequence of star formation

In this section we compare our X-ray-based work with previous measurements of the main sequence of star formation. The primary finding of our work is the identification of an “X-ray main sequence”: a linear relationship between $\log \mathcal{M}_*$ and $\log L_X^{\text{mode}}$, the mode of the distribution of X-ray luminosities, with a slope of $b = 0.63 \pm 0.03$ and a normalization that evolves as $(1+z)^{3.79 \pm 0.12}$ (see Section 4, Equation 3 and Figure 4). Our X-ray main sequence provides an independent tracer of the main sequence of star formation. However, to compare with previous measurements, we use Model 4 from Table 2, derived in Section 5, to convert L_X^{mode} to SFRs (indicated by the black crosses in Figure 8). We note that this L_X –SFR relation relies on our multiwavelength measurements of SFRs from our UV-to-IR SFR ladder. Thus, while our measurement of the main sequence of star formation is distinct from previous studies, it is not completely independent of the UV-to-IR data. We allow for the uncertainty in this conversion due to the potential contribution of LMXBs by using Model 5 (with a stellar-mass-dependent contribution to the X-ray luminosity, see Table 2) to convert the lower limit on L_X^{mode} to a conservative lower limit on SFR (indicated by the upward pointing arrows in Figure 8).

Using the conversions described above, our X-ray measurements correspond to a main sequence of star formation, given by

$$\log \text{SFR} [\mathcal{M}_\odot \text{ yr}^{-1}] \approx -7.6 + 0.76 \log \left(\frac{\mathcal{M}_*}{\mathcal{M}_\odot} \right) + 2.95 \log(1+z) \quad (6)$$

i.e. a main sequence with a constant slope $m \approx 0.76 \pm 0.06$ and a normalization that evolves with redshift as $\sim (1+z)^{2.95 \pm 0.33}$,

⁹ It is clear from our results that using a single luminosity limit of e.g. $L_X > 10^{42} \text{ erg s}^{-1}$ to distinguish AGN from star-forming galaxies is inappropriate when considering galaxies of various stellar masses and redshifts (see also discussion in Aird et al. 2015).

shown by the grey lines in Figure 8, where the uncertainties are propagated from our fit to the X-ray main sequence (Equation 3) and the L_X –SFR relation (Model 4 in Table 2). Based on our X-ray measurements, there is no evidence for a turnover or flattening of the slope of the main sequence at high stellar masses or low redshifts. However, we cannot rule out such a turnover at $z \lesssim 1$ and $\log \mathcal{M}_*/\mathcal{M}_\odot \gtrsim 10.5$ because of the uncertainties in how accurately the X-ray luminosity traces SFR in this regime (due to the potential contribution from LMXBs). Furthermore, our estimate of the slope and evolution of the main sequence in the SFR– \mathcal{M}_* plane relies on our calibration of the L_X –SFR relation. Any systematic uncertainties in our UV-to-IR SFR estimates used for the cross-calibration (e.g. a stellar mass dependence for lower metallicity galaxies) could bias our estimate of the slope and evolution of the main sequence of star formation. However, there is currently no evidence of a significant impact on our results from such systematic biases.

In Figure 8, we compare our X-ray based estimates of the main sequence to a number of previous studies based on SFR estimators at other wavelengths. In general, there is good agreement between our work and the previous studies, revealing a main sequence that rises with stellar mass and evolves strongly with redshift. A general agreement with prior measurements is expected as we have calibrated L_X as a SFR tracer using our UV-to-IR SFR ladder. Nevertheless, our method is distinct from prior studies, relying only on the overall calibration between L_X and UV-to-IR based SFRs rather than using such estimates directly for each individual galaxy. Indeed, using the UV-to-IR SEDs to estimate *both* the stellar mass and SFR directly for individual galaxies can lead to correlated uncertainties that may bias any measurement of the main sequence (Reddy et al. 2012; Speagle et al. 2014). Measurements of the main sequence that are based on far-IR data (e.g. from *Herschel*: Schreiber et al. 2015; Lee et al. 2015) or radio data (e.g. the 1.4 Ghz stacking analysis of Karim et al. 2011) should be less susceptible to these issues, although the limited depth of such data can introduce additional selection effects and biases (see Speagle et al. 2014; Ilbert et al. 2015).

We find that the normalization of the main sequence evolves strongly with redshift as $\sim (1+z)^{2.95}$ to $z \sim 3$. This evolution is consistent with most recent studies (e.g. Whitaker et al. 2014; Schreiber et al. 2015; Tasca et al. 2015). Karim et al. (2011) found a somewhat stronger evolution $\sim (1+z)^{3.9}$ based on stacking of 1.4 Ghz radio data, although their measurements (cyan squares in Figure 8) are generally consistent with our X-ray estimates, given the uncertainties, falling slightly below our data at low redshifts and slightly above at higher redshifts due to the stronger evolution. Speagle et al. (2014) also noted the discrepancy in radio-based measurements of the main sequence and proposed a redshift-dependent correction factor to bring them in line with the $\sim (1+z)^{2.8}$ evolution found by other studies. The origin of this redshift-dependent systematic offset remains unclear. Interestingly, we find a very similar evolution to Karim et al. (2011), $\sim (1+z)^{3.8}$, in our original measurements of the X-ray main sequence (i.e. tracking L_X^{mode} as a function of redshift, see Section 4 and Figure 4). The strength of this evolution is counteracted by the redshift dependence of the L_X –SFR relation (derived in Section 5). Thus, the radio-based measurements may suffer from a systematic dependence on metallicity or a related physical property of star-forming galaxies in a similar way to our X-ray measurements. Alternatively, the UV-to-IR based SFR estimates could be systematically biased and the true evolution of the main sequence may be closer to $(1+z)^{3.8}$. Further study and comparison of X-ray, ra-

dio and UV-to-IR SFR estimates out to high redshifts is required to fully reconcile these discrepancies (see also Basu et al. 2015; Magnelli et al. 2015).

Shivaei et al. (2015) measure the main sequence at $z = 2.09 - 2.61$ using dust-corrected $H\alpha$ luminosities to trace the SFR. Their best fitting relation (shown by the small dark blue squares in the $2.0 < z < 2.5$ panel of Figure 8) falls significantly below our X-ray based estimates and the other measurements from the literature, although our main sequence does fall within the range spanned by their data points for individual galaxies, which have a relatively large intrinsic scatter (~ 0.3 dex), and our slope ($m = 0.76 \pm 0.06$) is statistically consistent with their measurement ($m = 0.65 \pm 0.08$). Furthermore, Shivaei et al. find that using SED-based estimates of the SFRs yielded a steeper slope and a higher normalization, which is closer to our X-ray-based measurements and the other studies shown in Figure 8. However, using SED-based measurements of both SFR and \mathcal{M}_* may underestimate the intrinsic scatter due to underlying correlations between the parameters, and bias measurements of the slope of the main sequence, as essentially the same data are used to estimate both \mathcal{M}_* and SFR. A different fitting method (treating both \mathcal{M}_* and SFR as dependent variables) also led to a steeper slope. These findings highlight the potential issues in any measurements of the slope of the main sequence based on UV-to-IR data only. A major advantage of our study is that our X-ray-based measurements provide independent constraints on the average SFRs of galaxies (at a given stellar mass), with realistic uncertainties.

The slope of the main sequence (m) and the existence of any turnover or flattening at the highest stellar masses has been the subject of much debate (e.g. Whitaker et al. 2014; Shivaei et al. 2015; Johnston et al. 2015). The slope reflects the efficiency at which stars are forming relative to the total stellar mass. A sub-linear slope indicates that specific SFRs are lower, on average, in higher mass galaxies and thus the star-formation efficiency is lower. Recent studies probing lower mass galaxies have instead suggested a roughly linear slope (i.e. $m \approx 1$) at low stellar masses, with a break or turnover at $\mathcal{M}_* \gtrsim 10^{10} \mathcal{M}_\odot$. The slope of the main sequence above this break becomes progressively flatter as redshift decreases (e.g. Whitaker et al. 2014; Schreiber et al. 2015; Tomczak et al. 2016). Earlier studies (sensitive to a more limited range of stellar mass and redshift) may have effectively found the same behaviour but parametrized this pattern with a single slope that becomes flatter toward lower redshifts (e.g. Speagle et al. 2014). A linear slope at low masses, indicating a constant star-formation efficiency per unit stellar mass built up, is in agreement with some theoretical predictions (e.g. Somerville et al. 2008; Sparre et al. 2015), while a break or turnover at higher masses could indicate a reduction in star formation efficiency or the onset of quenching in higher mass galaxies.

Our X-ray based measurements are consistent with a main sequence of star formation with a constant slope of $m = 0.76 \pm 0.06$ (see Equation 6 above) across a wide range of stellar masses ($8.5 \lesssim \log \mathcal{M}_*/\mathcal{M}_\odot \lesssim 11.5$) and out to at least $z \sim 3$, ostensibly ruling out a linear ($m = 1$) slope at the $\gtrsim 3\sigma$ level. However, the uncertainty in this slope is somewhat larger at low masses ($\log \mathcal{M}_*/\mathcal{M}_\odot \lesssim 10.5$) and it is clear in Figure 8 that our X-ray based measurements are consistent with previous studies that measure a linear ($m \approx 1$) slope in this regime (e.g. Whitaker et al. 2014; Schreiber et al. 2015).

At higher stellar masses ($\log \mathcal{M}_*/\mathcal{M}_\odot \gtrsim 10.5$) and lower redshifts ($z \lesssim 1$) our X-ray-based main sequence is significantly steeper than prior estimates and lacks the strong turnover found

by some recent studies (e.g. Lee et al. 2015; Schreiber et al. 2015; Tomczak et al. 2016). The impact of dust on other SFR estimators is expected to be more severe at higher stellar masses; thus our X-ray based measurements may be more robust and reliable in this regime, revealing a higher level of star formation and thus a steeper main sequence that does not turn over. Our conservative lower limits (black arrows in Figure 8), however, are consistent with the high-mass turnover seen in prior studies. Nevertheless, it is important to note that our measurement of an X-ray main sequence (in terms of L_X) with a constant slope that continues to rise to high stellar masses and does not turn over is *independent* of our calibration relative to UV-to-IR SFRs. The existence of the X-ray main sequence indicates a fundamental connection between the total stellar mass and the X-ray luminosity of galaxies. Interpreting this relation relies on an understanding of how the HMXB and LMXB populations are related to other physical galaxy properties (e.g. SFR, stellar mass, metallicity, star formation history). In broad terms, however, the X-ray luminosity will provide a tracer of the star formation within a galaxy that is insensitive to dust and may effectively average star formation over longer time-scales than estimators at other wavelengths. Based on this probe, there is no evidence for a change in the efficiency of star formation at high stellar masses or the onset of quenching within the star-forming galaxy population.

Fully reconciling the various measurements of the main sequence of star formation could provide important insights into the physical nature and evolution of the galaxy population (but is beyond the scope of this paper). Different SFR estimators will be sensitive to the stellar populations of galaxies in different ways. Thus, measurements of the main sequence are affected by metallicities, star-formation histories and dust attenuation curves, as well as selection biases inherent to any individual method. Our measurements of the X-ray main sequence, providing an additional probe of the SFR with a unique dependence on these key parameters, are thus a vital addition to the current observational landscape.

7 SUMMARY AND CONCLUSIONS

In this paper, we present measurements of the ‘‘X-ray main sequence’’, tracing the main sequence of star formation out to $z \sim 4$ using new, novel methods. Here, we summarize our work and overall conclusions.

- We construct stellar-mass limited samples of star-forming galaxies out to $z = 4$, based on deep NIR-selected catalogues from the CANDELS/3DHST and UltraVISTA surveys. We extract X-ray data for all sources from the deep (~ 160 ks – 4 Ms) *Chandra* imaging.
- We measure the intrinsic probability distribution functions of X-ray luminosities at a given stellar mass and redshift. We implement a flexible Bayesian modelling technique to fully utilize the X-ray data from all galaxies, correct for incompleteness, probe below the nominal detection limits and accurately recover the underlying distribution. Our measured distributions all exhibit a clear peak at low luminosities ($L_X \lesssim 10^{42}$ erg s $^{-1}$), which we associate with star formation processes, and a broad tail to higher luminosities (associated with AGN activity).
- We track the position of the low-luminosity peak (L_X^{mode}) as a function of stellar mass and redshift, revealing an ‘‘X-ray main sequence’’ with a constant slope ($b = 0.63 \pm 0.03$) and a normalization that evolves strongly with redshift as $(1+z)^{3.79 \pm 0.12}$.

There is no evidence for a break or turnover in the X-ray main sequence at high stellar masses.

- To relate our X-ray main sequence to the main sequence of star formation, we compare our measurements of L_X^{mode} with UV-to-IR based estimates of the SFR. For our star-forming galaxy sample, we find a sub-linear scaling between the X-ray luminosity and SFR, $L_X \propto \text{SFR}^{0.83 \pm 0.05}$, that evolves with redshift as $\sim (1+z)^{1.3 \pm 0.2}$ and may be related to changes in the stellar populations of galaxies with stellar mass and redshift (that affect the number and luminosities of HMXBs). We also consider all galaxies (both star-forming and quiescent) as a function of SFR, stellar mass and redshift and find evidence for a contribution to the X-ray luminosity that appears to scale more directly with the total stellar mass and is likely due to the contribution from longer-lived LMXBs. We use these relations to convert our X-ray luminosities to SFRs to provide an X-ray based measurement of the main sequence of star formation.
- Our X-ray based measurements are consistent with an SFR– \mathcal{M}_* relation (i.e. a main sequence of star formation) with a constant slope of $m \approx 0.76 \pm 0.06$ and a normalization that evolves with redshift as $\sim (1+z)^{2.95 \pm 0.33}$. However, the stellar-mass-dependent contribution to the X-ray luminosity introduces an uncertainty such that we are unable to rule out a turnover in the main sequence at high stellar masses and low redshifts.

Our X-ray based measurements provide a robust, independent tracer of the main sequence that is not significantly affected by dust and may trace star formation over longer time-scales than probes at other wavelengths. Our work illustrates the potential of X-ray surveys to probe star formation over cosmic time. The deepest *Chandra* surveys are now reaching flux limits (over small areas) that enable direct detection of normal star-forming galaxies out to at least $z \sim 1.5$. In the future, the *Athena* X-ray observatory (Nandra et al. 2013) will reach comparable X-ray flux limits over areas of several square degrees (Aird et al. 2013), providing large, statistical samples of X-ray detected star-forming galaxies that will enable vital constraints on both the slope and scatter of the main sequence of star formation over \sim half of cosmic time. The all-sky X-ray surveys carried out with *eROSITA* (Merloni et al. 2012) will also directly detect large numbers of star-forming galaxies in the local Universe, enabling detailed studies of X-ray binary populations as a function of galaxy properties and providing a critical baseline for higher redshift studies. In the meantime, combining the full range of deep and wide *Chandra* surveys and using advanced statistical techniques to push to the limits of these data, as in our work, will allow further studies of the relation between the X-ray emission and the physical properties of galaxies, placing unique constraints on the star formation history of the universe.

ACKNOWLEDGEMENTS

We thank the referee for helpful comments that have improved this paper. We acknowledge helpful discussions with Kirpal Nandra. JA acknowledges support from ERC Advanced Grant FEEDBACK 340442. ALC acknowledges support from NSF CAREER award AST-1055081. AG acknowledges the THALES project 383549 that is jointly funded by the European Union and the Greek Government in the framework of the programme “Education and lifelong learning”. This work is based in part on observations taken by the 3D-HST Treasury Program (GO 12177 and 12328) with the NASA/ESA HST, which is operated by the Association of Uni-

versities for Research in Astronomy, Inc., under NASA contract NAS5-26555. Based in part on data obtained with the European Southern Observatory Very Large Telescope, Paranal, Chile, under Large Program 185.A-0791, and made available by the VUDS team at the CESAM data center, Laboratoire d’Astrophysique de Marseille, France. The scientific results reported in this article are based to a significant degree on observations made by the *Chandra* X-ray Observatory.

REFERENCES

- Abbott, B. P., et al. 2016, *ApJ*, 818, L22
- Aird, J., Coil, A. L., & Georgakakis, A. 2016, in prep.
- Aird, J., Coil, A. L., Georgakakis, A., Nandra, K., Barro, G., & Pérez-González, P. G. 2015, *MNRAS*, 451, 1892
- Aird, J., et al. 2012, *ApJ*, 746, 90
- . 2013, An Athena+ Supporting Paper (arXiv:1306.2325)
- Aird, J., Nandra, K., Georgakakis, A., Laird, E. S., Steidel, C. C., & Sharon, C. 2008, *MNRAS*, 387, 883
- Aird, J., et al. 2010, *MNRAS*, 401, 2531
- Alexander, D. M., et al. 2003, *AJ*, 126, 539
- Basu, A., Wadadekar, Y., Beelen, A., Singh, V., Archana, K. N., Sirothia, S., & Ishwara-Chandra, C. H. 2015, *ApJ*, 803, 51
- Basu-Zych, A. R., et al. 2013a, *ApJ*, 762, 45
- . 2013b, *ApJ*, 774, 152
- . 2007, *ApJS*, 173, 457
- Behroozi, P. S., Wechsler, R. H., & Conroy, C. 2013, *ApJ*, 770, 57
- Belczynski, K., Kalogera, V., Rasio, F. A., Taam, R. E., Zezas, A., Bulik, T., Maccarone, T. J., & Ivanova, N. 2008, *ApJS*, 174, 223
- Bell, E. F., et al. 2005, *ApJ*, 625, 23
- . 2004, *ApJ*, 608, 752
- Blanton, M. R., & Moustakas, J. 2009, *ARA&A*, 47, 159
- Bongiorno, A., et al. 2012, *MNRAS*, 427, 3103
- . 2016, *A&A*, 588, A78
- Brammer, G. B., van Dokkum, P. G., & Coppi, P. 2008, *ApJ*, 686, 1503
- Brammer, G. B., et al. 2012, *ApJS*, 200, 13
- . 2011, *ApJ*, 739, 24
- Brandt, W. N., & Alexander, D. M. 2015, *A&A Rev.*, 23, 1
- Brorby, M., Kaaret, P., Prestwich, A., & Mirabel, I. F. 2016, *MNRAS*, 457, 4081
- Brown, M. J. I., Dey, A., Jannuzi, B. T., Brand, K., Benson, A. J., Brodwin, M., Croton, D. J., & Eisenhardt, P. R. 2007, *ApJ*, 654, 858
- Brusa, M., et al. 2007, *ApJS*, 172, 353
- Buchner, J., et al. 2015, *ApJ*, 802, 89
- Capak, P., et al. 2007, *ApJS*, 172, 99
- Chabrier, G. 2003, *PASP*, 115, 763
- Ciliegi, P., Zamorani, G., Hasinger, G., Lehmann, I., Szokoly, G., & Wilson, G. 2003, *A&A*, 398, 901
- Civano, F., et al. 2016, *ApJ*, 819, 62
- Coil, A. L., et al. 2011, *ApJ*, 741, 8
- Conroy, C. 2013, *ARA&A*, 51, 393
- Conroy, C., & Gunn, J. E. 2010, *ApJ*, 712, 833
- Conroy, C., Gunn, J. E., & White, M. 2009, *ApJ*, 699, 486
- Cool, R. J., et al. 2013, *ApJ*, 767, 118
- Cowley, M. J., et al. 2016, *MNRAS*, 457, 629
- Daddi, E., et al. 2007, *ApJ*, 670, 156
- Dale, D. A., & Helou, G. 2002, *ApJ*, 576, 159
- David, L. P., Jones, C., & Forman, W. 1992, *ApJ*, 388, 82
- Domínguez Sánchez, H., et al. 2014, *MNRAS*, 441, 2
- Douna, V. M., Pellizza, L. J., Mirabel, I. F., & Pedrosa, S. E. 2015, *A&A*, 579, A44
- Elbaz, D., et al. 2007, *A&A*, 468, 33
- . 2011, *A&A*, 533, A119
- . 2010, *A&A*, 518, L29
- Elvis, M., et al. 2009, *ApJS*, 184, 158

- Erb, D. K., Shapley, A. E., Pettini, M., Steidel, C. C., Reddy, N. A., & Adelberger, K. L. 2006, *ApJ*, 644, 813
- Fabbiano, G. 1989, *ARA&A*, 27, 87
- Faber, S. M., et al. 2007, *ApJ*, 665, 265
- Feroz, F., & Hobson, M. P. 2008, *MNRAS*, 384, 449
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, 398, 1601
- Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2013, arXiv:1306.2144
- Fragos, T., et al. 2013, *ApJ*, 764, 41
- Georgakakis, A., et al. 2015, *MNRAS*, 453, 1946
- Georgakakis, A., Nandra, K., Laird, E. S., Aird, J., & Trichas, M. 2008, *MNRAS*, 388, 1205
- Georgakakis, A., et al. 2014, *MNRAS*, 440, 339
- Georgakakis, A. E., Chavushyan, V., Plionis, M., Georgantopoulos, I., Koulouridis, E., Leonidaki, I., & Mercado, A. 2006, *MNRAS*, 367, 1017
- Ghosh, P., & White, N. E. 2001, *ApJ*, 559, L97
- Gordon, K. D., Clayton, G. C., Witt, A. N., & Misselt, K. A. 2000, *ApJ*, 533, 236
- Grimm, H.-J., Gilfanov, M., & Sunyaev, R. 2003, *Astronomische Nachrichten*, 324, 171
- Grogin, N. A., et al. 2011, *ApJS*, 197, 35
- Hao, C.-N., Kennicutt, R. C., Johnson, B. D., Calzetti, D., Dale, D. A., & Moustakas, J. 2011, *ApJ*, 741, 124
- Hsu, L.-T., et al. 2014, *ApJ*, 796, 60
- Ilbert, O., et al. 2015, *A&A*, 579, A2
- Johnston, R., Vaccari, M., Jarvis, M., Smith, M., Giovannoli, E., Häußler, B., & Prescott, M. 2015, *MNRAS*, 453, 2540
- Jones, M. L., Hickox, R. C., Black, C. S., Hainline, K. N., DiPompeo, M. A., & Goulding, A. D. 2016, *ApJ*, 826, 12
- Kaaret, P., Schmitt, J., & Gorski, M. 2011, *ApJ*, 741, 10
- Karim, A., et al. 2011, *ApJ*, 730, 61
- Kelly, B. C., Fan, X., & Vestergaard, M. 2008, *ApJ*, 682, 874
- Kennicutt, R. C., & Evans, N. J. 2012, *ARA&A*, 50, 531
- Kennicutt, Jr., R. C. 1998, *ARA&A*, 36, 189
- Kewley, L. J., & Ellison, S. L. 2008, *ApJ*, 681, 1183
- Koekemoer, A. M., et al. 2011, *ApJS*, 197, 36
- Kriek, M., & Conroy, C. 2013, *ApJ*, 775, L16
- Kriek, M., et al. 2015, *ApJS*, 218, 15
- Kriek, M., van Dokkum, P. G., Labbé, I., Franx, M., Illingworth, G. D., Marchesini, D., & Quadri, R. F. 2009, *ApJ*, 700, 221
- Labbé, I., et al. 2005, *ApJ*, 624, L81
- Laird, E. S., et al. 2009, *ApJS*, 180, 102
- Laird, E. S., Nandra, K., Hobbs, A., & Steidel, C. C. 2006, *MNRAS*, 373, 217
- Le Fèvre, O., et al. 2015, *A&A*, 576, A79
- Lee, N., et al. 2015, *ApJ*, 801, 80
- Lehmer, B. D., Alexander, D. M., Bauer, F. E., Brandt, W. N., Goulding, A. D., Jenkins, L. P., Ptak, A., & Roberts, T. P. 2010, *ApJ*, 724, 559
- Lehmer, B. D., et al. 2016, *ApJ*, 825, 7
- 2008, *ApJ*, 681, 1163
- 2012, *ApJ*, 752, 46
- Luo, B., et al. 2010, *ApJS*, 187, 560
- Lupton, R. 1993, *Statistics in theory and practice*
- Lutz, D., et al. 2011, *A&A*, 532, A90
- Madau, P., & Dickinson, M. 2014, *ARA&A*, 52, 415
- Magnelli, B., et al. 2015, *A&A*, 573, A45
- 2013, *A&A*, 553, A132
- Maraston, C., Pforr, J., Renzini, A., Daddi, E., Dickinson, M., Cimatti, A., & Tonini, C. 2010, *MNRAS*, 407, 830
- Marchesi, S., et al. 2016, *ApJ*, 817, 34
- Martin, D. C., et al. 2005, *Astrophys.J.*, 619, L1
- McCracken, H. J., et al. 2012, *A&A*, 544, A156
- Merloni, A., et al. 2012, (arXiv:1209.3114)
- Mineo, S., Gilfanov, M., Lehmer, B. D., Morrison, G. E., & Sunyaev, R. 2014, *MNRAS*, 437, 1698
- Mineo, S., Gilfanov, M., & Sunyaev, R. 2012, *MNRAS*, 426, 1870
- Momcheva, I. G., et al. 2016, *ApJS*, 225, 27
- Moustakas, J., et al. 2013, *ApJ*, 767, 50
- Moustakas, J., Kennicutt, Jr., R. C., & Tremonti, C. A. 2006, *ApJ*, 642, 775
- Muzzin, A., et al. 2013a, *ApJ*, 777, 18
- 2013b, *ApJS*, 206, 8
- Nandra, K., et al. 2013, (arXiv:1306.2307)
- 2015, *ApJS*, 220, 10
- Nandra, K., Mushotzky, R. F., Arnaud, K., Steidel, C. C., Adelberger, K. L., Gardner, J. P., Teplitz, H. I., & Windhorst, R. A. 2002, *ApJ*, 576, 625
- Noeske, K. G., et al. 2007, *ApJ*, 660, L43
- Persic, M., & Rephaeli, Y. 2007, *A&A*, 463, 481
- Ranalli, P., Comastri, A., & Setti, G. 2003, *A&A*, 399, 39
- Rangel, C., Nandra, K., Laird, E. S., & Orange, P. 2013, *MNRAS*, 428, 3089
- Reddy, N. A., et al. 2015, *ApJ*, 806, 259
- Reddy, N. A., Pettini, M., Steidel, C. C., Shapley, A. E., Erb, D. K., & Law, D. R. 2012, *ApJ*, 754, 25
- Rodighiero, G., et al. 2011, *ApJ*, 739, L40
- Rovilos, E., Georgantopoulos, I., Tzanavaris, P., Pracy, M., Whiting, M., Woods, D., & Goudis, C. 2009, *A&A*, 502, 85
- Salim, S., et al. 2007, *ApJS*, 173, 267
- Salvato, M., et al. 2009, *ApJ*, 690, 1250
- 2011, *ApJ*, 742, 61
- Sanders, D. B., et al. 2007, *ApJS*, 172, 86
- Santini, P., et al. 2014, *A&A*, 562, A30
- Schreiber, C., et al. 2015, *A&A*, 575, A74
- Scoville, N., et al. 2007, *ApJS*, 172, 1
- Shivaei, I., et al. 2016, *ApJ*, 820, L23
- Shivaei, I., et al. 2016, (arXiv:1609.04814)
- Shivaei, I., et al. 2015, *ApJ*, 815, 98
- Skelton, R. E., et al. 2014, *ApJS*, 214, 24
- Somerville, R. S., Hopkins, P. F., Cox, T. J., Robertson, B. E., & Hernquist, L. 2008, *MNRAS*, 391, 481
- Sparre, M., et al. 2015, *MNRAS*, 447, 3548
- Speagle, J. S., Steinhardt, C. L., Capak, P. L., & Silverman, J. D. 2014, *ApJS*, 214, 15
- Straatman, C. M. S., et al. 2016, *ApJ*, 830, 51
- Symeonidis, M., et al. 2014, *MNRAS*, 443, 3728
- 2011, *MNRAS*, 417, 2239
- Tal, T., et al. 2014, *ApJ*, 789, 164
- Tasca, L. A. M., et al. 2015, *A&A*, 581, A54
- The LIGO Scientific Collaboration et al. 2016, (arXiv:1606.04856)
- Tomczak, A. R., et al. 2016, *ApJ*, 817, 118
- Tremonti, C. A., et al. 2004, *ApJ*, 613, 898
- Tzanavaris, P., et al. 2013, *ApJ*, 774, 136
- Viero, M. P., et al. 2013, *ApJ*, 779, 32
- Whitaker, K. E., et al. 2014, *ApJ*, 795, 104
- 2011, *ApJ*, 735, 86
- Whitaker, K. E., van Dokkum, P. G., Brammer, G., & Franx, M. 2012, *ApJ*, 754, L29
- Williams, R. J., Quadri, R. F., Franx, M., van Dokkum, P., & Labbé, I. 2009, *ApJ*, 691, 1879
- Wuyts, S., et al. 2011, *ApJ*, 738, 106
- Wuyts, S., Labbé, I., Förster Schreiber, N. M., Franx, M., Rudnick, G., Brammer, G. B., & van Dokkum, P. G. 2008, *ApJ*, 682, 985
- Wuyts, S., et al. 2007, *ApJ*, 655, 51
- Xue, Y. Q., et al. 2011, *ApJS*, 195, 10
- Zhang, Z., Gilfanov, M., & Bogdán, Á. 2012, *A&A*, 546, A36

APPENDIX A: MULTIWAVELENGTH ESTIMATES OF STAR FORMATION RATES AND SPECTRAL ENERGY DISTRIBUTION FITTING FOR GALAXIES

In this appendix, we describe our construction of an “SFR ladder” (e.g. Wuyts et al. 2011) to provide estimates of the SFR for every galaxy in our sample based on either the UV+IR emission or (when

a 24 μ m detection is unavailable) from fitting the UV-to-MIR spectral energy distributions (SEDs) with stellar population synthesis (SPS) models. Our SED fitting is performed using the FAST code (Kriek et al. 2009), subject to a number of modifications that are described in this appendix.

A widely used and reliable method of estimating the SFR of a galaxy, that we adopt when possible, is to sum the unobscured star formation traced by the UV light and the obscured star formation traced by the total IR luminosity (e.g. Kennicutt 1998; Gordon et al. 2000; Wuyts et al. 2008; Domínguez Sánchez et al. 2014; Ilbert et al. 2015). The total SFR, assuming a Chabrier (2003) initial mass function (IMF), is given by

$$\text{SFR}_{\text{UV+IR}}[\mathcal{M}_{\odot} \text{ yr}^{-1}] = 1.09 \times 10^{-10} (L_{\text{IR}} + 2.2L_{\text{UV}}) [L_{\odot}] \quad (\text{A1})$$

where L_{IR} is the total rest-frame 8–1000 μ m IR luminosity and L_{UV} is the total rest-frame 1216–3000 Å UV luminosity (Bell et al. 2005; Whitaker et al. 2014). The additional factor 2.2 accounts for UV emission outside the 1216–3000 Å range. We estimate L_{UV} from the monochromatic UV luminosity at rest-frame 2800 Å, $L_{\text{UV}} \approx 1.5\nu L_{\nu 2800}$, where $L_{\nu 2800}$ is interpolated from the observed photometry of a galaxy using our SED fits (described below, although this quantity is insensitive to the details of the SED modelling). To estimate L_{IR} , we require a $> 3\sigma$ detection in the *Spitzer*/MIPS 24 μ m imaging¹⁰ and scale the observed 24 μ m flux density to L_{IR} assuming a single template, the log average of the Dale & Helou (2002) templates (see Wuyts et al. 2008; Elbaz et al. 2010; Muzzin et al. 2013b; Whitaker et al. 2014). These estimates should be relatively insensitive to changes in the level of dust. Decreased levels of dust (e.g. due to lower metallicities) result in less obscured star formation, traced at 24 μ m, which is counteracted by an increase in the unobscured star formation, traced in the UV. However, a recent study (Shivaei et al. 2016) proposed that the extrapolation of the 24 μ m emission to a total IR luminosity is dependent on metallicity at higher redshifts ($z \sim 2$), when polycyclic aromatic hydrocarbon (PAH) features at rest-frame $\sim 7.7\mu$ m start to enter the observed 24 μ m band. As there is currently no consensus on this effect, however, we retain the standard approach of adopting a single, luminosity-independent IR template to extrapolate from the observed 24 μ m flux to the total IR luminosity at all stellar masses and redshifts (see Wuyts et al. 2008; Tomczak et al. 2016; Straatman et al. 2016).

Combining the observed UV and IR(24 μ m) emission should provide an accurate estimate of the total SFR. However, 24 μ m detections are only available for ~ 40 per cent of our star-forming galaxy sample. For the remainder, we must estimate the SFR based on the UV-to-MIR SED and apply a correction for any obscuration due to dust. This is achieved by fitting the SEDs with SPS models using FAST. The SED fits are also used to estimate the stellar masses for all of the galaxies in our sample.

FAST determines galaxy physical properties by generating SPS model templates over a grid of stellar population parameters (age, star formation time-scale τ , dust content A_V , metallicity, and redshift), comparing the predicted broadband fluxes to the observed

photometry, and adopting the template from the grid that gives the lowest χ^2 . We choose to use Flexible Stellar Population Synthesis (FSPS) models (Conroy et al. 2009; Conroy & Gunn 2010) with fixed solar metallicity¹¹ and a Chabrier (2003) Initial Mass Function (IMF). We allow for dust reddening of $A_V = 0 - 4$ magnitudes, assuming the dust attenuation curves of Kriek & Conroy (2013). We adopt “delayed- τ ” models with τ in the range 0.1–10 Gyr, which allows for both exponentially declining star formation histories and linearly rising star formation histories that are more appropriate for modelling the SEDs of high-redshift ($z \gtrsim 2$) galaxies (e.g. Maraston et al. 2010; Behroozi et al. 2013). Initially, we allow for galaxy ages in the range ~ 100 Myr – 13 Gyr (but see further discussion of a redshift-dependent minimum age below). FAST automatically excludes any templates where the age is *older* than the observable universe at a given z .

We identified a number of issues when running the standard FAST code with the parameters described above. First, we found that the recovered SFRs were restricted to a range of relatively discrete values. This discreteness is due to the coarse grid of stellar population parameters. Of more concern, we found that a non-negligible fraction of galaxies with very blue optical-to-NIR colours (classified as star-forming galaxies in the UVJ diagram) were assigned very low SFRs (in poor agreement with $\text{SFR}_{\text{UV+IR}}$, where available). Conversely, red (quiescent) galaxies could be assigned relatively large SFRs that appear at odds with their overall SEDs. Further investigation revealed that this issue was also related to the coarse parameter grid. A single stellar population template can provide the best χ^2 , but many other templates—corresponding to more realistic SFRs—give a comparable (but higher) χ^2 , especially when the photometric uncertainties are large. Further refinement of the grid for each individual galaxy is computationally prohibitive for the large number of galaxies in our sample. Instead, FAST includes a template error function, which introduces an additional uncertainty to account for the coarse parameter grid, intended to partly mitigate this issue. We include the template error function but still find that a single, unrealistic template can produce the best χ^2 . We therefore take a “semi-Bayesian” approach (see also Moustakas et al. 2013), where we weight each template according to the χ_i^2 value (for a given galaxy) and calculate a weighted average of the SFR over the entire parameter space, effectively applying a prior that prefers well-occupied regions of the parameter grid. Thus,

$$\text{SFR}_{\text{SED}} = \frac{\sum_i [\exp(-\chi_i^2/2) \times \psi_i]}{\sum_i [\exp(-\chi_i^2/2)]} \quad (\text{A2})$$

where ψ_i denotes the SFR associated with template i and the sum is taken over the entire grid of stellar population parameters. Applying this weighting scheme not only solves the issue of a single, outlier template being erroneously assigned to a galaxy but also eliminates the discreteness in the distribution of SFRs for our overall galaxy sample, producing a continuous range of values. For consistency, we also apply this χ^2 -weighted averaging when determining the stellar mass, although this has less of an impact (changing \mathcal{M}_*

¹⁰ We note that *Herschel* imaging at longer wavelengths with PACS and SPIRE is available in our fields (Elbaz et al. 2011; Lutz et al. 2011; Magnelli et al. 2013) that could be used to probe the peak of the IR emission or fit the far-IR SED. However, given the limited depths, only a very small fraction of the galaxies in our sample are actually detected by *Herschel* ($\lesssim 7$ per cent are detected in the PACS 100 μ m imaging) and thus for this work we only consider the deeper *Spitzer* 24 μ m data.

¹¹ We tested the effect of allowing a broader range of metallicities in our SED fitting (~ 0.1 – 1.5 times solar metallicity) for a representative subset of our galaxy sample. We found that our estimates of \mathcal{M}_* changed by < 0.02 dex and that our estimates of SFR_{SED} changed by < 0.05 dex, on average, with no strong systematics with \mathcal{M}_* or redshift. Thus, our SED fitting is relatively unaffected by the assumed metallicity, at least when combining large samples of galaxies, and we retain fixed solar metallicity in our final SED fitting for the full galaxy sample.

by less than 0.1 dex for the majority of galaxies and with no overall systematic shift) as the stellar masses are primarily determined by the NIR flux and are much less sensitive to the details of the SPS models.

To test the accuracy and reliability of our SED-based SFR estimates, in Figure A1 we compare with the UV+IR estimates, when available (i.e. for sources with $24\mu\text{m}$ detections). The top panel compares our initial SFR_{SED} estimates with the parameters described above, setting the minimum galaxy age to $t_{\text{min}} = 100$ Myr. In general, there is a reasonable agreement between the two estimates. However, at higher redshifts ($z \gtrsim 2$) the SED fitting tends to underestimate the SFR (compared to SFR_{UV+IR}) by up to ~ 0.5 dex. Our limited range of SPS templates are unable to produce the young, high SFR galaxies observed at these early times (see also Reddy et al. 2012, 2015; Shivaeei et al. 2016). Reducing the minimum allowed galaxy age to 10 Myr (middle panel of Figure A1) solves this issue and improves the agreement between SFR_{SED} and SFR_{UV+IR} at $z \gtrsim 2$. However, allowing for younger templates results in an increased scatter at lower redshift and a tendency to over-estimate the SFR from the SED. To solve this issue, we implement a redshift-dependent minimum age,

$$\log t_{\text{min}}(z)[\text{Myr}] = 2.2 - 0.4z \quad (\text{A3})$$

which corresponds to setting a minimum allowed age of 100 Myr at $z = 0.5$ and 10 Myr at $z = 3$. We exclude any templates with ages less than the limit given by Equation A3 when marginalizing over the parameter grid to estimate SFR_{SED} (i.e. applying Equation A1). The bottom panel of Figure A1 shows that our redshift-dependent t_{min} ensures good agreement (to within a factor ~ 2 , on average) between SFR_{SED} and SFR_{UV+IR} over the full redshift range without the large scatter introduced when allowing for young galaxy ages at all redshifts.

Comparing SFR_{SED} and SFR_{UV+IR} allows us to confirm the reliability of our SED fitting and calibrate our SFR ladder. In practice, we only adopt SFR_{SED} as our best estimate of the SFR when SFR_{UV+IR} is not available, i.e. for sources that lack $24\mu\text{m}$ detections. The good agreement between the two SFR estimates indicates that SFR_{SED} is reliable, even at high redshifts and for the highest SFRs. Sources with high levels of dust (where the reliability of SED-based estimates may break down) will be detected at $24\mu\text{m}$, ensuring our SFR ladder provides a reliable estimate.

APPENDIX B: BAYESIAN MIXTURE MODELLING OF THE DISTRIBUTION OF X-RAY LUMINOSITIES

In this appendix we describe the Bayesian mixture modelling approach that we use to estimate the intrinsic distribution of X-ray luminosities for galaxies in a given stellar mass–redshift bin, $p(\log L_X | \mathcal{M}_*, z)$.

For a single galaxy in our sample (with a given redshift, z_i , and stellar mass, \mathcal{M}_i), the observed X-ray data consist of the total observed counts, N_i , the estimated background count rate, b_i , and the effective exposure at the position of the source, t_i . In a Bayesian framework, our knowledge of the X-ray luminosity of the galaxy, L_X ,¹² is described by a probability distribution,

$$p(L_X | D_i)dL_X = \mathcal{L}(N_i | b_i, t_i, z_i) \pi(L_X | \mathcal{M}_i, z_i)dL_X \quad (\text{B1})$$

¹² Following Aird et al. (2015), L_X refers to the rest-frame 2–10 keV luminosity, regardless of the observed energy band (here, the 0.5–2 keV band).

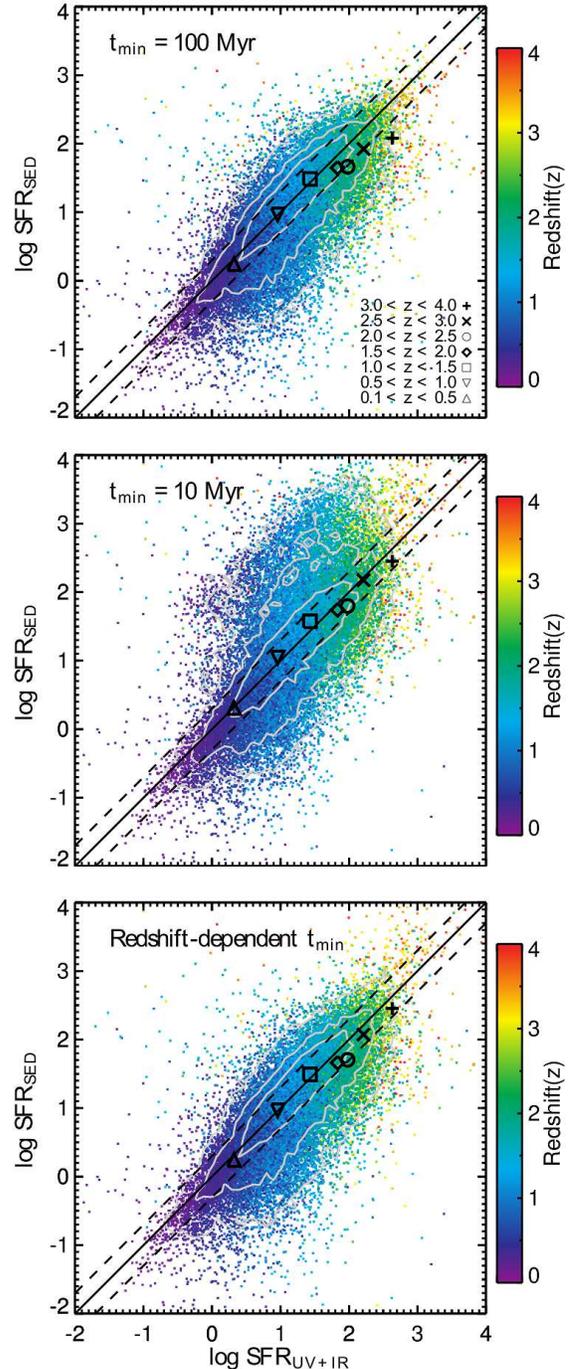


Figure A1. Comparison of UV+IR estimates of the SFR (SFR_{UV+IR}) to SED-based estimates (SFR_{SED}) for our star-forming galaxy sample with different assumptions on the minimum allowed age of the stellar population, t_{min} , in the SED fitting. Small coloured points indicate individual galaxies (with an estimate of SFR_{UV+IR} i.e. detected at $24\mu\text{m}$), colour-coded by redshift. The large black points show the median SFR in a redshift bin, as indicated. Grey contours enclose 50, 80, and 95 per cent of the galaxies. The solid black line indicates a 1:1 relation whereas the dashed black lines indicate ± 0.3 dex (i.e. a factor ~ 2). With a standard $t_{\text{min}} = 100$ Myr (*top panel*) the SED fitting tends to under-estimate the SFR (compared to UV+IR) at the highest redshifts ($z \gtrsim 2$). Allowing for younger ages (e.g. $t_{\text{min}} = 10$ Myr, *middle panel*) improves the SFR estimates for high- z galaxies but leads to over-estimates at lower z , resulting in an increased scatter. Adopting a redshift-dependent t_{min} (*bottom panel*, see Equation A3), whereby very young galaxies are not allowed at lower redshifts, provides a reasonable compromise.

where D_i indicates the observed data from source i , $\mathcal{L}(N_i | L_X, b_i, t_i, z_i)$ is the likelihood of observing N_i counts from a source with luminosity L_X , and $\pi(L_X | \mathcal{M}_i, z_i)$ acts as a prior, describing the true underlying distribution of luminosities of star-forming galaxies with mass \mathcal{M}_i and redshift z_i . The likelihood of observing N_i X-ray counts can be described by a Poisson process, thus,

$$\mathcal{L}(N_i | b_i, t_i, z_i) = \frac{(k_i L_X + b_i)^{N_i}}{N_i!} e^{-(k_i L_X + b_i)} \quad (\text{B2})$$

where

$$k_i = \eta(z_i) t_i \quad (\text{B3})$$

and $\eta(z_i)$ is a factor that converts between L_X and the observed count rate in the 0.5–2 keV band. In this paper, we assume the X-ray emission has a fixed spectral shape that is described by a power-law with photon index $\Gamma = 1.9$ (subject to absorption due to Galactic N_{H} only) and thus the conversion factor depends on redshift only. A correction for the size of the aperture used to extract X-ray counts (which corresponds to a 70 per cent EEF, see Section 2.2 above) is included in this term. We note that Equations B1 and B2 are valid for sources that are associated with significant X-ray detections (above our detection threshold) and for galaxies that are “non-detections” in the X-ray imaging where we have extracted X-ray information at the galaxy position.

As a mathematical convenience, we re-write Equation B2 as

$$\mathcal{L}(N_i | b_i, t_i, z_i) = \sum_{S=0}^{N_i} \left(\frac{(k_i L_X)^S}{S!} e^{-k_i L_X} \frac{b_i^{N_i-S}}{(N_i-S)!} e^{-b_i} \right). \quad (\text{B4})$$

This equation can be interpreted as describing two independent Poisson processes that produce the observed counts: the source of luminosity L_X which produces an integer number of observed counts, S ; and a background component with expected rate b_i that also produces an integer number of observed counts, $B = N_i - S_i$. The summation can then be seen as integrating over the possible values of the unknown nuisance parameter, S . Equation B4 is also mathematically equivalent to the binomial expansion of Equation B2.

The likelihood of observing all sources within a stellar mass–redshift bin is given by integrating Equation B1 over all possible L_X for an individual source and then taking the product of the probabilities for all sources in a bin. Thus,

$$\begin{aligned} \mathcal{L}(\mathbf{D}_{\text{bin}}) &= \prod_{i=1}^{n_{\text{source}}} \int_0^{\infty} p(L_X | D_i) dL_X \quad (\text{B5}) \\ &= \prod_{i=1}^{n_{\text{source}}} \int_0^{\infty} \mathcal{L}(N_i | b_i, t_i, z_i) \pi(L_X | \mathcal{M}_{\text{bin}}, z_{\text{bin}}) dL_X \end{aligned}$$

where \mathbf{D}_{bin} indicates the data from all sources within a bin and n_{source} is the number of sources in the bin. Here, we assume that the intrinsic distribution of L_X is the same throughout the stellar mass–redshift bin and described by $\pi(L_X | \mathcal{M}_{\text{bin}}, z_{\text{bin}})$.

Our next task is to define a model for the intrinsic probability distribution function of luminosities for galaxies of a given stellar mass and redshift,

$$\pi(L_X | \mathcal{M}_{\text{bin}}, z_{\text{bin}}) dL_X \equiv p(\log L_X | \mathcal{M}_*, z) d \log L_X \quad (\text{B6})$$

where we have re-written our target function as $p(\log L_X | \mathcal{M}_*, z)$ to indicate a probability density per unit $\log L_X$ at a given stellar mass and redshift.

We choose to adopt a flexible, non-parametric description of

$p(\log L_X | \mathcal{M}_*, z)$ based on a Bayesian mixture modelling approach. The overall probability function is described by the sum of K mixture components,

$$p(\log L_X | \mathcal{M}_*, z) d \log L_X = \sum_{j=1}^K f(\log L_X | \theta_j) d \log L_X \quad (\text{B7})$$

where $f(\log L_X | \theta_j)$ is a single function that describes each mixture component with a different set of parameters, θ_j . We choose to model our distribution by a series of Gamma distributions as this function is the conjugate of our Poisson likelihood and thus significantly simplifies our computation compared to using e.g. a sum of Gaussian distributions that is more generally adopted in Bayesian mixture modelling (e.g. Kelly et al. 2008). A single Gamma distribution component is given by

$$f(\log L_X | \theta_j) d \log L_X = A_j \frac{\ln 10}{\Gamma(\alpha_j)} \left(\frac{L_X}{L_j} \right)^{\alpha_j} e^{-L_X/L_j} d \log L_X \quad (\text{B8})$$

where α_j controls the shape of the distribution, L_j is a scale parameter (which determines the position of the peak of the distribution), and $\Gamma(\alpha_j)$ is the Gamma function. Each component is normalised such that

$$\int_{-\infty}^{+\infty} f(\log L_X | \theta_j) d \log L_X = A_j \quad (\text{B9})$$

and we required that

$$\sum_{j=1}^K A_j = 1 \quad (\text{B10})$$

and thus each A_j represents the fraction of the overall probability distribution function that is associated with component j of the Bayesian mixture model. We note that a single component of our Gamma distribution is analogous to the Schechter distribution, although we require $\alpha_j > 0$ to ensure that the integral of a component does not diverge; a rising “faint-end slope” is instead described by the combination of multiple components.

With the form of our mixture components defined by Equation B8 and the Poisson likelihood for each source given by Equation B4, our overall likelihood (Equation B5) can be written as

$$\begin{aligned} \mathcal{L}(\mathbf{D}_{\text{bin}}) &= \prod_{i=1}^{n_{\text{source}}} \left[\int_{-\infty}^{+\infty} \sum_{S=0}^{N_i} \left(\frac{(k_i L_X)^S}{S!} e^{-k_i L_X} \frac{b_i^{N_i-S}}{(N_i-S)!} e^{-b_i} \right) \right. \\ &\quad \times \left. \sum_{j=1}^K \left(A_j \frac{\ln 10}{\Gamma(\alpha_j)} \left(\frac{L_X}{L_j} \right)^{\alpha_j} e^{-\frac{L_X}{L_j}} \right) d \log L_X \right] \\ &= \prod_{i=1}^{n_{\text{source}}} \left[\sum_{j=1}^K A_j w_{ij} \right] \quad (\text{B11}) \end{aligned}$$

where

$$\begin{aligned} w_{ij} &= \sum_{S=0}^{N_i} \left[\frac{b_i^{N_i-S}}{(N_i-S)!} e^{-b_i} \frac{\ln 10}{\Gamma(\alpha_j)} \times \right. \\ &\quad \left. \int_{-\infty}^{+\infty} \frac{(k_i L_X)^S}{S!} e^{-k_i L_X} \left(\frac{L_X}{L_j} \right)^{\alpha_j} e^{-\frac{L_X}{L_j}} d \log L_X \right]. \quad (\text{B12}) \end{aligned}$$

The integral in Equation B12 has an analytic solution for a given L_j , α_j and S . To ease our computation (while retaining sufficient flexibility), we thus choose to fix all $\alpha_j = 3.0$ and adopt a fixed, logarithmically spaced grid for L_j spanning $38 \leq$

$\log L_j$ (erg s^{-1}) ≤ 45 in steps of 0.2 dex.¹³ Evaluation of our overall likelihood (Equation B11) then reduces to the summation of a series of pre-computed weights, w_{ij} , multiplied by each A_j that constitute the free parameters to be determined.

The final step in our Bayesian inference is to combine our overall likelihood (Equation B11) with a prior on the free parameters (the set of A_j) to recover the posterior distribution, which in turn describes our estimate of the intrinsic distribution of luminosities, $p(\log L_X | \mathcal{M}_*, z)$. We apply an additional “smoothness” prior (see also Buchner et al. 2015) which requires that the values of A_j do not vary rapidly between adjacent mixture components in our grid of L_j , as well as ensuring that our recovered distribution is not dominated by a small number of the mixture components. We adopt a log-normal prior that penalises large deviations in logarithmic space between adjacent components. Thus,

$$\pi(\log A_{j+1}) \sim \text{Normal}(\log A_j, \sigma) \quad (\text{B13})$$

where we have assumed an ordering of the mixture components in increasing L_j . The level of smoothness is controlled by σ , which we fix at 0.3, providing a reasonable balance between a smoothness requirement and ensuring we are able to recover distinct features. No smoothness prior is applied for A_1 , corresponding to the lowest luminosity component, which is instead allowed to take any value between 0 and 1. The remaining components are subject to the additional constraint that

$$\sum_{j=2}^K A_j = 1 - A_1. \quad (\text{B14})$$

We adopt the MULTINEST algorithm (Feroz & Hobson 2008; Feroz et al. 2009, 2013) to efficiently explore the multi-modal parameter space and produce posterior draws of the set of parameters, $\{A_j\}$. Our final estimates in Figure 3 are derived by evaluating $p(\log L_X | \mathcal{M}_*, z)$ for each posterior draw of $\{A_j\}$ and determining the median (solid lines) and 90 per cent confidence intervals (hatched regions) at a given $\log L_X$.

APPENDIX C: SIMULATIONS TO VERIFY OUR BAYESIAN METHODOLOGY

We have performed a number of simulations to assess the ability of our Bayesian methodology (described above) to recover different underlying distributions and identify a low-luminosity peak (as seen in Figure 3).

Our initial simulations are shown in Figure C1. For these simulations, we start with the sample of 6937 galaxies in our $10.0 < \log \mathcal{M}_*/\mathcal{M}_\odot < 10.5$ and $1.0 < z < 1.5$ bin and select a random subset of $N_{\text{gal}} = 5000, 500$ and 100 galaxies. We adopt the observed values of the redshift, background rate, and effective exposure (z_i, b_i, t_i) for the randomly selected galaxies, ensuring our simulations adopt a realistic distribution for depths of the X-ray data. We then simulate the observed total counts, N_i , from each galaxy, assuming different underlying distributions of $p(\log L_X | \mathcal{M}_*, z)$, as shown by the black dashed lines in Figure C1. Finally, we attempt to recover the underlying distribution

using our Bayesian mixture modelling algorithm for each of the different N_{gal} samples.

Our recovered constraints on $p(\log L_X | \mathcal{M}_*, z)$ for each of the N_{gal} samples are shown by the coloured regions in Figure C1. We are able to recover a wide range of distributions. With a large sample of galaxies ($N_{\text{gal}} = 5000$), our modelling only deviates significantly from the input distributions at the faintest luminosities, where the data provide little constraint and thus our prior for a smoothly varying distribution tends to dominate. Such discrepancies are reasonable and should have a minimal impact on our results based on real data. As expected, with fewer galaxies in a sample (e.g. $N_{\text{gal}} = 500$ or 100) the recovered constraints on $p(\log L_X | \mathcal{M}_*, z)$ become progressively weaker. Furthermore, our prior, which prefers a smoothly varying distribution, has a stronger impact for the smaller sample sizes and thus smooths out the more acute features of the input distribution. This effect is most apparent in Panel b) for the “Two Gaussians” input distribution. We note that the identification of a distinct, narrow peak at lower luminosities in $p(\log L_X | \mathcal{M}_*, z)$, as seen in the real data, tends to go *against* our prior for a smoothly varying distribution, indicating that these features are real and driven by the observed data, rather than the prior.

The final panel (e) of Figure C1 assumes that there is no X-ray emission from any of the galaxies in a given sample and that the data are purely due to Poisson fluctuations of the background. Our recovered posterior distributions for $p(\log L_X | \mathcal{M}_*, z)$ are consistent with a minimal population of X-ray emitting sources. The posterior distributions set upper limits on the fraction of galaxies at a given luminosity. With fewer galaxies in a sample, the posterior upper limits on $p(\log L_X | \mathcal{M}_*, z)$ shift to higher probability densities. For the smallest sample considered ($N_{\text{gal}} = 100$) the upper limit from the “Background only” simulation (shown by the thin, long-dashed blue line, which is replicated in every panel) is roughly consistent with the posterior distributions in panels c) and d), indicating that with such a small sample the posterior constraints are consistent with the data coming from background fluctuations and we are unable to place strong constraints on the underlying distribution.

The ability of our Bayesian modelling to recover a distinct, low-luminosity peak in $p(\log L_X | \mathcal{M}_*, z)$ that can be attributed to the bulk of the population (and thus to star formation processes) in a given sample of galaxies will depend on a number of different factors: i) the size of the galaxy sample, ii) the depth of the X-ray data for such a sample, iii) the extent of AGN activity within the sample, and iv) the true luminosity of any such star-formation peak. To explore these effects in more detail, and thus verify the robustness of the results presented in this paper, we have carried out a large suite of additional simulations. We start with the observed sample of galaxies in a given stellar mass–redshift bin (as shown in Figure 3). We retain the observed X-ray data (total counts, background counts and effective exposure) for any galaxies with significant X-ray detections *and* with an observed $L_X > 10^{41} \text{ erg s}^{-1}$. This ensures we retain a realistic AGN contribution to the observed $p(\log L_X | \mathcal{M}_*, z)$ that will vary with stellar mass and redshift. For the remaining galaxies, we simulate new X-ray data assuming the X-ray luminosity is drawn from a log-Gaussian distribution centred at $L_X^{\text{mode}} = 10^{39} \text{ erg s}^{-1}$ with a width of 0.25 dex.

We then run our Bayesian analysis on the simulated data and attempt to identify the mode of the recovered luminosity distribution. The process is repeated for a grid of input L_X^{mode} (increasing in steps of 0.5 dex) and for every stellar mass–redshift bin. The results are shown in Figure C2, which compares the in-

¹³ An additional mixture component with $\log L_j$ (erg s^{-1}) = 36 is also included to account for the possibility that some fraction of our sources produce essentially no X-ray luminosity (in practice, modelled as a luminosity ~ 2 orders of magnitude lower than the rest of the population and significantly below the limits probed by our data).

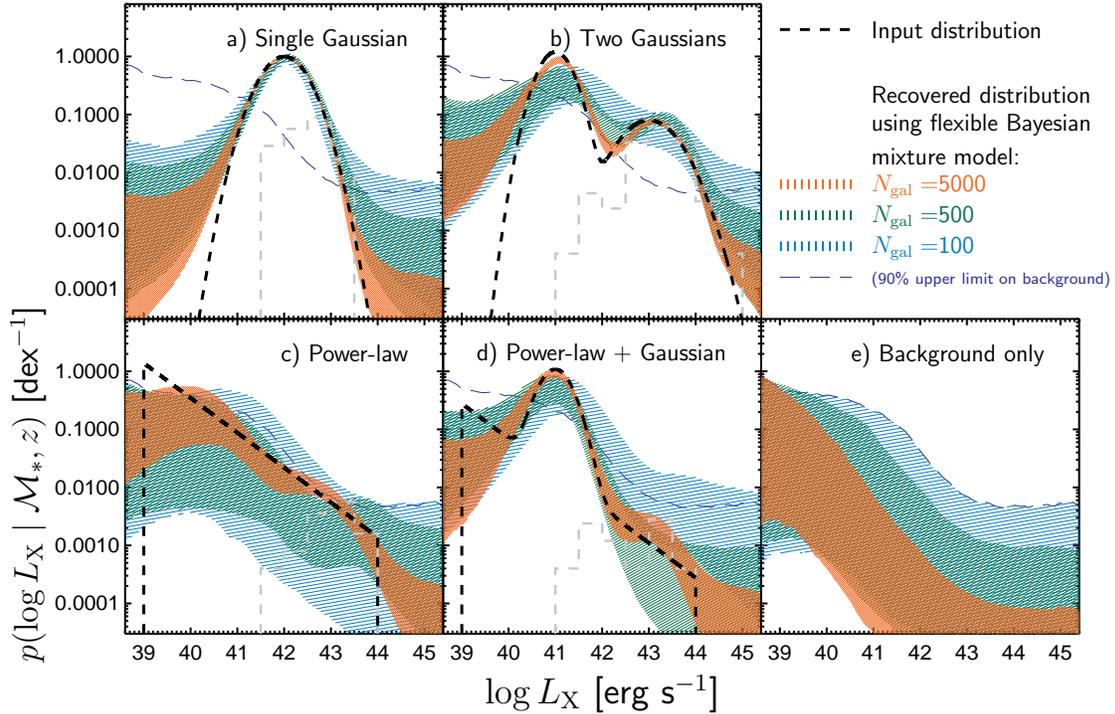


Figure C1. Demonstration of the reconstruction of $p(\log L_X | \mathcal{M}_*, z)$ with our Bayesian mixture modelling (coloured regions, corresponding to different galaxy sample sizes, as indicated) for simulating various underlying distributions (black dashed lines): a) a single Gaussian (log-normal) function; b) two Gaussian (log-normal) distributions centred at different $\log L_X$; c) a power-law distribution over $39 < \log L_X \text{ (erg s}^{-1}\text{)} < 44$; d) the combination of a Gaussian and a power-law; and e) assuming the observed counts are purely from the background for all galaxies in a sample. The grey dashed histogram shows the distribution of sources in our $N_{\text{gal}} = 5000$ simulated sample that would satisfy the nominal X-ray detection limits. Our flexible Bayesian method is able to recover distributions with a range of underlying shapes, correcting for incompleteness and probing below the nominal detection limits. However, with smaller samples of galaxies the constraints become progressively weaker and our prior has a stronger impact, smoothing out features in the underlying distribution. For the smallest sample size considered ($N_{\text{gal}} = 100$), it is difficult to distinguish between our recovered distribution and the upper limits based on the background (shown by the blue long-dashed line in every panel).

put L_X^{mode} in the simulation with that recovered by our Bayesian modelling. We set two conditions for the identification of a significant low-luminosity peak given the posterior constraints on $p(\log L_X | \mathcal{M}_*, z)$ in our simulations:

- (i) The probability density at the mode of the measured $p(\log L_X | \mathcal{M}_*, z)$ lies above the 90 per cent upper limit from the random sampling of the background.
- (ii) Integrating $p(\log L_X | \mathcal{M}_*, z)$ down to 1 dex below the peak accounts for > 50 per cent of galaxies.

The first condition corresponds to requiring that any observed peak cannot be accounted for purely by background fluctuations. The second condition is required for bins with substantial AGN contributions at higher L_X ; in such bins the mode of the recovered distribution can correspond to a broad, high-luminosity peak that is associated with AGN activity, but only accounts for at most ~ 20 per cent of the galaxies. These conditions correspond to our original identification of significant, low-luminosity peaks that account for the bulk of galaxies in the real data. Simulations where the recovered L_X^{mode} does not satisfy these conditions are indicated by grey open symbols in Figure C2.

The results from the simulations shown in Figure C2 illustrate that our Bayesian modelling is able to accurately recover the position of L_X^{mode} over a wide range of luminosities. The simulations reveal a complex relationship between the number of galaxies in a

bin, the level of AGN activity in such galaxies, and our ability to identify a distinct peak in the underlying distribution at different luminosities. In all stellar mass–redshift bins there is a lower limit on the luminosity of the peak that can be recovered (i.e. the lowest luminosity black point) that depends in a fairly complex manner on the number of galaxies in a bin and the level of AGN activity in such galaxies. The errors on L_X^{mode} below this limit (grey points) become very large as the mode either tracks a broad, poorly defined peak at higher luminosities (associated with AGN in a small fraction of the galaxies) or the recovered distribution is consistent with background noise. The position of the observed L_X^{mode} in the real data (indicated by the large orange squares) varies relative to this effective limit but always lies within the range where a reliable measurement is possible. In many stellar mass–redshift bins, it would be possible to measure an L_X^{mode} at ~ 0.5 –1 dex fainter in luminosity than the observed value, providing further evidence that our measurements are robust and not an artefact of the data. These simulations thus verify the reliability of our Bayesian method and the results presented in the paper.

In stellar mass–redshift bins with sufficiently large number of galaxies (i.e. for $\log \mathcal{M}_*/\mathcal{M}_\odot < 10.5$), the simulations show that our method can robustly probe a factor ~ 10 –100 below the nominal point source detection limits of our deepest *Chandra* data, indicated by the dotted blue lines in Figure C2 which correspond to the on-axis flux limit for the ~ 4 Ms data in the GOODS-S field.

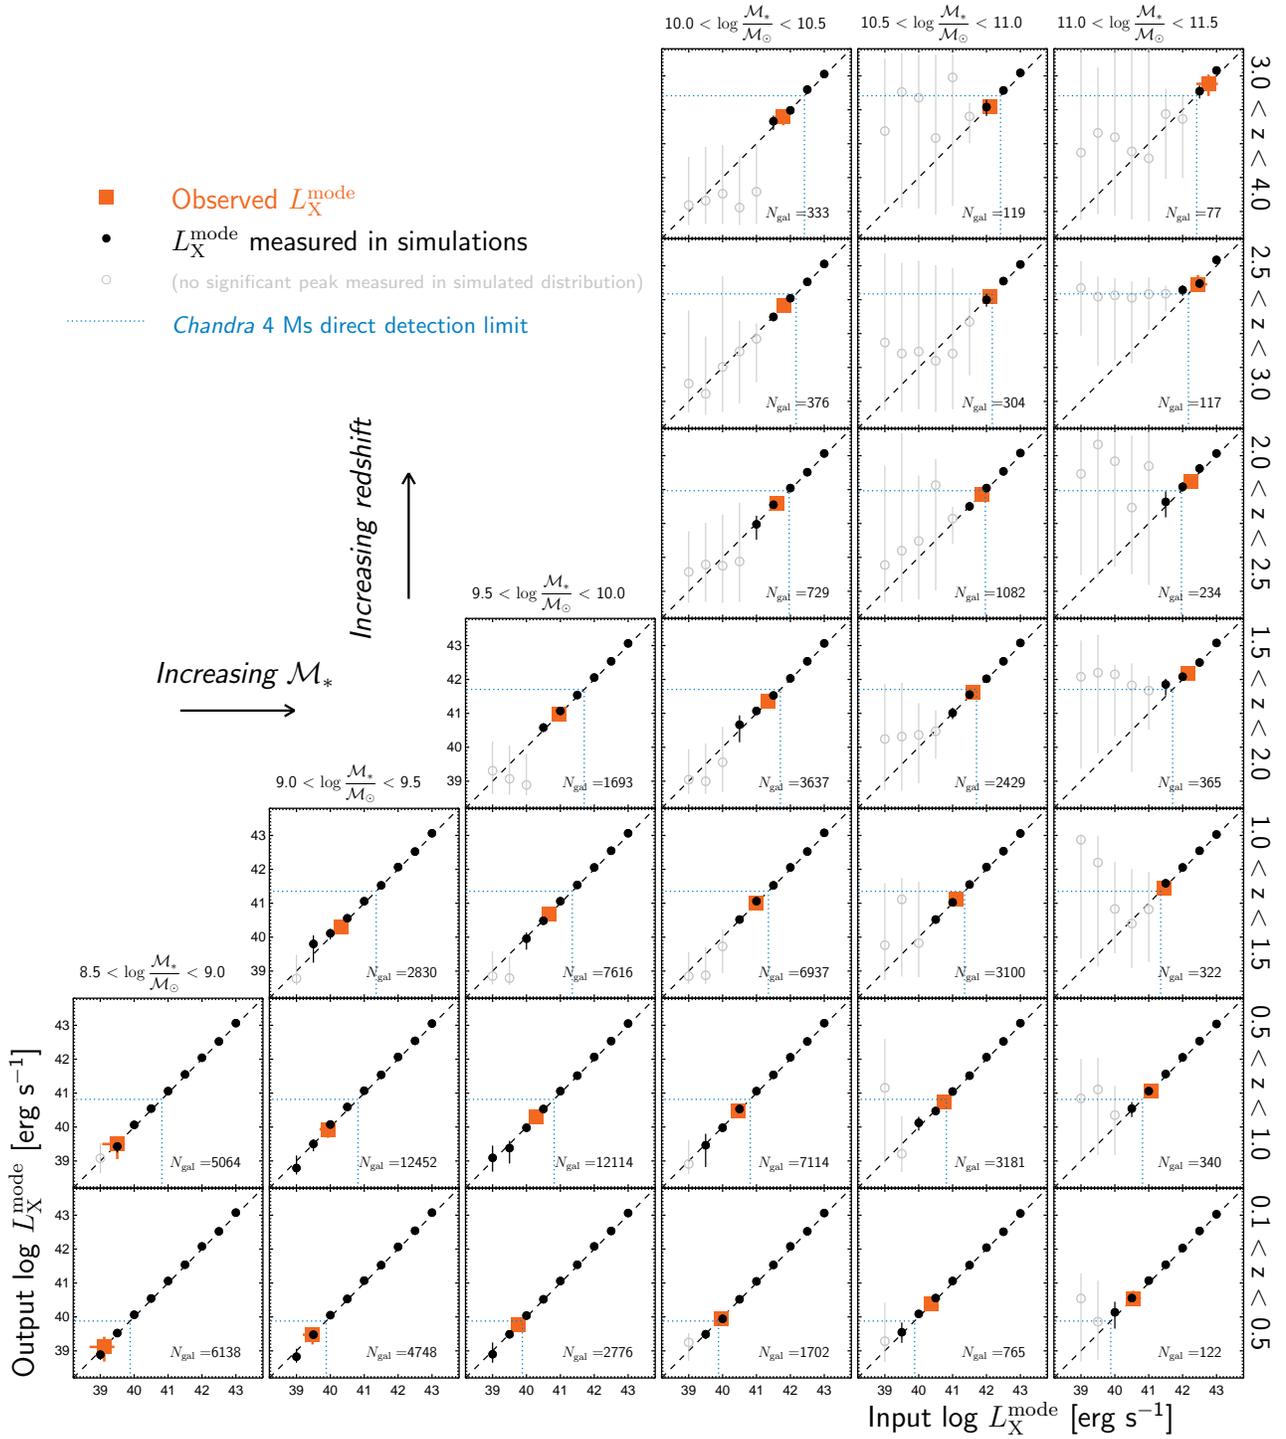


Figure C2. Comparison of the input L_X^{mode} in simulations (x -axis) to that recovered using our Bayesian modelling (y -axis), running simulations based on the observed galaxy samples in each stellar mass–redshift bin for a range of input luminosities (black circles). Grey open symbols indicate bins where the mode of the recovered distribution either lies below the background constraints or only accounts for a small fraction of higher luminosity sources (i.e. AGN) and thus we do not have a significant detection of a low-luminosity peak of X-ray emission from galaxies in our simulated data. Large orange squares show the observed L_X^{mode} in the real data. The blue dotted lines indicate the luminosity corresponding to the flux limit for direct detection of point sources in the field with our deepest *Chandra* data (~ 4 Ms in the GOODS-S field). With sufficiently large galaxy samples we can reliably probe below these limits.

Such sensitivity limits are only achieved over < 1 per cent of our total area coverage. Our method is thus a powerful way of fully exploiting *Chandra* survey data.

APPENDIX D: THE DISTRIBUTION OF X-RAY LUMINOSITIES FOR QUIESCENT GALAXIES AND THE RELATION TO STELLAR MASS

The main results of this paper focus on the X-ray emission from star-forming galaxies, revealing the “X-ray main sequence”. In this appendix, we apply the same analysis to sources that lie in the quiescent¹⁴ region of the UVJ diagram. Figure D1 presents measurements of $p(\log L_X | \mathcal{M}_*, z)$ at moderate-to-high stellar masses ($10 < \log \mathcal{M}_*/\mathcal{M}_\odot < 11.5$) and lower redshifts ($z < 2$) where we are able to place good constraints on the distributions. A low-luminosity peak can be identified in all of these distributions that we associate with galactic processes, as well as a tail to higher luminosities associated with AGN accretion activity. This low-luminosity peak is found despite the apparently quiescent nature of these galaxies (based on their UVJ colours). The X-ray emission is likely to come from a combination of HMXBs (tracing low levels of ongoing star formation that may be obscured at optical wavelengths), LMXBs (a longer lived source population that provide a delayed tracer of star formation activity in the galaxy), and any hot gas component. At lower stellar masses ($\log \mathcal{M}_*/\mathcal{M}_\odot < 10$) or higher redshifts ($z > 2$), not shown in Figure D1, we are unable to identify a clear peak (significantly above the constraints from a sampling of the *Chandra* background) in our measurements of $p(\log L_X | \mathcal{M}_*, z)$, most likely due to the relatively small numbers of quiescent galaxies in our samples and the limited depths probed by our X-ray data.

In Figure D2 we plot the luminosity of the peak (the mode of the distribution) as a function of stellar mass for the quiescent galaxy samples (cf. Figure 4 for star-forming galaxies). The dashed lines show the “X-ray main sequence” for star-forming galaxies identified in Section 4. The dotted lines show the best fit to a linear relationship between L_X^{mode} and \mathcal{M}_* for quiescent galaxies that evolves with redshift, of the form

$$\log L_X^{\text{mode}} [\text{erg s}^{-1}] = a + \left(\log \frac{\mathcal{M}_*}{\mathcal{M}_\odot} - 10.2 \right) + b \log \frac{1+z}{1+z_0} \quad (\text{D1})$$

where we find $a = 40.19 \pm 0.05$ and $b = 3.83 \pm 0.33$. Allowing for a non-linear scaling between L_X^{mode} and \mathcal{M}_* (i.e. a non-unity mass-dependent slope in Equation D1) does not significantly improve the fit. Thus, the relationship between L_X^{mode} and \mathcal{M}_* for quiescent galaxies has a steeper slope than the “X-ray main sequence” for star-forming galaxies. At high stellar masses our measured L_X^{mode} in quiescent galaxies is comparable to the luminosity observed from star-forming galaxies, indicating that for galaxies in this regime there may be a significant contribution from longer-lived LMXBs to the observed X-ray luminosities. This issue is dis-

cussed further in Section 5 above where we explore the scaling between SFR and L_X .

¹⁴ We note that this definition of “quiescent” selects galaxies on the basis of red optical-to-NIR colours. Such sources may not be completely quiescent and have low-levels of ongoing star formation. Furthermore, a non-negligible fraction have relatively high levels of star formation that is identified in the IR (and traced by our SFR estimates). However, such star formation does not have a significant impact on the observed UVJ colours, which are dominated by the light from old stellar populations, especially in the most massive galaxies.

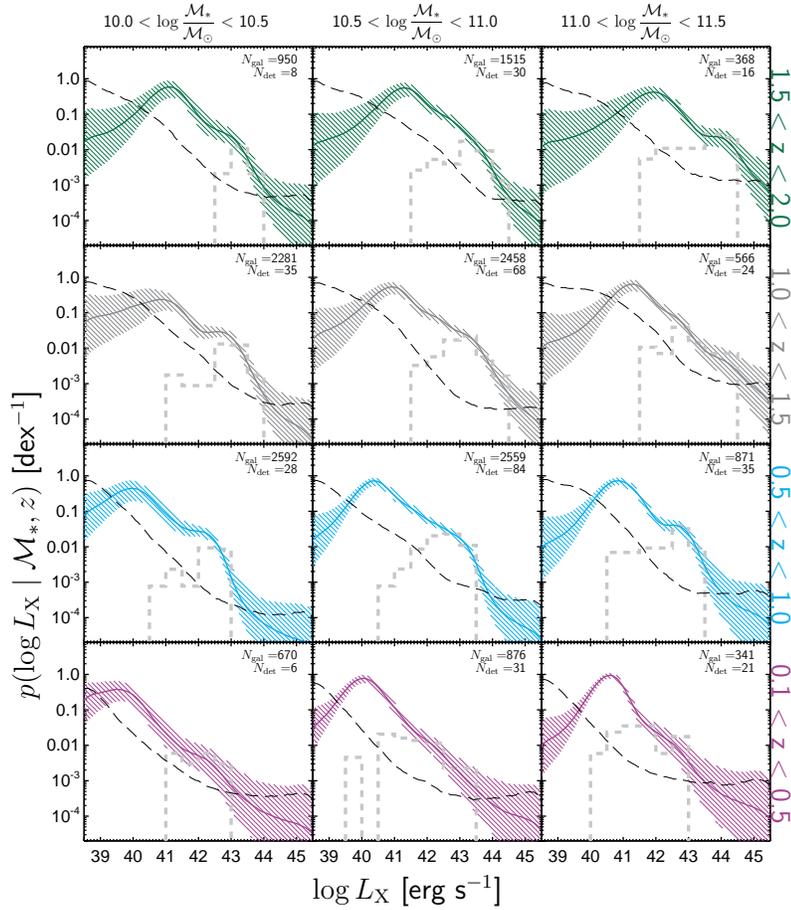


Figure D1. Measurements of the intrinsic probability distributions of X-ray luminosities for samples of quiescent galaxies (selected using the UVJ diagram), as a function of stellar mass and redshift (cf. Figure 3 for star-forming galaxies). We only show results at lower redshifts and higher stellar masses where we are able to place robust constraints on the distribution and identify a peak at low luminosities. As in Figure 3, the thick coloured lines show our best estimates of $p(L_X | M_*, z)$ and shaded regions indicate the 90 per cent confidence interval. The grey histograms correspond to significant detections in the 0.5–2 keV X-ray band (without any corrections for incompleteness). The thin dashed black curve is the 90 per cent upper limit based on shifting the galaxy positions to randomly sample the background.

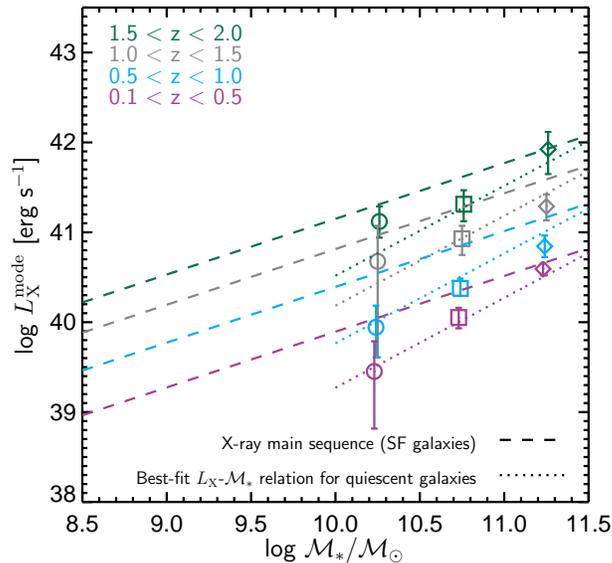


Figure D2. The peak of the distribution of X-ray luminosities from quiescent galaxies, L_X^{mode} , as a function of stellar mass and redshift (data points, dotted lines indicated the best fit to these data using Equation D1), compared to the “X-ray main sequence” for star-forming galaxies (dashed lines) measured in Section 4. At high masses and lower redshifts the average X-ray luminosity from quiescent galaxies is comparable to that observed from star-forming galaxies, indicating longer time-scale processes (e.g. LMXBs) that track the overall stellar mass (rather than the current SFR seen at other wavelengths) can make a significant contribution to the observed X-ray luminosities.