




Constraints on Axion-like Particles from X-Ray Observations of NGC1275

Marcus Berg¹, Joseph P. Conlon², Francesca Day², Nicholas Jennings², Sven Krippendorff² ,
Andrew J. Powell², and Markus Rummel²

¹ Department of Physics, Karlstad University, SE-651 88 Karlstad, Sweden

² Rudolf Peierls Centre for Theoretical Physics, 1 Keble Road, Oxford, OX1 3NP, UK

Received 2017 March 20; revised 2017 August 25; accepted 2017 September 5; published 2017 September 27

Abstract

Axion-like particles (ALPs) can induce localized oscillatory modulations in the spectra of photon sources passing through astrophysical magnetic fields. Ultra-deep *Chandra* observations of the Perseus cluster contain over 5×10^5 counts from the active galactic nucleus (AGN) of the central cluster galaxy NGC1275 and represent a data set of extraordinary quality for ALP searches. We use this data set to search for X-ray spectral irregularities from the AGN. The absence of irregularities at the $\mathcal{O}(30\%)$ level allows us to place leading constraints on the ALP-photon mixing parameter $g_{a\gamma\gamma} \lesssim 1.4\text{--}4.0 \times 10^{-12} \text{ GeV}^{-1}$ for $m_a \lesssim 10^{-12} \text{ eV}$, depending on assumptions on the magnetic field realization along the line of sight.

Key words: elementary particles – X-rays: galaxies: clusters – X-rays: individual (NGC1275)

1. Introduction

The Perseus galaxy cluster (A426) is the brightest X-ray cluster in the sky. It is a cool-core cluster at redshift $z = 0.0176$, centered around the Seyfert galaxy NGC1275 and its active galactic nucleus (AGN). Due to its proximity and brightness, the Perseus cluster has been a standard target for all X-ray satellites. The X-ray spectrum and emission specifically from the AGN are described in Churazov et al. (2003), Balmaverde et al. (2006), Yamazaki et al. (2013), and Fabian et al. (2015).

Axion-like particles (ALPs) are a well-motivated extension of the Standard Model, and arise generically in string compactifications (for example, see Conlon 2006; Svrcek & Witten 2006; Cicoli et al. 2012). As the potential and interactions of ALPs are protected by shift symmetries, they can naturally have extremely small masses or even be massless.³ If it exists, an ALP a interacts with the Standard Model via a coupling to electromagnetism,

$$\frac{a}{M} \mathbf{E} \cdot \mathbf{B}, \quad (1)$$

where $M \equiv g_{a\gamma\gamma}^{-1}$ parameterizes the strength of the interaction and \mathbf{E} and \mathbf{B} are the electric and magnetic fields. A general review of ALPs and their physics can be found in Ringwald (2012). We concern ourselves in this paper with the case $m_a \lesssim 10^{-12} \text{ eV}$, for which nonobservation of gamma-ray photons coincident with the SN1987A neutrino burst constrains $M \gtrsim 2 \cdot 10^{11} \text{ GeV}$ (Brockway et al. 1996; Grifols et al. 1996; Payez et al. 2015).

If ALPs exist, then the interaction of Equation (1) causes ALPs and photons to interconvert in the presence of a background magnetic field $\langle B \rangle$ (Sikivie 1983; Raffelt & Stodolsky 1988). Starting with a pure photon spectrum, this $\gamma \leftrightarrow a$ interconversion results in modulations in the spectrum of arriving photons.

Galaxy clusters are particularly efficient photon-ALP converters (Burrage et al. 2009; Conlon & Marsh 2013), and for

the electron densities and magnetic fields present within galaxy clusters, it is a result that at X-ray energies the $\gamma \leftrightarrow a$ conversion probability is both energy-dependent and quasi-sinusoidal (Conlon & Marsh 2013; Wouters & Brun 2013; Angus et al. 2014; Conlon et al. 2016). Compared to the source spectrum, the spectrum of the arriving photons has oscillatory modulations imprinted on it. By searching for such modulations, we can place constraints on the coupling parameter M .

For this purpose, quasars or AGNs that are either behind or embedded in galaxy clusters provide attractive sources. The original photon spectrum is reasonably well described by an absorbed power law, and all photons arise from a single sight line passing through the cluster. As bright sources, AGNs can also provide the large number of counts necessary for statistical leverage in searching for oscillatory modulations of the photon spectrum.

This method is indeed already largely described and was used in Wouters & Brun (2013; see also Wouters & Brun 2012; Brax et al. 2015; and Collaboration 2016 for a recent analysis of NGC1275 in gamma rays using this approach). However, Wouters & Brun (2013) only applied these ideas to the study of the AGN at the center of the Hydra A galaxy cluster (redshift $z = 0.052$), for which 200 ks of *Chandra* observation time exist. But in this case, the combination of the intrinsic AGN brightness and the redshift of $z = 0.052$ results in only a few thousand counts in total, limiting the ability to produce bounds.

In contrast, the AGN at the center of NGC1275 is both exceedingly bright and the subject of enormous observational time. NGC1275 has been observed for 1 Ms by *Chandra* ACIS-S, with observations taken in 2002 and 2004, and for a further 0.5 Ms by *Chandra* ACIS-I in 2009. Taken together, these generate over half a million X-ray counts from the central AGN—a data set of extraordinary quality for searching for spectral irregularities. Furthermore, there are also 180 ks of observation time with *XMM-Newton* taken in 2001 and 2006. This data set is not as rich as the *Chandra* data set but allows us to cross-check our analysis with a different instrument. In this paper, we use these data sets to search for and constrain ALPs.

The paper is organized as follows. Section 2 provides further details on the physics of ALPs and the attractiveness of bright quasars or AGNs for searching for ALPs. Section 3 describes

³ The shift symmetry also implies that, even if massless, ALPs are not constrained by searches for fifth forces or modifications of general relativity.

the *Chandra* and *XMM-Newton* observations that we have used and their properties. Section 4 describes the *Chandra* analysis, the effects of pileup, and strategies taken to mitigate pileup. Section 5 describes the analysis of the *XMM-Newton* data. In Section 6, we describe bounds on the ALP-photon coupling, and we conclude in Section 7. In Appendix B we discuss irregularities in the data at the 10% level and mention possible instrumental, astrophysical, or new physics explanations.

2. AGNs and ALP Physics

If ALPs exist, they interconvert with photons in a background magnetic field $\langle B \rangle \neq 0$ (Sikivie 1983, 1985; Raffelt & Stodolsky 1988). This occurs because the magnetic field generates, via Equation (1), a two-particle interaction between the ALP a and the photon γ , resulting in a mixing of the mass eigenstates. In a mathematically identical fashion to neutrino oscillations, photons and ALPs then have a finite probability of interconversion as they pass through the magnetic field. This is a quantum-mechanical effect and results in an evolution of the quantum state

$$|\psi_{\text{init}}\rangle = |\gamma(E)\rangle \longrightarrow |\psi_{\text{final}}\rangle = \alpha|\gamma(E)\rangle + \beta|a(E)\rangle,$$

where $|\alpha|^2 + |\beta|^2 = 1$ and $\gamma(E)[a(E)]$ denotes a photon (ALP) with energy E . The conversion probability $|\beta|^2$ depends on the free electron density (which sets the effective photon mass), the magnetic field (which sets the strength of the mixing), and the magnetic field coherence length (which determines the region over which mixing applies).

The calculational details of this are standard. Following the original papers (Sikivie 1983, 1985; Raffelt & Stodolsky 1988), the physics of ALPs passing through magnetic fields has been described in many works. An incomplete list of articles studying aspects of photon-ALP interconversion in astrophysical magnetic fields includes Brockway et al. (1996), Grifols et al. (1996), Burrage et al. (2009), Mirizzi & Montanino (2009), Fairbairn et al. (2011), Tavecchio et al. (2012), Payez et al. (2012), Horns et al. (2012), Meyer et al. (2013), Wouters & Brun (2013), Conlon & Marsh (2013), Angus et al. (2014), Payez et al. (2015), Powell (2015), Dobrynina et al. (2015), Schlederer & Sigl (2016), and Conlon et al. (2016).

In this respect, galaxy clusters at X-ray energies sit at a sweet spot for photon-ALP physics. This is due to two key results. First, galaxy clusters are particularly efficient environments for photon-ALP interconversion. The electron densities are relatively low. Clusters have magnetic fields that are not significantly smaller than in galaxies but in which the \mathbf{B} field extends over megaparsec scales, far greater than the tens of kiloparsecs applicable for galactic magnetic fields. The magnetic field coherence lengths in clusters are also larger than those in galaxies, comfortably reaching tens of kiloparsecs. For massless ALPs, this feature singles out galaxy clusters as providing the most suitable environment in the universe for ALP-photon interconversion.⁴

For convenience, we will restrict our work to massless ALPs in this paper. As the efficiency of ALP-photon conversion depends on $|m_a^2 - \omega_{pl}^2|^{-2}$, where ω_{pl} is the plasma frequency in galaxy clusters, in practice, all masses $m_a \lesssim 10^{-12}$ eV are equivalent for our purposes. For $m_a \sim 10^{-12}$ eV, resonant

conversion can occur, allowing weaker ALP-photon couplings to be probed, but we will not consider that case explicitly in this paper.

The second key result is that, for the electron densities and magnetic field structures present within galaxy clusters, the photon-ALP conversion probability is energy-dependent, with a quasi-sinusoidal oscillatory structure at X-ray energies. This provides distinctive spectral features to search for. We illustrate this in Figure 1, where we plot a typical photon survival probability as a function of energy along a single line of sight modeled on that from NGC1275 to us.

The precise form of the survival probability is not predictable. It depends on the actual magnetic field structure along the line of sight, and so differs for each line of sight. Faraday rotation measures can give statistical information about the strength and extent of the magnetic fields, as for the Coma cluster in Bonafede et al. (2010), but the actual magnetic field along any one line of sight is unknown. However, the form shown in Figure 1—a quasi-sinusoidal structure with a period that increases with energy—is generic and arises for any reasonable choice of central magnetic field value or range of coherence lengths. The inefficiency of conversion at energies $E \lesssim 0.2$ keV is also generic, implying that the effects of photon-ALP conversion are not visible in the optical (and below) range.

Figure 1 (right) also shows the same survival probability convolved with a Gaussian with a full width at half maximum (FWHM) of 150 eV, representing the approximate energy resolution of the CCD detectors present on the *Chandra* and *XMM-Newton* satellites (the precise figure of 150 eV is taken from the in-orbit performance of the ACIS-I detectors on *Chandra*; see table 6.4 of the *Chandra* proposer’s guide⁵). While at the lowest energies the oscillations are too rapid to be resolved by CCD detectors and would require microcalorimeters such as those that were present on *Hitomi*,⁶ in general it is fortuitous that the scales of the oscillations match those of the X-ray telescopes extensively used to observe galaxy clusters.

If ALPs exist, then, for photons arriving from a single location, this conversion imprints a particular quasi-sinusoidal modulation on the actual photon spectrum. There is also an overall reduction in luminosity, but this can be absorbed into the overall normalization of the spectrum. For unpolarized light, the $\gamma \rightarrow a$ conversion probability cannot exceed 50% and, in the limit of strong coupling, saturates at an average value of $\langle P(\gamma \rightarrow a) \rangle = 1/3$ (for example, see Wouters & Brun 2012). It therefore follows that, expressed as a ratio of data to model, the maximal allowed range of ALP-induced modulations is approximately $\pm 30\%$.⁷

There are three effects that can wash out these modulations. The first is the finite energy resolution of the telescope; as shown in Figure 1, this removes any structure present at the lowest energies. The second is when emission arises from an extended source involving many different lines of sight. In this case, the peaks and troughs from different lines of sight undergo destructive interference, reducing any signal. The third

⁵ <http://cxc.harvard.edu/proposer/POG/html/chap6.html>

⁶ The energy resolution from *Hitomi* was around ~ 5 eV.

⁷ We reemphasize here that photon-ALP conversion involves quantum oscillations between states rather than absorption. Therefore, for passage from $A \rightarrow C$, the survival probability $P(A \rightarrow C)$ does not equal $P(A \rightarrow B) \times P(B \rightarrow C)$.

⁴ Although they appear appealing, magnetars and related objects do not provide efficient environments for ALP-photon conversion (Raffelt & Stodolsky 1988).




Figure 1. Left: photon survival probability along a line of sight modeled on that from NGC1275 to us for a randomly generated magnetic field. A central magnetic field of $B_0 = 25\mu\text{G}$ was assumed, with a radial scaling of $\langle B(r) \rangle \sim n_e(r)^{0.7}$. There were 200 domains, with lengths drawn randomly between 3.5 and 10 kpc, and the total propagation length was 1200 kpc. The ALP-photon coupling is $g_{a\gamma\gamma} = 1.5 \times 10^{-12} \text{ GeV}^{-1}$ (roughly a factor of three beyond the current upper limit $g_{a\gamma\gamma} < 5 \times 10^{-12} \text{ GeV}^{-1}$ from SN1987A). One should not take too seriously the particular details of the magnetic field model used, as there is no observational data to constrain the precise range of magnetic field coherence lengths in the Perseus cluster, but the quasi-sinusoidal structure arises generically. Right: photon survival probability for the same magnetic field convolved with a Gaussian with a FWHM of 150 eV.

is insufficient photon statistics, when localized oscillations become indistinguishable from Poisson fluctuations.

Bright pointlike sources either behind or embedded in a galaxy cluster are particularly attractive for searching for ALP-induced modulations. The galaxy cluster provides a good environment for ALP-photon conversion; the bright point source ensures that there are many photons, all passing along the same line of sight.

These factors make quasar or AGN spectra attractive for searching for ALPs. Emission from an AGN arises from matter accreting onto the central black hole. As evidenced by the rapid time variability of AGN luminosities, the physical region sourcing the X-ray AGN emission is tiny—of order a few Schwarzschild radii of the central black hole. As cluster magnetic fields are ordered on kiloparsec scales, this implies that, for all practical purposes, every photon arising from the AGN has experienced an identical magnetic field structure during its passage to us.

To first approximation, at X-ray energies an AGN spectrum can be described as an absorbed power law. The effect of ALPs is then to imprint a quasi-sinusoidal modulation on this power law of relative amplitude at most $\mathcal{O}(30\%)$ and with a modulation period of order a few hundred eV. As the fractional Poisson error on N counts is $\frac{1}{\sqrt{N}}$ and CCD detectors such as those on *Chandra* and *XMM-Newton* have intrinsic energy resolutions of around $\mathcal{O}(100 \text{ eV})$, it therefore requires large numbers of counts to be able to distinguish any ALP-induced modulations from normal statistical fluctuations.

All of the above facts make the AGN of the Seyfert galaxy NGC1275 an excellent candidate for searching for ALP-photon interconversion. NGC1275 is the central galaxy of the Perseus cluster, which, as a cool-core cluster, should have a high central magnetic field (estimated as $25\mu\text{G}$ in Taylor et al. 2006)—implying that the sight line from NGC1275 to us should be efficient at ALP-photon conversion.

The nucleus is intrinsically bright and unobscured, with a spectrum that is well characterized by a power law and narrow Fe $K\alpha$ line, absorbed by the galactic n_{H} column density (Churazov et al. 2003). Furthermore, there is enormous *Chandra* observation time on NGC1275, encompassing 1.5 Ms in total. This results in over half a million photon counts

originating from the central AGN, although, for the on-axis observations, quite a number are contaminated by pileup. This is a level two orders of magnitude larger than that of the study in Wouters & Brun (2013) involving Hydra A.

3. The Observations

3.1. Chandra

The deep *Chandra* observations involving NGC1275 can be divided into three main groups. The first involves 200 ks of ACIS-S observations taken in 2002 together with 800 ks of ACIS-S observations taken in 2004.⁸ In these observations, NGC1275 is close to the aimpoint. The second group involves 300 ks of ACIS-I observations carried out in 2009, where NGC1275 is approximately midway between the edge of the chips and the aimpoint. The third group involves 200 ks of ACIS-I observations also taken in 2009, in which NGC1275 is close to the edge of one of the chips, around 8 arcmin from the aimpoint. Finally, there are also some brief pre-2002 observations that we do not include. A list of the *Chandra* observations used in this paper can be found in Table 1.

The relevance of this classification is that the point-spread function of the telescope grows off-axis. In the first group, the photons from the AGN suffer little dispersion and are highly concentrated on a few pixels. In the third group, the arriving photons are scattered over many pixels, whereas the second group is intermediate.

The consequence of this is that, despite having the shortest observation time, it is the third group (the 200 ks in which NGC1275 is at the edge of the detector) that provides the cleanest data set. In the first case, the superb optics of *Chandra* work against it; almost all of the photons are concentrated on a few pixels, and these central pixels are highly contaminated by pileup. In this case, a clean spectrum can only be obtained by extracting from the wings of the point-spread function, which reduces the photon count. In the last case, however, being highly off-axis causes degradation of the optics sufficient that

⁸ Note for particle theorists: the ACIS instrument has two main modes, ACIS-S and ACIS-I. One ACIS-S chip leads to an 8 arcminute by 8 arcminute field of view, while ACIS-I will result in a 16 arcminute by 16 arcminute field of view.




Figure 2. NGC1275 in three types of observation (left to right): the center of the chip in 2004 (*Chandra* ObsID 4952), midway between the edge and center of the chip in 2009 (ObsID 11714), and the edge of the chip in 2009 (ObsID 11713). The color coding is adjusted to account for the different observation times such that each color corresponds to the same count rate across images.

the arriving photons are scattered over many pixels, resulting in greatly reduced pileup. The second grouping is intermediary in quality between these two. This is illustrated in Figure 2, which shows images of NGC1275 for each of the different observation types.

Another relevant factor is that the brightness of the AGN varies substantially with time. As described in Fabian et al. (2015), the NGC1275 AGN was brightest from 1970 to 1990 before rapidly declining by an order of magnitude until around 2000. Since then, it appears to have brightened significantly from 2003 onward, although it has not yet returned to the luminosities it had pre-1990.

The combination of pileup and the intrinsic brightening of the AGN implies that, despite the shorter observation times, each of the stacked 200 and 300 ks ACIS-I edge and midway observations from 2009 has more counts than the 1 Ms of ACIS-S observations taken in 2002 and 2004. Furthermore, when we consider cleaned spectra that exclude regions of high pileup, the total count in the ACIS-I edge spectra is more than in the ACIS-I midway and ACIS-S observations put together. These factors make the ACIS-I edge observations optimal for analysis purposes, even though they involve the shortest observational time.

3.2. *XMM-Newton*

There are two observations of NGC1275 with significant exposure time; see Table 2. The first one, 0085110101, was taken in 2001 when the AGN had its lowest emissivity in observational history. The second one, 0305780101, was taken when the emissivity was still relatively low but nevertheless almost twice as bright as in 2001. NGC1275 is on-axis in both observations.

The 2001 observation was taken in full-frame mode, while the 2006 observation was taken in extended full-frame mode. This affects the frame time of the pn camera, which is 73.4 ms in full-frame mode and 199.1 ms in extended full-frame mode. For the MOS camera, the frame time is 2.6 s. The pixel size is $4''.1$ for the pn and $1''.1$ for MOS. This means that the pn in extended full-frame mode is more susceptible to pileup than MOS. In general, pileup is an issue for all *XMM-Newton* observations of NGC1275, as we will discuss in the *XMM-Newton* analysis in Section 5.

A significant difference between *XMM-Newton* and *Chandra* is the angular resolution. For both MOS and pn, the radius of the disk containing 50% of the photons collected in the focal

Table 1
The *Chandra* Observations Used in This Paper

ObsID	Exposure (ks)	Year	Instrument	Location of NGC1275
3209	95.77	2002	ACIS-S	Central
3404	5.31	2002	ACIS-S	Central
4289	95.41	2002	ACIS-S	Central
4946	23.66	2004	ACIS-S	Central
4947	29.79	2004	ACIS-S	Central
6139	56.43	2004	ACIS-S	Central
6145	85	2004	ACIS-S	Central
4948	118.61	2004	ACIS-S	Central
4949	29.38	2004	ACIS-S	Central
6146	47.13	2004	ACIS-S	Central
4950	96.92	2004	ACIS-S	Central
4951	96.12	2004	ACIS-S	Central
4952	164.24	2004	ACIS-S	Central
4953	30.08	2004	ACIS-S	Central
11715	73.36	2009	ACIS-I	Midway
11716	39.64	2009	ACIS-I	Midway
12037	84.63	2009	ACIS-I	Midway
11714	91.99	2009	ACIS-I	Midway
11713	112.24	2009	ACIS-I	Edge
12025	17.93	2009	ACIS-I	Edge
12033	18.89	2009	ACIS-I	Edge
12036	47.92	2009	ACIS-I	Edge

Note. The last column shows the location of NGC1275 in the respective observation relative to the focal point.

Table 2
The *XMM-Newton* Observations Used in This Paper

ObsID	Exposure (ks)	Year	Instrument	Location of NGC1275
0085110101	53.08	2001	EPIC	Central
0305780101	123.3	2006	EPIC	Central

Note. The last column shows the location of NGC1275 in the respective observation relative to the focal point.

plane (half energy width) is around $8''.5$ at 1.5 keV, while for *Chandra* it is much smaller at $\sim 0''.5$ arcseconds. As the central region of the Perseus cluster is also intrinsically bright, this makes it harder to separate AGN and cluster emission for the *XMM-Newton* observations.

The effective area at 1 keV is 922 cm² for MOS and 1227 cm² for the pn, compared to 340 cm² for *Chandra*. This allows the MOS and pn combined to collect roughly 7 times more photons in a given observation time than *Chandra* (although this also increases the amount of pileup). However, there is roughly 10 times more exposure time in the *Chandra* data set, and the significantly better angular resolution of *Chandra* allows a much better contrast to be attained between the AGN and the cluster emission. This leads to much better statistics in the *Chandra* data set in terms of the total number of counts after background subtraction.

4. *Chandra* Analysis

We use CIAO 4.7 (Sherpa Freeman et al. 2001 Fruscione et al. 2006) and HEASOFT 6.17 for the *Chandra* data analysis.⁹ After the data is reprocessed using CALDB 4.6.9, it is cleaned from time periods that are polluted by flares using the program *chips*. We find that only ObsID 4950 is affected by flares, and the cleaning reduces the observation time slightly from 96.12 to 89.23 ks for this observation.

4.1. Pileup: General Comments

As the brightness of the AGN makes pileup a significant feature of all observations, we will make some general comments on it here.

The energy recorded on the individual ACIS pixels (each approximately 0.5 arcsec square) is read out approximately every 3 s (one frame time) in these observations. Based on groupings of 3×3 pixels, events are graded. Pileup refers to the arrival of more than one photon in this grouping within the same readout frame. This can lead to the energy of the two (or more) incident photons being summed and either treated incorrectly as a single photon event of higher energy or assigned a bad grade (grade migration). For an on-axis bright source (as in the ACIS-S observations of NGC1275), the level of pileup can be high, and the resulting spectrum contains events with two, three, and more photons. As pileup is a statistical feature of the number of arriving photons, some level of pileup is inevitable in any observation. The question is always whether the magnitude of pileup is sufficient to corrupt the science analysis being undertaken.

In terms of the measured photon distribution, the general effect of pileup is to cause a hardening of the spectrum: two or more lower-energy photons are misidentified as a single higher-energy photon. This implies that, for a fit of a single power law to a photon distribution, as pileup increases, the best-fit power-law index will decrease. In a spectrum contaminated by pileup, this makes it harder to determine the correct original power-law index.

What about searches for and constraints on ALPs? As we have seen in Section 2, the distinct signal of ALPs is a quasi-sinusoidal modulation in the spectrum—a local excess or deficit in the photon count rate compared to the nearby continuum. For sufficiently small levels of pileup, localized modulations will remain localized modulations, as an overall global continuum redistribution of photons is unable to create or remove localized spikes (or dips) relative to the continuum.¹⁰ However, for sufficiently heavy pileup, the majority

of photons will be redistributed, and such local features will be lost.

While pileup is always a contaminant on the spectrum, what this implies is that a search for localized spectral irregularities is more robust against pileup than, for example, a measurement of the overall power-law spectral index of a source. While some movement of counts from low to high energies will reduce the number of low-energy photons compared to high-energy photons, it will be less likely to affect the presence or absence of sharp localized features.

This robustness is more applicable at low energies. The effective area of the *Chandra* telescopes starts falling rapidly above around 5 keV.¹¹ As at higher energies, a power-law distribution also produces intrinsically fewer photons; it only requires a small amount of pileup of lower-energy photons in the $E > 5$ keV region to cause a significant distortion of the spectrum there.

In contrast, at lower energies, the effective area is larger and there are far more photons, so small amounts of pileup will not affect any spectral features. For observations toward Perseus, there is an additional benefit: the high galactic absorbing column density ($n_H = 1.5 \times 10^{21}$ cm⁻²) removes the lowest-energy photons, resulting in an effective minimal value for a piled-up energy $E_1 + E_2$ (this will be relevant when we discuss in detail below a feature at 2–2.2 keV).

There is an important caveat to this that requires careful treatment. In the presence of rapid variations in the effective area, a failure to account for pileup can result in significant spectral distortion. This is because the distribution of pileup photons around $E_1 + E_2$ does not respect the behavior of the effective area at $E_1 + E_2$ but rather only the effective areas near E_1 and E_2 . If the effective area is rapidly varying near $E_1 + E_2$ and pileup photons make up a large fraction of the total photon count at $E_1 + E_2$, then a fitted spectrum will introduce unphysical features at $E_1 + E_2$.

In the case of *Chandra*, there is a sharp dip in the effective area at around 2 keV arising from an iridium edge, and the effective area also falls off rapidly for energies above 5 keV. Between 2.5 and 5 keV, the effective area is relatively constant and smoothly varying. With *XMM-Newton*, there is also significant structure in the effective area up to around 2.4 keV, but beyond that it is smooth and slowly varying up to 7 keV.

The level of pileup differs greatly between the three observational groups. In the ACIS-S observations, the central pixels are very heavily affected by pileup. As an illustration, for an extraction of the ACIS-S spectrum with no exclusion of a central region, over 5% of the photon counts have energies in the 10–15 keV range. As the effective area of the telescope is zero at these energies, these counts all arise from piled-up events in which several photons have arrived in the same readout frame. In contrast, pileup is far more moderate for the ACIS-I edge observations (for which a similar extraction across the entirety of the image results in only 0.1% of the counts lying in the 10–15 keV range).

Given the data, there are two basic methods we can employ to reduce the effects of pileup on the analysis. The first is to reduce the amount of pileup by using an annular extraction region and extracting the spectrum only from the wings of the point-spread function. The point-spread function of the

⁹ The update to CIAO 4.8 affects data taken in continuous-clocking mode, which does not apply to these observations.

¹⁰ In the case of a strong low-energy emission line, this is not true, as its pileup may result in spikes at integer multiples of the original line.

¹¹ cf. cxc.harvard.edu/proposer/POG/html/ACIS.html.

telescope causes the arriving photons to be spread out on the detector, with the degree of spread increasing as one moves progressively off-axis. Furthermore, the point-spread function is mildly energy-dependent: photons with higher energy are spread out further than photons with smaller energy. While fewer photons arrive in the wings of the point-spread function, those that do suffer less from pileup than those arriving on the central pixels. The second is to model the pileup, and the distortion it induces on the spectrum, explicitly. We shall utilize both approaches below.

The overall balance here is between enhanced photon statistics—but with more pileup and so with worse data quality—and fewer photons but better data. In particle physics language, this is the trade-off between efficiency and purity.

4.2. Fitting Accuracy and Bounds

Before we describe the data in detail, we comment on the relationship between the accuracy of the fit and the bounds on ALP parameters that we are able to produce. This is relevant, as the data involves a bright point source at the middle of a bright and complex continuum background (the center of the Perseus cluster).

Constraints on ALPs arise because they induce modulations that would make the conventional fit (an absorbed power law) a bad one. The constraints on ALPs come from requiring that the badness-of-fit that ALPs would induce is no worse than that present in the actual data—a coupling $g_{a\gamma\gamma}$ is excluded when it would lead to a fit that is worse than actually occurs. To place bounds on ALPs in this way, it is not necessary that the actual data be a perfect fit. It is true that the better the fit of the actual data to an absorbed power law, the stronger the bounds that can be placed on ALPs. However, even if the actual data is only, for example, within 10% of the standard fit, then this still implies the exclusion of ALP couplings large enough to give 30% modulations in the data.

This method of producing bounds is conservative. The simulated data we use when considering the effects of ALPs is more idealized than for real data. In particular, the simulated data do not account for any possible observational or instrumental inaccuracies arising from the complexity of the source, such as pileup from a high photon flux or possible mis-subtraction of the cluster background when obtaining the AGN spectrum.

By contrast, in the actual data, there are various observational, instrumental, or astrophysical effects that could cause the standard fit to fail to be a perfect description of the measured data, even without any exotic physics. These include the following.

1. As mentioned above, pileup can lead to photons being registered at the wrong energy. This can lead to fake excesses at the location of effective area dips.
2. The X-ray emission at the center of the Perseus cluster has an intensity that varies spatially away from the nucleus. As the background regions differ from the source region used, the process of background subtraction may lead to residual cluster emission being present in addition to the AGN spectrum.
3. The response files may not deal fully accurately with cases where the center of the image has significant pileup.

We will exclude ALP couplings for which idealized simulated data are still a worse fit compared to the noisier

actual data. “Noise” in the actual data reduces the bounds that can be placed on ALPs (as small genuine ALP-induced modulations could hide under the noise) but still robustly allows the exclusion of large ALP-induced modulations. The bounds that we place here are therefore conservative but have the potential of further improvement from better modeling of the actual data or through observing NGC1275 with cleaner settings (such as a shorter readout time).

4.3. Spectral Analysis

We first extract the spectrum for NGC1275 for the ACIS-I edge observations (11713, 12025, 12033, and 12036), where the AGN is around 7 arcmin off the optical axis, using the full extraction region without any exclusion of the central core. This is only valid for the edge observations where pileup is relatively low throughout the whole image. Spectra and responses were created using `specextract` for each observation and then stacked using `combine_spectra`. An ellipse around NGC1275 of radii $11''.6$ and $7''.2$ was used for the extraction region. The boundaries of the ellipse were set by the location where the image of the AGN ceased to dominate the background cluster emission. The background was taken from an elliptical annulus around NGC1275, with the outer radii being 19 and 13 arcsec and the inner radii $13''.3$ and $9''.3$ (for the two short observations 12025 and 12033, this region goes beyond the edge of the chip, and a rectangular box was used instead for the background).

The resulting stacked spectrum contains around 266,000 counts, reducing to 230,000 after background subtraction, giving a ratio of 6.5:1 for the AGN against the cluster emission. The resulting spectrum was binned to ensure a minimum of 2000 counts per bin¹² and fitted between 0.8 and 5 keV with an absorbed power law, `xswabs × powlaw1d`,

$$AE^{-\gamma} \times e^{-n_H \sigma(E)}. \quad (2)$$

Here A denotes the normalization of the power law, γ is the power-law index, and n_H is the effective hydrogen column density.

The resulting fit is shown in Figure 3, together with the fractional ratio of data to model. The best-fit value of n_H is $2.3 \times 10^{21} \text{ cm}^{-2}$, and the power-law index is $\gamma = 1.83 \pm 0.01$. While the absorbed power law is a reasonable characterization of the data, there are two large localized residuals: one positive between 2 and 2.2 keV and one negative around 3.4–3.6 keV. There is an upward trend at 5 keV. As the effective area of *Chandra* begins to fall off rapidly here, and there are also intrinsically fewer photons expected, pileup plays a proportionately more important role. This rising trend continues beyond 5 keV, and we attribute it to the effects of pileup.

For the 300 ks of ACIS-I midway observations with the AGN around 3 arcmin from the aimpoint, a comparable full extraction results in a spectrum that is still highly piled up (the best-fit index is an unphysically hard $\gamma = 1.30$). This is even more the case for the ACIS-S observations, where a similar full extraction produces an index of $\gamma = 0.66$. For these cases, it is necessary to account for pileup to extract physical results.

¹² In general, we bin so that there are approximately 100 bins in total. If there are too few counts per bin, then the fit is insensitive to localized modulations as can be produced by ALPs.




Figure 3. Stacked spectrum of the ACIS-I edge observations from the entirety of the extraction region, involving 230,000 counts after background subtraction. The fit is to an absorbed power law and results in $n_H = 2.3 \times 10^{21} \text{ cm}^{-2}$ and a spectral index of $\gamma = 1.83$. Here σ refers to the standard deviation from the model expectation for a Poissonian count rate.

Given the high flux levels of the source, a cleaner spectrum requires reducing the effects of pileup in the spectrum. There are two ways of reducing the effects of pileup. The first method involves excluding a central region of highest pileup from the analysis. While this reduces the photon counts—and thus the statistical significance of any features in the spectrum—it also produces purer spectra of higher quality. The second method is to keep the full spectrum but model the pileup explicitly. For the purpose of placing bounds on ALPs and constraining the existence of ALP-induced spectral modulations, the most useful spectra are those where photon events occur at their true energies. We therefore only describe in the main text the creation and analysis of spectra with a central region of high pileup excluded. In Appendix A, we discuss the explicit modeling of pileup in cases where we retain the entirety of the image.

We analyze the images and count rates of each observation using the image software `ds9`. The precise method of central exclusion is somewhat arbitrary; the method used to produce the spectra is as follows. We set the `ds9` binning such that the `ds9` pixels are essentially the same size as the physical *Chandra* pixels (half an arcsecond across). We then create an extraction region manually excluding all `ds9` pixels that are neighbors (either side-by-side or diagonally) to any pixel with total counts greater than 1% of the overall exposure time. Using the `pileup_map` tool, this results in almost all retained pixels having a pileup fraction lower than 5%.¹³ Using `arccorr`, the

responses reflect the presence of a central exclusion in the extraction region.

We first show in Figure 4 the ACIS-I edge observations with a central exclusion according to the above method. As these observations are relatively clean to begin with, the fractional change in the number of counts is relatively small. There are now 187,000 counts before background subtraction and 153,000 counts afterward. The fraction of counts in the 7–10 and 10–15 keV regions are 1.2% and 0.04%. The best-fit spectral index is $\gamma = 1.85 \pm 0.015$ (the errors quoted here and in the remainder of this paper are statistical 1σ errors; see Appendix C.2 for more details) with $n_H = 2.2 \times 10^{21} \text{ cm}^{-2}$.

While the statistical significance of the features at 2–2.2 and 3.4–3.5 keV reduces (consistent with the reduced photon counts), the magnitudes of the data to model fluctuations remain the same.

We now perform a similar cleaning of the ACIS-I observations with the source midway on the chip. As the central pixels are heavily piled up, in this case the cleaning procedure significantly increases the quality of the fit. This is at the cost of a significant reduction in photon statistics: there are now only 88,000 counts before background subtraction and 74,000 after background subtraction. We group counts so that there are 700 counts per bin, leading to a spectral index of $\gamma = 1.64 \pm 0.02$ with $n_H = 1.3 \times 10^{21} \text{ cm}^{-2}$. With a Q -value of 0.18 and a reduced χ^2 of 1.14 for 86 degrees of freedom, this is now an overall good fit to the data (see Figure 5).¹⁴

We note that the spectral index γ is clearly different from that found with the ACIS-I edge data. However, this spectrum was obtained by excluding the central part of the image and extracting only from the wings of the point-spread function. By comparison with the brightness of the AGN measured in the contemporaneous ACIS-I edge observations, we can infer that

¹³ For these observations, the telemetry limit was 15 keV and any (piled-up) photons recorded with energy larger than this were not telemetered to the ground. For there to be a significant number of such events, there would also exist a tail of photons with energies close to but below the telemetry limit. From the $<0.2\%$ fraction of photons measured in the 10–15 keV band for the cleaned spectra, we can infer that these contained very few untelemetered events. The same analysis shows that there are very few untelemetered events for the ACIS-I edge observations, even without excising the central part of the image, but the central high pileup regions of the ACIS-I midway and ACIS-S observations did have significant numbers of untelemetered events.

¹⁴ Fitting the different observations individually produces results for γ that are consistent with the stacked result both for the midway and the edge observations.




Figure 4. Stacked spectrum of the ACIS-I edge observations with central pixels excluded according to the prescription in the text. There are 153,000 counts after background subtraction. The fit is to an absorbed power law and results in $n_H = 2.2 \times 10^{21} \text{ cm}^{-2}$ and a spectral index of $\gamma = 1.85 \pm 0.015$.




Figure 5. Stacked spectrum of the ACIS-I midway observations with central pixels excluded according to the prescription in the text. There are 74,000 counts after background subtraction. The fit is to an absorbed power law and results in $n_H = 1.3 \times 10^{21} \text{ cm}^{-2}$ and a spectral index of $\gamma = 1.64 \pm 0.02$.

this spectrum retains only around a quarter of the total number of photons from the AGN. As the ACIS point-spread function increases with energy, this biases the analyzed photons toward a harder spectrum (and a lower γ). In principle, this is taken into account in the analysis procedure, as the aperture correction of `arccorr` modifies the response to reflect the exclusion of the central part of the image. However, it is also the case that the image has a complex form with neither a spherical nor an elliptical shape. Instead, the image of the AGN—traced by the excluded central region—resembles a “Maltese cross” shape, as can be seen in Figure 1.

The objective of this work is to constrain large, localized modulations in the spectrum of the AGN that could arise from ALP-photon conversion. For this purpose, we do not regard these different values of γ as problematic, particularly given the large difference in observational conditions between the two sets of observations and the large fraction of photons that are excluded in the ACIS-I midway observations.

We perform a similar procedure for the ACIS-S observations. In this case, we reduce pileup by removing a central square of 16 pixels (2 arcsec square) from an extraction region of a circle of radius $3''/4$. The resulting cleaned spectrum has




Figure 6. Stacked spectrum of the ACIS-S observations with a central 16 pixel square excluded. The fit is to an absorbed power law and soft thermal component.

117,000 counts before background subtraction and 74,000 counts after background subtraction. After background subtraction, 0.1% of the counts are in the 10–15 keV band and 1.1% of the counts are in the 7–10 keV band, indicating that this spectrum is now substantially cleaner. We group the counts so that there are at least 700 counts per bin. In this case, an absorbed power law is not sufficient for a good fit, and we supplement this by a soft thermal component using `xsape` (the presence of a thermal component for NGC1275 was also reported in Yamazaki et al. 2013). The presence of a soft thermal component substantially improves the fit, resulting in an acceptable Q -value of 10^{-2} . The resulting fit (see Figure 6) has $n_H = 1.6 \times 10^{21} \text{ cm}^{-2}$ and a power-law index of $\gamma = 1.84 \pm 0.03$. The thermal component has a temperature $T = 0.92 \text{ keV}$. At this temperature, the amplitude and abundance of the thermal component are largely degenerate in the fit. Fixing the abundance at solar abundance, the relative amplitude of the thermal component to the power law is 0.15.

Why is a soft thermal component necessary for a good fit in the ACIS-S observations but not in the ACIS-I observations? There are two main reasons. First, the effective area of ACIS-S has more support at low energies than does ACIS-I. The net result is that for ACIS-S, the peak of the observed photon spectrum is around 1 keV, whereas for the ACIS-I observations, the peak is around 1.6 keV. This implies that the ACIS-S observations have a proportionally higher fraction of low-energy photons and so are more sensitive to the presence of a soft thermal component.

Second, the ACIS-S and ACIS-I observations were taken at different times: the former in 2002 and 2004 and the latter in 2009. Over this period, the intrinsic brightness of the AGN, as measured by the power-law component, increased significantly. Assuming that the soft thermal component arises differently from the hard power law, and that the absolute strength of the soft thermal component has remained the same over this period, its relative significance would be much greater in the ACIS-S observations.

We have checked that one can add a soft thermal component to the ACIS-I spectra of the same absolute magnitude as that for the ACIS-S spectra without affecting the quality of the fit. Because, for the ACIS-I observations, the thermal component neither significantly improves nor worsens the fit and does not affect the above results, we have not included it explicitly for the ACIS-I data.

Finally, we also considered *Chandra*’s grating observations of NGC1275 (ObsIDs 333 and 428 with HETG) as a means to reduce pileup, but this reduced the counts to the point where it did not seem worthwhile.

4.4. Summary of Fits for ALP Constraints

In terms of constraining ALPs and their parameter space, the main consistent feature observed from these plots is that there are no residuals seen above a level of 10%. This result is found across all three distinct classes of observational conditions. If they exist with strong enough coupling, ALPs can induce localized fractional modulations of up to 50% in the spectrum. It is unambiguously the case that no modulations of this size are present in the spectrum, and therefore ALP-photon couplings large enough to generate such modulations are excluded. As any mis-modeling of the actual data would worsen the fit to the real data and increase the size of the residuals in the real data, this ability to constrain ALP parameter space is robust against any such mis-modeling that may be present. Before we determine constraints on ALPs, however, for completeness, we describe the *XMM-Newton* observations of NGC1275.

5. *XMM-Newton* Analysis

We use SAS version 15.0.0. The 2001 and 2006 data sets are reprocessed with up-to-date calibrations via `cifbuild`. To remove flares, we apply the standard filters of counts $\text{s}^{-1} < 0.35$ for MOS and counts $\text{s}^{-1} < 0.4$ for pn. There is a significant flare toward the end of the 2006 observational period and various shorter flares during the 2001 observation.

Removing the polluted time intervals results in reducing the effective exposure time from 53 to 49 ks for MOS 2001, 25 to 7.4 ks for pn 2001, 123 to 117 ks for MOS 2006, and 76 to 49 ks for pn 2006.

For the extraction regions, we first choose circles of radii 13''8 (MOS1 2001), 17''5 (MOS2 2001), 14''3 (MOS1 2006), 17''5 (MOS2 2006), 15''5 (pn 2001), and 19'' (pn 2006). We then check if pileup is an issue using the SAS tool `epatplot` and find that pileup is in general present in all the observations. Since the AGN was less bright by a factor of ~ 2 in 2001, pileup is slightly less of an issue for these observations. By contrast, the smaller pixel size of 1''.1 of the MOS cameras makes it less susceptible to pileup than the 4''.1 pixel size of the pn camera. The pn's time resolution of 73.4 ms in full-frame mode is advantageous with respect to pileup compared to the 2.6 s time resolution of MOS. However, in extended full-frame mode, the pn's 199.1 ms time resolution makes it less advantageous with respect to pileup compared to MOS. Note that after flare removal there are only 7.4 ks of pn data in full-frame mode, while there are 49 ks of more piled-up data in extended full-frame mode. As the MOS data is expected to be more sensitive, we do not present the pn data from either 2001 or 2006 and instead focus on the MOS data only.

In order to make a compromise between avoiding the most heavily piled-up regions and not too small signal over background ratio, we choose the following inner annuli radii: 2''.5 (MOS1 2001), 3'' (MOS2 2001), 7''.5 (MOS1 2006), and 7''.5 (MOS2 2006). For the background regions, we choose annuli of 13''.8–22'' (MOS1 2001), 17''.5–22''.5 (MOS2 2001), 14''.3–22''.5 (MOS1 2006), and 17''.5–22''.5 (MOS2 2006). The worse angular resolution of *XMM-Newton* compared to *Chandra* implies that the discrimination of the AGN emission against the cluster is relatively poor (for the 2006 data, the AGN:cluster data ratio is 1:3, whereas for *Chandra* spectra, this ratio is not worse than 1.7:1).

After applying `eveselect` to generate the spectra, we use `rmfgen` and `arfgen` to generate the response files. We use `epicspeccombine` to combine the MOS1 and MOS2 spectra of the 2001 and 2006 observations, respectively. We also use the `ftools` routines `mathpha`, `addrmf`, and `addarf` (`epicspeccombine` does not produce an `arf` file) and find that the results are comparable within statistical errors to using `epicspeccombine`. In order to avoid combining different systematic errors, we do not combine the 2001 and 2006 data.

The spectra are further analyzed using the spectral fitting programs *Xspec* (Arnaud 1996) and *Sherpa* (Freeman et al. 2001), requiring a minimum count rate per bin of about 1.5% of the overall count rate. The energy interval is 0.5–7.5 keV.

For NGC1275, we use a spectral model of a power law as in our *Chandra* analysis,

$$A(E) = K E^{-\gamma}, \quad (3)$$

where K is the normalization in photons $\text{keV}^{-1} \text{cm}^{-2} \text{s}^{-1}$ and γ is the photon index. To model a soft excess of the AGN, we add an *apec* model whose temperature we expect to be around 1 keV. The abundance parameter of this model is set to unity.¹⁵ We also include the well-known iron line at 6.4 keV in the cluster rest frame in our spectral fit.

¹⁵ Note that around 1 keV, this *apec* spectrum is dominated by atomic lines. Hence, there is a degeneracy between the abundance parameter and the normalization of the *apec* model.

To model photoelectric absorption, the source spectrum is multiplied by the *zwabs* model,

$$M(E) = e^{-n_H \sigma(E)}, \quad (4)$$

where n_H is the hydrogen column density and σ is the photoelectric cross-section (Morrison & McCammon 1983). Photoelectric absorption is mostly relevant below approximately 1 keV if $n_H \gtrsim 10^{21} \text{ atoms cm}^{-2}$. (The galactic column density is measured as $n_H = 1.5 \cdot 10^{21} \text{ atoms cm}^{-2}$.)

The fitted spectra are shown in Figures 7 and 8. While the fit is good (see Table 3), there are both fewer counts and a worse contrast against the cluster background than for the *Chandra* data. From the perspective of ALP constraints, we can say that modulations of the spectrum are not allowed beyond the 20% level—a weaker constraint than that for the *Chandra* data.

6. Bounds on ALPs

From the perspective of constraints on ALPs, the overall summary of the data is that the spectrum from NGC1275 is fit by an absorbed power law with residuals from the power law not larger than 10%. We will further discuss residuals at the 10% level in Appendix B (in particular features around 2–2.2 and at 3.5 keV). However, the absence of any modulations in the spectrum at a level greater than 10% allows powerful constraints to be placed on ALP parameter space.

In particular, this allows us to say that on passing from NGC1275 through the Perseus cluster and toward us, $\langle P(\gamma \rightarrow a) \rangle \lesssim 20\%$. It follows that ALP-photon couplings large enough to generate the saturated limit of $\langle P(\gamma \rightarrow a) \rangle = 1/3$ are robustly excluded, as these would produce larger residuals in the data than are actually observed.

Note that this approach to placing bounds is conservative. The simulated data (with ALPs) is idealized—it has no pileup present and assumes a perfect subtraction of the cluster background. Couplings are excluded when they lead to simulated data that would still be a worse fit than the actual data, even though the simulated data is observationally ideal while the real data contains effects such as pileup and cluster background that lead to increased residuals and make it far noisier than simulated data.

To obtain approximate bounds on $g_{a\gamma\gamma}$, we compare two models for the flux $F(E)$ observed from NGC1275:

1. Model 0: an absorbed power law $F_0(E) = AE^{-\gamma} \times e^{-n_H \sigma(E)}$, as described in Equation (2).
2. Model 1: an absorbed power law multiplied by the photon survival probability assuming the existence of ALPs with coupling $g_{a\gamma\gamma}$. In this case, the predicted flux also depends on the magnetic field \mathbf{B} along the line of sight. We have $F_1(E, \mathbf{B}) = AE^{-\gamma} \times e^{-n_H \sigma(E)} \times P_{\gamma \rightarrow \gamma}(E(1+z), \mathbf{B}, g_{a\gamma\gamma})$.

Although we have (limited) empirical estimates of the strength of the magnetic field in Perseus, the exact structure is unknown. In practice, we randomly generate many instances of the field from a given power spectrum. The parameter most relevant to ALP-photon conversion is the central magnetic field strength B_0 , estimated as 25 μG in Taylor et al. (2006). Based on results for the Coma cluster, we assume that B decreases with radius as $B \propto n_e^{0.7}$ (Bonafede et al. 2010). The electron density n_e has the radial distribution found in




Figure 7. MOS 2001 spectral fit. The best-fit parameters are listed in Table 3.

Churazov et al. (2003),

$$n_e(r) = \frac{3.9 \times 10^{-2}}{\left[1 + \left(\frac{r}{80 \text{ kpc}}\right)^2\right]^{1.8}} + \frac{4.05 \times 10^{-3}}{\left[1 + \left(\frac{r}{280 \text{ kpc}}\right)^2\right]^{0.87}} \text{ cm}^{-3}.$$

We simulate each field realization with 600 domains. The length l of each domain is between 3.5 and 10 kpc, randomly drawn from a power-law distribution with minimum length 3.5 kpc and power 0.8. We therefore have

$$P(l = x) = N \begin{cases} 0 & \text{for } x > 10 \text{ kpc} , \\ x^{-1.2} & \text{for } 3.5 \text{ kpc} < x < 10 \text{ kpc} , \\ 0 & \text{for } x < 3.5 \text{ kpc} , \end{cases} \quad (5)$$

with normalization constant N .

The coherence length and power spectrum of the magnetic field in the center of Perseus is not observationally determined. Instead, these parameters are motivated by those found for the cool-core cluster A2199 (Vacca et al. 2012), taking a conservative value for the magnetic field radial scaling. The magnetic field and electron density are constant in each domain, with $B(r)$ and $n_e(r)$ evaluated at the center of the domain and the direction of \mathbf{B} chosen at random.

We compute 95% confidence limits on $g_{a\gamma\gamma}$ by generating fake data from Model 1 and assessing how well it is fit by Model 0, i.e., how well the oscillations due to ALP-photon conversion can hide in the Poisson noise. As it has minimal pileup, we use the clean ACIS-I edge observations for this analysis. We fit the spectrum between 1 and 4 keV (a region unaffected by pileup) and bin such that there are 1000 counts in each energy bin. We use Sherpa’s Levenberg–Marquardt fitting method with Poisson errors derived from the value of the

data in each bin. Our procedure to determine whether ALPs with coupling $g_{a\gamma\gamma}$ are excluded at the 95% confidence level is as follows.

1. Fit Model 0 to the real data and find the corresponding reduced χ^2 , χ_{data}^2 .
2. Randomly generate 50 different magnetic field realizations \mathbf{B}_i for the line of sight to NGC1275.
3. For each \mathbf{B}_i , compute $P_{\gamma \rightarrow \gamma}(E, \mathbf{B}_i, g_{a\gamma\gamma})$ by numerically propagating photons at different energies through \mathbf{B}_i , as described, for example, in Angus et al. (2014). We take 300 photon energies equally spaced between 1 and 4 keV.
4. For each \mathbf{B}_i , generate 10 fake data sets from Model 1, using Sherpa’s fake pha method.
5. Fit Model 0 to each of the 500 fake data sets and find the corresponding reduced χ^2 , χ_i^2 , for each.
6. If fewer than 5% of the χ_i^2 are lower than χ_{data}^2 , $g_{a\gamma\gamma}$ is excluded at the 95% confidence level.

We scan over $g_{a\gamma\gamma}$ in steps of $10^{-13} \text{ GeV}^{-1}$. For the $g_{a\gamma\gamma}$ value excluded, we also check that the three values above it in our grid are also excluded. For the magnetic field parameters described above, we find $g_{a\gamma\gamma} \lesssim 1.4 \times 10^{-12} \text{ GeV}^{-1}$, which is shown in Figure 9.

If we consider a more pessimistic scenario with $B_0 = 15 \mu\text{G}$ and a minimum coherence length of 0.7 kpc, we instead find $g_{a\gamma\gamma} \lesssim 2.7 \times 10^{-12} \text{ GeV}^{-1}$. If we take an even more pessimistic scenario in which the central field is $B_0 = 10 \mu\text{G}$ and the minimum coherence length is 0.7 kpc, the bound increases further to $g_{a\gamma\gamma} \lesssim 4.0 \times 10^{-12} \text{ GeV}^{-1}$.

We finally note that if we excluded the feature at 2.2 keV from the fit, removing the 1.8–2.3 keV region based on this being the location of an effective area edge, the bounds would




Figure 8. MOS 2006 spectral fit. The best-fit parameters are listed in Table 3.

Table 3

Fit Results for the *XMM-Newton* MOS Observations from 2001 and 2006

	MOS 2001	MOS 2006
Exposure (ks)	49	118
Counts	~47000	~63000
γ	1.65 ± 0.03	1.65 ± 0.03
T_{soft}	0.84 ± 0.06	0.80 ± 0.06
n_{H} (10^{22} atoms cm^{-2})	0.13 ± 0.03	0.13 ± 0.03
χ^2/dof	154 / 131	205 / 178
Q -value	0.08	0.08

Note. These results are in agreement with the fit parameters found in Churazov et al. (2003) for the 2001 MOS data of NGC1275.

improve significantly, to $g_{a\gamma\gamma} \lesssim 1.1 \times 10^{-12} \text{ GeV}^{-1}$ for the case of a central magnetic field of $B_0 = 25 \mu\text{G}$ (see Appendix B). This illustrates the conservative nature of the bounds compared to any mis-modeling of the actual data.

6.1. Validation with MARX Simulation

As a validation of these bounds, we also perform MARX simulations (Davis et al. 2012) of the ACIS-I edge observations to confirm that the results are robust against any pileup effects. Significantly piled-up data is known to produce complex effects, in particular with a bright background (see, for instance, Harris et al. 2006).

As the actual spectrum is a sum of the AGN power law and thermal emission from the cluster, we first need to determine

the magnitude of the thermal emission with the ACIS-I edge extraction region.

To extract the level of thermal emission within the ACIS-I edge observations, we follow two strategies. First, we take the longest on-axis ACIS-S observation (*Chandra* ObsID 4952) and use the same extraction region as for the ACIS-I edge data (also excluding the AGN with a central circle of radius 2 arcsec). We fit an *apepec* model to this, obtaining a temperature of $T \sim 4 \text{ keV}$, an abundance of $Z \sim 0.4$, and a normalization of 0.0022. While such a single-temperature *apepec* model is not an accurate description of the gas dynamics at the center of the cluster ($\chi^2/N = 474/375$), it is adequate for our purposes, as we are more concerned with the level of pileup and thus the overall magnitude of the thermal emission, rather than a fully perfect spectral shape. Using this model, together with the exposure and responses for the ACIS-I edge observations, suggests the presence of around $\sim 30,000$ thermal counts present in the ACIS-I spectrum.

Second, as a check on this, we also use purely the ACIS-I edge data and directly fit it with a sum of an absorbed power law and absorbed thermal emission. In doing this, we freeze the temperature and abundance of the thermal model to the values obtained from the *Hitomi* spectrum of the central region of the Perseus cluster, $T = 3.48 \text{ keV}$ and $Z = 0.54$ (Aharonian et al. 2017). This produces a best-fit normalization for the *apepec* norm of 0.0027, comparable to that attained when using the ACIS-S data to determine the magnitude of the thermal emission. This slightly higher value would correspond to a total of $\sim 36 \text{ k}$ counts from thermal emission within the ACIS-I edge spectrum. As both results are consistent, they show that of the $\sim 270 \text{ k}$ counts in the ACIS-I edge spectrum, only around 10%–15% arise from thermal cluster emission.




Figure 9. For a central magnetic field strength $B_0 = 25 \mu\text{G}$, we show the 95% upper bounds on the ALP-photon coupling in comparison to previously obtained bounds on ALPs.

We then implement a MARX simulation of the ACIS-I edge observations. For this, we use the observational parameters (date, on-axis location, and roll angle) for the longest ACIS-I edge observation of the AGN (*Chandra* ObsID 11713). As all of the four actual ACIS-I edge observations have similar properties, we set the exposure of this MARX simulation to 200 ks, the total exposure of the actual data.

We first simulate the data without any axions present.

1. The thermal emission is simulated as a spatially extended Gaussian with a width of 30 arcmin centered on the AGN. It is modeled as an `apec` model with a temperature $T = 3.5 \text{ keV}$ and an abundance $Z = 0.48$. The normalization of the model is adjusted so that it has the same fitted strength as that for real data (as determined above) when extracted within the region used for the ACIS-I edge data.
2. The AGN emission is simulated as a point-source power law. The normalization of the power law is adjusted so that—after it has been combined with the thermal emission and the pileup processing applied—it has the same normalization as the real data when extracted over the same extraction region.
3. The two individual simulations are then combined using `marxcat` and the MARX pileup processing applied to the combined simulation.
4. Counts are grouped to 500, and the combined spectrum is fitted with the sum of a power law plus thermal emission. The previous two steps are iterated until the fitted strength and index of the power law matches that of the real data.

We generated a total of 50 fake data samples in this way. To match the fitted strength of the AGN ($4.7 \times 10^{-3} \text{ ph cm}^{-2} \text{ s}^{-1} \text{ keV}^{-1}$

at 1 keV) required an intrinsic strength of $6.4 \times 10^{-3} \text{ ph cm}^{-2} \text{ s}^{-1} \text{ keV}^{-1}$ at 1 keV. Comparison of the number of counts within the extraction region in the MARX event files pre- and post-pileup processing suggests an overall pileup fraction of 15% in the data set. Nonetheless, the pileup is not sufficient to make the spectral fit unsatisfactory (the reduced χ^2 values cluster around unity for ~ 250 degrees of freedom, with the largest reduced χ^2 being 1.37).

As a test of an even more piled-up spectrum with a similar number of photon counts, we also simulated data using a 100 ks MARX exposure but with the intrinsic AGN strength doubled. As this normalization of $12.8 \times 10^{-3} \text{ ph cm}^{-2} \text{ s}^{-1} \text{ keV}^{-1}$ at 1 keV is much larger than that present in the 2016 *Hitomi* data, and as the AGN has been increasing in strength from 2001 onward (see Fabian et al. 2015), pileup is much greater than in the actual 2009 data. The fit now tends to be worse, but not terrible. Among the 50 fake data samples, the average reduced χ^2 is ~ 1.25 for ~ 250 degrees of freedom, but there are still a few cases with a reduced $\chi^2 < 1$.

We now repeat the MARX simulations including the effects of axions in the data. We multiply the AGN power law with a $P(\gamma \rightarrow \gamma)$ survival probability coming from photon-ALP conversion. As this leads to a net reduction in the number of photons present, we adjust the intrinsic normalization upward so that the fitted normalization matches the actual data.

For a fixed ALP coupling, we generate 50 fake data samples. To compare these results to those described above using Sherpa’s `fake_pha` command, we simulate three separate ALP couplings: $g_{a\gamma\gamma} = 1, 1.5, 2 \times 10^{-12} \text{ GeV}^{-1}$. In all cases, we assume the “optimistic” magnetic field model. We directly fit the absorbed sum of a power law and `apec` thermal emission to the ACIS-I edge extraction region for the fake data




Figure 10. Example of simulated data using *ChaRT* and *MARX* reproducing the observational conditions for the ACIS-I edge observations for an axion coupling of $g_{a\gamma\gamma} = 2 \times 10^{-12} \text{ GeV}^{-1}$. In this case, the spectrum is directly fit with the sum of a power law and thermal *ap_{ec}* emission (the ingredients of the simulation). Despite the pileup fraction of $\sim 15\%$, the axion-induced modulations are clearly visible and cannot be fit by any conventional astrophysical model.

and compared to the quality of the fit when doing so using the real data.¹⁶ We first group counts to 500 and fit between 1 and 5 keV using the `chi2dataavar` statistic. Applied to the real data, this gives a reduced χ^2 of 1.83 (coming from the excess around 2 keV).

We then apply this fit to the data involving axions. Of the 50 fake data sets produced using $g_{a\gamma\gamma} = 2 \times 10^{-12} \text{ GeV}^{-1}$, the lowest reduced χ^2 was 4.07. When $g_{a\gamma\gamma} = 1.5 \times 10^{-12} \text{ GeV}^{-1}$, the average reduced χ^2 is still ~ 2.5 , but now with 3 out of 50 spectra with better fits than the actual data ($\chi^2/N < 1.83$). When $g_{a\gamma\gamma} = 1.0 \times 10^{-12} \text{ GeV}^{-1}$, more than half of the fake spectra have better fits than the actual data. Using 95% confidence level exclusions, this shows that $g_{a\gamma\gamma} = 2 \times 10^{-12} \text{ GeV}^{-1}$ is strongly excluded (the best scenario having a reduced χ^2 of 4.07), while $g_{a\gamma\gamma} = 1.5 \times 10^{-12} \text{ GeV}^{-1}$ is on the boundary of exclusion, and $g_{a\gamma\gamma} = 1 \times 10^{-12} \text{ GeV}^{-1}$ is clearly not excluded.

This extremely sharp behavior of the reduced χ^2 with $g_{a\gamma\gamma}$ can be understood qualitatively. As $P(\gamma \rightarrow a) \propto g_{a\gamma\gamma}^2$, a linear increase in $g_{a\gamma\gamma}$ leads to a quadratic increase in the amplitudes of modulations in the data. For a Poissonian process, the amount of data required to detect a fixed fractional deviation grows quadratically with the size of the deviation. So, roughly, a decrease in $g_{a\gamma\gamma}$ by a factor of two requires a sixteen-fold increase in the quantity of data for the same statistical sensitivity.

We also repeat this using a 100 ks exposure with an AGN that is twice as bright (to give a data sample that has a similar number of photons but is substantially more piled up than the actual data). It is again the case that $g_{a\gamma\gamma} = 2 \times 10^{-12} \text{ GeV}^{-1}$ is

strongly excluded, $g_{a\gamma\gamma} = 1 \times 10^{-12} \text{ GeV}^{-1}$ is not excluded, and $g_{a\gamma\gamma} = 1.5 \times 10^{-12} \text{ GeV}^{-1}$ is marginal.

We note that there are, of course, several differences between the *MARX* simulation and actual data processing. The *MARX* treatment of pileup is simpler than the real physics of deposited electron clouds from the interaction of photons with CCD chips. The *MARX* simulation also does not take into account the charge transfer inefficiencies present on the actual chips, which obstruct the flow of charge to the readout. However, these results imply that, in terms of bounding $g_{a\gamma\gamma}$, uncertainties due to pileup appear to be far smaller than uncertainties due to the cluster magnetic field.

The above bounds used *MARX*'s simulation of *Chandra*'s optics. The most advanced simulation of *Chandra*'s optics is through the *ChaRT* simulator (Carter et al. 2003), rather than *MARX*. As *ChaRT* requires the original source spectrum to be manually uploaded to a website, it is not possible to automate this process to produce axion bounds (as every different ALP conversion template represents a different source spectrum). However, for each of the couplings $g_{a\gamma\gamma} = 1, 1.5, 2 \times 10^{-12} \text{ GeV}^{-1}$, we simulated some individual *ChaRT* spectra. For the spectra simulated, the results are similar to those of the *MARX* simulations. We plot in Figure 10 a sample spectrum for $g_{a\gamma\gamma} = 2 \times 10^{-12} \text{ GeV}^{-1}$. The ALP-induced modulations are clearly visible, leading to a large badness-of-fit.

7. Conclusions

The most basic point of this paper is that X-ray observations of NGC1275 are a superb way to search for ALPs. NGC1275 is an extremely bright—and brightening—X-ray point source shining through a galaxy cluster environment. If ALPs exist, they can lead to oscillatory modulations within the energy ranges probed by *Chandra* and *XMM-Newton*. Although there is uncertainty on the precise magnetic field structure along the line of sight to NGC1275, it is almost certainly the case that

¹⁶ This is marginally different from the analysis procedure described in earlier sections, as it aims at fitting the thermal emission rather than subtracting it. We also did this analysis by subtracting background emission from the real data set and fitting only with a power law. The results, however, were very similar, with no significant changes to the excluded coupling.

current and future observations of NGC1275 will provide a greater reach in searches for light ALPs with $m_a \lesssim 10^{-12}$ eV than even the proposed dedicated experiment the International AXion Observatory (IAXO) (Armengaud et al. 2014).

Previous *Chandra* observations of NGC1275 already provide a data set of extraordinary quality. Three factors contribute to this. First, *Chandra*’s angular resolution allows a large contrast between the AGN emission and cluster. Second, the existing observations are very deep and include three independent locations of NGC1275 on the telescope. Third, it is a fortuitous fact that in one set of observations, NGC1275 was located on the edge of the chip, thereby providing a clean observation that minimizes pileup.

In this paper, we have used all existing observations of NGC1275 with *Chandra* and *XMM-Newton* to search for spectral modulations induced by ALPs. If they exist, ALPs give rise to spectral modulations, and we have used the absence of modulations at the $\mathcal{O}(30\%)$ level to place leading current bounds on ALP-photon couplings.

At the $\mathcal{O}(10\%)$ level, there are two main modulations in the data: one upward around 2–2.2 keV and one downward around 3.4–3.5 keV. The former can be attributed to the iridium edge in the mirror. In the main body of the paper, we have focused on constraints that can be placed on ALP parameter space. In Appendix B, we discuss the residuals in greater detail. In particular, the residual around 3.5 keV can be well described as an absorption feature at (3.54 ± 0.02) keV, the same energy as the diffuse cluster excess observed in Bulbul et al. (2014) and Boyarsky et al. (2014).

The existing data set already places powerful constraints on ALPs. Nonetheless, from the perspective of ALP physics, the data set could be substantially improved with even relatively modest further observations. NGC1275 is now brighter than it was in 2009 by a factor of two. Observations with either *XMM-Newton* in small-window mode or *Chandra* off-axis and in reduced frame time could give a larger and cleaner data set than the best used in this paper, the 2009 ACIS-I edge observations. X-ray observations of the center of the Perseus cluster therefore provide an outstanding way to look for new physics beyond the Standard Model.

We thank David Marsh and the referee for comments on the paper, Alexis Finoguenov and Jeroen Franse for helpful discussions, and the Bethe Centre for Theoretical Physics for hospitality while part of this work was carried out. We would like to thank Gray Rybka for providing his exclusion plot template. We also thank the staff and science teams of the *Chandra* and *XMM-Newton* X-ray observatories for friendly helpdesks and documentation that is sufficiently clear that particle theorists can follow it. This project is funded in part by the European Research Council starting grant ‘‘Supersymmetry Breaking in String Theory’’ (307605). JC is also funded by a Royal Society University Research Fellowship.

Software: CIAO v4.7 (Fruscione et al. 2006), Sherpa (Freeman et al. 2001), HEASOFT v6.17, SAS v15.0.0, MARX 5.3.2, ChaRT v2.

Appendix A Modeling Pileup

The main text addressed pileup by removing central, high-pileup regions from the extraction region. Here we take a complementary approach, aiming to model the pileup directly.

While this modeling does not provide a complete picture of pileup, it ameliorates its effect. A full treatment would require customized tools going beyond the scope of this paper (the pileup tool *jdpileup* (Davis 2001) provided with Sherpa is optimized for on-axis sources, while NGC1275 is off-axis in all ACIS-I observations).

We first describe our pileup modeling for the ACIS-S observations, where NGC1275 is on-axis. This results in a high degree of pileup—the central pixels have large numbers of counts above 10 keV, all of which arise from multi-photon pileup events.

We use the *jdpileup* model as described in Davis (2001). As this model assumes that the source is on-axis, it is directly appropriate here. The model assumes Poisson statistics to calculate the probability of different numbers of photons hitting an event-detection cell (a 3×3 pixel region in *Chandra*) within the readout time adjusted to these respective observations. For the ACIS-S and ACIS-I midway observations, the readout time was 3.1 s, and for the ACIS-I edge observations, it was 3.2 s. The model then convolves this probability with the probability of such events being assigned a ‘‘good’’ grade and the conversion from photon energy to pulse height by the detector. Obviously, this is a nonlinear process: the pileup in any particular bin depends on the energy spectrum for all energies below the bin. The parameters of the model therefore need to be determined together with those for the spectrum under consideration. This can lead to degeneracies in parameter space, in particular for a simple power law.

Per Davis (2001), the two parameters of the pileup model that we allow to vary are α and f , where α^{p-1} is the probability that p piled photons will be assigned a ‘‘good’’ grade and f is the fraction of events to which pileup will be applied. The other parameter that will prove important is n , the number of regions to which *jdpileup* will be applied independently. For a point source, this should be set to 1, as was done for the ACIS-S observations. For an extended source, it should roughly correspond to the number of 3×3 pixel islands in the region. The reason for this is that *jdpileup* assumes spatial uniformity across the extraction region. For the ACIS-I midway observations, where the AGN is smeared out across several pixel islands, the value of n proved difficult to determine for large extraction regions, with the value corresponding to the best fit being unrelated to the number of pixel islands. We therefore constrained ourselves to a smaller central region with little variation in count rate between pixels to give us more control over the pileup model.

It is also worth noting that the *jdpileup* model is set to zero for energies less than 0.5 keV and greater than 10 keV. As there are many events above the 10 keV range for the ACIS-S observations, we only model the spectrum up to 10 keV; extending the fit beyond that would result in the model parameters being sent to unphysical values.

For the ACIS-S pileup analysis, we used a circular region around NGC1275 of radius $3''/5$, with the background taken from an annular region of inner radius 4 arcsec and outer radius 7 arcsec. Counts were grouped to 1500 per bin. The spectrum was modeled using an absorbed power law and thermal emission with temperature $T = 0.85$ keV and fitted with the *jdpileup* model for energy values between 1 and 10 keV. The data were fit using the *moncar* Monte Carlo method and run several times to ensure the global minimum had been found.




Figure 11. ACIS-S observations, involving 177,000 counts after background subtraction. The fit is to an absorbed power law with a thermal component of $T = 0.85$ keV, and pileup is modeled with `jdpileup`. The ratio of data to model is shown at the bottom of the figure.




Figure 12. ACIS-I midway observations from a central 6×6 pixel region, involving 134,000 counts after background subtraction. The fit is to an absorbed power law, and pileup is modeled with `jdpileup`.

The resulting fit and the ratio of data to model are shown in Figure 11. The best-fit parameters involve $n_H = 2.6 \times 10^{21} \text{ cm}^{-2}$ and a spectral index $\gamma = 1.81$.¹⁷ The best-fit α and f parameters of `jdpileup` were $\alpha = 0.660$ and $f = 0.943$, respectively. We can clearly see that this model provides a reasonable description of the data all the way up to 10 keV and produces physically sensible values for n_H and γ , despite estimating a pileup fraction of 82%. The reduced statistic of the

fit is 1.75, with a Q -value of 10^{-5} . While we reemphasize that this will not represent a perfect account of pileup, it does capture the relevant physics, producing a sensible fit with physical parameters.

For the ACIS-I midway observations, we considered a central 6×6 pixel extraction region (this is almost exactly the complement of the clean ACIS-I midway spectrum used in the previous subsection). The resulting spectrum contains around 136,000 counts, reducing to 134,000 after background subtraction, giving a very high AGN-to-cluster contrast of 67:1. After background subtraction, the fraction of counts in

¹⁷ As the uncertainties on the pileup modeling are hard to quantify, we only quote best-fit parameters and do not include errors.

the 7–10 and 10–15 keV bands is 5.3% and 1.2%, respectively. Counts were grouped to 1000 per bin. The spectrum was modeled using an absorbed power law and fitted with the `jdpileup` model, this time for energy values between 1 and 9 keV, to ensure no counts with energy greater than 10 keV were included in the final bin. The `jdpileup` parameter n was set to 4, the number of 3×3 pixel islands in the extraction region.

The resulting fit and ratio of data to model are shown in Figure 12. The best-fit parameters involve $n_H = 2.5 \times 10^{21} \text{ cm}^{-2}$ and a spectral index $\gamma = 1.93$. The best-fit α and f parameters of `jdpileup` were $\alpha = 0.324$ and $f = 0.975$, respectively, and the estimated pileup fraction was 35%. While the fit is not perfect, it does give a reasonable characterization of the data.

We finally consider the case of the ACIS-I edge observations, where pileup is relatively weak. Although a reasonable fit can be made with no pileup modeling, we now aim at modeling the pileup in these observations for completeness. There is a rule of thumb¹⁸ that 0.007 counts s^{-1} per 3×3 pixel island is about 1% pileup and 0.07 counts s^{-1} is about 10%. We will assume that the ACIS-I edge observations with mild cleaning have pileup fractions of 10% or lower.

For the ACIS-I edge observations, we used as a source an elliptical annulus of radii 3 and 5 arcsec, with the 12 brightest pixels removed by a contour. The background was extracted from an elliptical annulus of outer radii 18 and 23 arcsec, with inner radii 8 and 12 arcsec. The resulting spectrum contains around 128,000 counts, reducing to 111,000 after background subtraction, giving an AGN-to-cluster ratio of 6:1. After background subtraction, the fraction of counts in the 7–10 and 10–15 keV bands is 1.1% and 0.1%, respectively. We grouped counts to 1000 per bin. Here there are fewer counts out at higher energies than in the previous two sets of observations, so we fit only out to 5 keV. The `jdpileup` parameter n was set to 16, the number of 3×3 pixel islands in the extraction region. This should only be thought of as a rough estimate, as the `jdpileup` model is stated to be accurate for on-axis point sources, and the edge observations are also not spatially uniform in terms of counts. The fit is to an absorbed power law and results in $n_H = 2.7 \times 10^{21} \text{ cm}^{-2}$ and a spectral index of $\gamma = 1.89$. The best-fit α and f parameters of `jdpileup` were $\alpha = 0.52$ and $f = 0.87$, respectively. The pileup fraction was 3%, consistent with estimates and very mild, as expected. The resulting fit and ratio of data to model are shown in Figure 13.

Appendix B

Analysis of 10% Residuals

No residuals in the spectra exceed 10%. At the 10% level, the *Chandra* spectra presented in the main body of the paper show two main features departing from a power law: an excess at 2–2.2 keV, present in all *Chandra* observations with overwhelming statistical significance, and a deficit at 3.4–3.5 keV that is not as strong but is still present at almost 5 (local) σ in the ACIS-I edge data. These features ensure that an absorbed power law (plus thermal component) is not a good fit to the data.

First, we perform a statistical analysis of these features at $E \sim 2$ keV and $E \sim 3.5$ keV. Second, we mention potential instrumental, astrophysical, or new physics explanations.

B.1. Statistical Significance of Features at 2 and 3.5 keV

We first consider the $E \sim 2$ keV feature. Considering the clean ACIS-I edge data, we include a positive Gaussian (`xswabs*(xspowerlw + xsgaussian)`) and analyze its effect on the fit. For a zero-width Gaussian, this leads to an improvement in the fit by $\Delta\chi^2 = 33.7$ with a best-fit energy of 2.13 keV. For a finite-width Gaussian, the improvement is $\Delta\chi^2 = 48.0$ with a best-fit energy of 2.02 keV and a best-fit width of 0.2 keV.

For the ACIS-I midway data, a similar additional Gaussian improves the fit by $\Delta\chi^2 = 18.6$ with a best-fit energy of 2.17 keV (the fitted width is much smaller than the detector resolution, so the results are identical for zero-width and finite-width Gaussians). For the ACIS-S data, the additional Gaussian improves the fit by $\Delta\chi^2 = 20.3$ with a best-fit energy of 2.06 keV (again, fitting the width gives a result smaller than the detector resolution and so does not affect the result).

For the 3.5 keV feature, we consider the cleaned data sets for the three sets of observations and include a negative Gaussian (`xswabs*(xspowerlw - xsgaussian)`). We formally treat the Gaussian as zero width, but any finite width much narrower than the ACIS energy resolution gives an identical result. For the three data sets (ACIS-I edge, ACIS-I midway, and ACIS-S), we plot in Figures 14–16 the improvement in the χ^2 that can come from adding an additional negative Gaussian $\Delta\chi^2_{\text{NG}}$. The ACIS-I edge data show a strong preference for an additional negative Gaussian ($\Delta\chi^2 \sim 17$) at 3.5 keV. For both the ACIS-I midway and ACIS-S data, a negative Gaussian at 3.5 keV mildly improves the fit ($\Delta\chi^2 \sim 1.5$ and $\Delta\chi^2 \sim 0.4$) but is not required (as the ACIS-I edge data set is both larger and cleaner than the other two data sets, these results are consistent—in particular, the inferred strength of the dip is consistent within 2σ).

Across all three plots, the largest feature is clearly seen to be the residual in the ACIS-I edge data at 3.5 keV. The next strongest residuals are two features at $E \sim 1.2$ keV (ACIS-I edge) and $E \sim 1.4$ keV (ACIS-S) with $\Delta\chi^2 \sim 10$. In both cases, there are very strong atomic lines visible in the background spectra at precisely these energies (Fe XXI, Fe XXII, Fe XXIII, Fe XXIV, and Mg XII). These features can be reliably associated with these atomic lines, coming from a small mis-subtraction of the background (deep in the core of the Perseus cluster, the physical conditions of the gas in the signal region and background region will not be precisely identical).

In Figure 17, we produce a combined plot of the significance of the 3.5 keV deficit across all *Chandra* observations. In order to compare observations taken at different times and with different intrinsic AGN luminosities, to measure the strength of the deficit we have used the induced fractional deficit of photons in the 3.3–3.7 keV regime. The feature remains over 3σ significant globally and is also located at precisely the same energy as the diffuse cluster excess observed in Bulbul et al. (2014) and Boyarsky et al. (2014).

B.1.1. The 2.2 keV Residual

Despite the overwhelming statistical significance of the feature at 2–2.2 keV, it is at the same location as a sharp effective area edge from the iridium coating of the mirrors. As pileup arising from high flux levels can generate fake excesses at the location of such edges, we associate the feature with this edge and do not discuss it further.

¹⁸ cxc.harvard.edu/csc/memos/files/Davis_pileup.pdf




Figure 13. ACIS-I edge observations, involving 110,000 counts after background subtraction. The fit is to an absorbed power law, and pileup is modeled with `jdpileup`.




Figure 14. Improvement in χ^2 attainable by adding a negative Gaussian at the specific energy for ACIS-I edge observations.

B.1.2. The 3.5 keV Deficit

We now consider possible explanations for the feature at 3.4–3.5 keV that is present as a deficit of data compared to the model. This is present at very high significance in the ACIS-I edge observations. Given the smaller data samples present there, the ACIS-I midway and ACIS-S observations are compatible but not significant in themselves.

While this feature lacks the statistical significance of the one at 2–2.2 keV, it is harder to come up with instrumental or astrophysical explanations for such a deficit. Compared to the 2–2.2 keV region, at 3.4–3.5 keV the *Chandra* effective areas are smooth functions of energy, and in general this is a clean part of the spectrum, so the previous arguments for pileup contamination or effective area miscalibration do not apply here.

The simplest explanation is a statistical fluctuation—although, as this deficit is present as a local $\sim 4.5\sigma$ deficit in the ACIS-I edge observation, such an explanation is problematic. It is also notable that this occurs at precisely the same energy as that for the diffuse excess found in Bulbul et al. (2014) and Boyarsky et al. (2014).

In terms of conventional astrophysical explanations, the most obvious candidate for this deficit would be an atomic absorption line. However, it is difficult to see how this could work, as there are no strong lines around this energy. Even more seriously, absorption along the sight line to NGC1275 dominantly arises from the Milky Way—this is evidenced by the fitted values of n_H , which are all consistent with the galactic value toward Perseus of $n_H \simeq 1.5 \times 10^{21} \text{ cm}^{-2}$. This implies that, if the feature at 3.4–3.5 keV did arise from atomic absorption, a similar feature should also arise for the continuum cluster spectrum of Perseus, as the galactic n_H values have degree-scale gradients and thus n_H is approximately the same for all sight lines to Perseus. However, no such absorption feature at 3.5 keV is detected in the cluster spectra of Perseus. Indeed, precisely the opposite: a significant excess is found at 3.5 keV (Bulbul et al. 2014; Boyarsky et al. 2014; for a review, see Iakubovskiy 2016).

The most plausible conventional explanation of the 3.4–3.5 keV deficit is then as a localized detector effect or statistical fluctuation for the ACIS-I edge observations, combined with mild downward statistical fluctuations for the




Figure 15. Improvement in χ^2 attainable by adding a negative Gaussian at the specific energy for ACIS-I midway observations.




Figure 16. Improvement in χ^2 attainable by adding a negative Gaussian at the specific energy for ACIS-S observations.

ACIS-S and ACIS-I central observations. However, we also note that inspection of the background region for the ACIS-I edge observations does not show any deficit around 3.4–3.5 keV.

New Physics Explanations. Given the observations of Bulbul et al. (2014) and Boyarsky et al. (2014) of an unidentified line at 3.5 keV in diffuse emission from the Perseus cluster, it is an interesting fact that we observe a deficit at ~ 3.5 keV in our observations of the NGC1275 AGN at the center of the cluster. This of course may be just a coincidence, but in the context of new physics models, we mention two ways that these facts could be more than simply a coincidence.

The first involves models of excited dark matter invoked for the 3.5 keV line, where dark matter has a resonance at an energy ~ 3.5 keV above its ground state (for example, as in Cline & Frey 2014b, 2014a; Finkbeiner & Weiner 2016; Agrawal et al. 2015; Berlin et al. 2015; D’Eramo et al. 2016). In this case, there is an absorption cross-section of $E \sim 3.5$ keV photons on dark matter. While for an isotropic initial distribution of photons such absorption and re-emission would not affect the photon spectrum, for a directional beam dark matter absorption will result in an absorption hole in the spectrum (the presence of an absorbing torus around an AGN ensures that its outward radiation is indeed directional).

Would such an effect have been observed already elsewhere? We do not see why. In this scenario, the relevant quantity determining the fractional absorption rate is the dark matter column density along the line of sight. It is entirely plausible that the dark matter column density toward the NGC1275 AGN is larger than for almost any other direction in the universe. This is because the emission all originates very close to the central AGN; thus, the column density is sensitive to not just the Perseus cluster but also the central cluster galaxy NGC1275 right down to any subparsec-level dark matter spikes close to

the central supermassive black hole. The effect we observe is not large—a 10% reduction over around 100 eV in width—and requires a spectrum with $\mathcal{O}(10^5)$ counts for a statistically significant detection. With a smaller dark matter column density, the effect would reduce to an unobservable $\mathcal{O}(1\%)$ effect.

There is a second possible connection to the 3.5 keV line. An attractive scenario for the 3.5 keV line involves the decay of dark matter to 3.5 keV ALPs, which then convert to photons through axion-photon conversion in the cluster magnetic field (Cicoli et al. 2014; Conlon & Day 2014; Conlon & Powell 2015; Alvarez et al. 2015). In this scenario, the strength of the 3.5 keV line depends on the efficiency of ALP-photon interconversion and so is broadly expected to be larger in regions with large magnetic fields extended over wide areas: for example, in galaxy clusters as opposed to galaxies. The 3.5 keV line is observed to be stronger toward the center of Perseus than for other clusters; one way this could arise is if it fortuitously happens that ALP-photon interconversion is particularly efficient around 3.5 keV for sight lines toward the center of Perseus. In this case, the presence of a deficit of $E \sim 3.5$ keV photons from NGC1275 and an excess of $E \sim 3.5$ keV photons from the cluster as a whole could come from the same underlying physics—efficient photon-ALP interconversion at energies $E \sim 3.5$ keV along sight lines toward the center of Perseus.

B.1.3. ALP Interpretation of Features

The purpose of this paper is to use the extraordinary data set of counts from the NGC1275 AGN to search for spectral irregularities, with the intent of constraining ALP parameters. This search has resulted in two features being present in the data at high statistical significance, the former of which is consistent with arising from pileup around the iridium edge.




Figure 17. Overall significance and location of the 3.5 keV deficit, combined across all *Chandra* observations (cleaned data sets only). We have used the induced fractional deficit in photons in the 3.3–3.7 keV region to facilitate a uniform comparison for the strength of the dip, as the AGN luminosity varies between observations.




Figure 18. Fit to the clean ACIS-I edge observations with an absorbed power law multiplied by the photon survival probability for a specific example of ALP-photon conversion. We include ALPs with $g_{a\gamma\gamma} = 1 \times 10^{-12} \text{ GeV}^{-1}$ and assume a central field of $B_0 = 25 \mu\text{G}$. The reduced χ^2 is 1.51, compared to 1.65 for a fit to an absorbed power law without ALPs.

It is not possible for us to constrain ALP couplings beyond a level where they would produce residuals comparable to these features. We can estimate approximately that to produce such residuals would require an ALP-photon coupling of similar magnitude to the limits placed in Section 6: $g_{a\gamma\gamma} \sim 1\text{--}5 \times 10^{-12} \text{ GeV}^{-1}$. For example, Figure 18 shows a fit to the clean ACIS-I edge observations with an absorbed power law multiplied by the photon survival probability $P_{\gamma \rightarrow \gamma}$. In this

case, $P_{\gamma \rightarrow \gamma}$ was calculated assuming the existence of ALPs with $g_{a\gamma\gamma} = 1 \times 10^{-12} \text{ GeV}^{-1}$ and a central magnetic field $B_0 = 25 \mu\text{G}$ (the most optimistic field scenario used in Section 6). We see that the anomalies at 2.2 and 3.5 keV have been alleviated by the presence of ALPs (although, for this magnetic field, at the expense of creating similarly sized anomalies at higher energies). Figure 19 shows the corresponding photon survival probability spectrum.




Figure 19. Photon survival probability for the fit shown in Figure 18.

Table 4

Fit Parameters for a Power Law + Gaussian Fit between 5.5 and 7 keV

Data Set	Power-Law Index	E_{line} (keV)	$\Delta\chi^2$ (dof)
ACIS-I edge	$0.84^{+0.3}_{-0.1}$	$6.32^{+0.06}_{-0.05}$	6.8/2
ACIS-I midway	$-0.14^{+0.03}_{-0.02}$	$6.33^{+0.05}_{-0.06}$	2.2/2
ACIS-S	$0.49^{+0.07}_{-0.11}$	$6.29^{+0.02}_{-0.03}$	12.2/2

Appendix C Technical Details

C.1. The Presence of the 6.4 keV Iron $K\alpha$ Line

As the cleaned spectra in the main body of the paper are only extracted up to 5 keV, it is a useful consistency check to ensure that the reflected iron $K\alpha$ line at 6.4 keV (6.3 keV after redshifting) is present in our spectra. For this purpose, we use the cleaned spectra for each of the three sets of *Chandra* observations and fit the spectrum only between 5.5 and 7 keV. We first fit a power law and then the sum of a power law and a Gaussian (we do not fit n_{H} , as absorption is irrelevant at these energies).

We determine the best-fit central energy of the additional Gaussian and the resulting improvement in χ^2 . These fits are summarized in Table 4, and we explicitly show the fit for the ACIS-S data in Figure 20.

These show that the iron line is clearly present in the ACIS-S and ACIS-I edge data and marginally present in the ACIS-I midway data. As ACIS-S has better energy resolution than ACIS-I, it is unsurprising that the Fe line is found at the highest significance there.

We note that all spectral indices are unphysically hard (and even negative for the ACIS-I midway data). This is a clear consequence of pileup. For energies above 5 keV, the *Chandra* effective area is falling off rapidly. There are also intrinsically far fewer source photons than are present at lower energies, given the power-law spectrum. The effects of pileup at these


higher energies are therefore far more severe than at lower energies.

While these spectra have been cleaned, this does not mean that there is zero pileup. For the bulk of the spectra, the effects of pileup in the cleaned spectra are at the level of a few percent. However, for $E > 5$ keV, the effects of pileup are substantially more important for the above reasons: the combination of the intrinsic decline in the spectrum and the falloff in the *Chandra* effective area implies that pileup plays a proportionally far more severe role at high energies.

C.2. Errors in the Fitted Parameters

For fits of spectra to models, the paper includes the fitted statistical errors. However, significant caution should be applied when interpreting such a quantity (for example, a power-law index) as an error on the intrinsic spectrum of the source. Systematic errors on overall power-law indices are likely to be much larger. While hard to quantify, there are several clear sources of such systematic effects.

1. Although the cleaned spectra have significantly reduced levels of pileup, no spectrum can be entirely free of pileup. Pileup automatically leads to a hardening of spectral indices as it moves photons from low to high energies.
2. The point-spread function of *Chandra* is a function of energy, and more energetic photons tend to be spread out more. For spectra such as those used here, where a central and highly piled-up core is excluded, this will bias the analyzed photons to higher energy. While the *Chandra* analysis software aims to take this into account, this correction will not be perfect, particularly for off-axis sources where the form of the optical image is more complex (as for the ACIS-I midway data, where the AGN image takes on a “Maltese cross” form).
3. Generally, there are systematic errors on the overall power-law index that will arise because the optical conditions are significantly different on-axis and off-axis.

Figure 20. Fe K α line in the cleaned ACIS-S data.

These are hard to quantify but are certainly much larger than the purely statistical errors on the fitted parameters.

4. The AGN luminosity is time variable, and the power-law index will be time variable as well. For observations taken at different times, the intrinsic power law of the source may be different.

ORCID iDs

Sven Krippendorff  <https://orcid.org/0000-0001-6374-6828>

References

- Agrawal, P., Chacko, Z., Kilic, C., & Verhaaren, C. B. 2015, *JHEP*, **08**, 072
- Aharonian, F. A., Akamatsu, H., Akimoto, F., et al. 2017, *ApJL*, **837**, L15
- Alvarez, P. D., Conlon, J. P., Day, F. V., Marsh, M. C. D., & Rummel, M. 2015, *JCAP*, **1504**, 013
- Angus, S., Conlon, J. P., Marsh, M. C. D., Powell, A. J., & Witkowski, L. T. 2014, *JCAP*, **1409**, 026
- Armengaud, E., Avignone, F. T., Betz, M., et al. 2014, *JINST*, **9**, T05002
- Arnaud, K. A. 1996, in ASP Conf. Ser. 101, in *Astronomical Data Analysis Software and Systems V*, ed. G. H. Jacoby & J. Barnes (San Francisco, CA: ASP), 17
- Balmaverde, B., Capetti, A., & Grandi, P. 2006, *A&A*, **451**, 35
- Berlin, A., DiFranzo, A., & Hooper, D. 2015, *PhRvD*, **91**, 075018
- Bonafede, A., Feretti, L., Murgia, M., et al. 2010, *A&A*, **513**, A30
- Boyarsky, A., Ruchayskiy, O., Iakubovskiy, D., & Franse, J. 2014, *PhRvL*, **113**, 251301
- Brax, P., Brun, P., & Wouters, D. 2015, *PhRvD*, **92**, 083501
- Brockway, J. W., Carlson, E. D., & Raffelt, G. G. 1996, *PhLB*, **383**, 439
- Bulbul, E., Markevitch, M., Foster, A., et al. 2014, *ApJ*, **789**, 13
- Burridge, C., Davis, A.-C., & Shaw, D. J. 2009, *PhRvL*, **102**, 201101
- Carter, C., Karovska, M., Jerius, D., Glotfelty, K., & Beikman, S. 2003, in ASP Conf. Ser. 295, *Astronomical Data Analysis Software and Systems XII*, ed. H. E. Payne, R. I. Jedrzejewski, & R. N. Hook (San Francisco, CA: ASP), 477
- Churazov, E., Forman, W., Jones, C., & Bohringer, H. 2003, *ApJ*, **590**, 225
- Cicoli, M., Conlon, J. P., Marsh, M. C. D., & Rummel, M. 2014, *PhRvD*, **90**, 023540
- Cicoli, M., Goodsell, M., & Ringwald, A. 2012, *JHEP*, **10**, 146
- Cline, J. M., & Frey, A. R. 2014a, *PhRvD*, **90**, 123537
- Cline, J. M., & Frey, A. R. 2014b, *JCAP*, **1410**, 013
- Collaboration, T. F.-L. 2016, arXiv:1603.06978
- Conlon, J. P. 2006, *JHEP*, **05**, 078
- Conlon, J. P., & Day, F. V. 2014, *JCAP*, **1411**, 033
- Conlon, J. P., & Marsh, M. C. D. 2013, *PhRvL*, **111**, 151301
- Conlon, J. P., Marsh, M. C. D., & Powell, A. J. 2016, *PhRvD*, **93**, 123526
- Conlon, J. P., & Powell, A. J. 2015, *JCAP*, **1501**, 019
- Davis, J. E. 2001, *ApJ*, **562**, 575
- Davis, J. E., Bautz, M. W., Dewey, D., et al. 2012, *Proc. SPIE*, **8443**, 84431A
- D'Eramo, F., Hambleton, K., Profumo, S., & Stefaniak, T. 2016, *PhRvD*, **93**, 103011
- Dobrynina, A., Kartavtsev, A., & Raffelt, G. 2015, *PhRvD*, **91**, 083003
- Fabian, A. C., Walker, S. A., Pinto, C., Russell, H. R., & Edge, A. C. 2015, *MNRAS*, **451**, 3061
- Fairbairn, M., Rashba, T., & Troitsky, S. V. 2011, *PhRvD*, **84**, 125019
- Finkbeiner, D. P., & Weiner, N. 2016, *PhRvD*, **94**, 083002
- Freeman, P., Doe, S., & Siemiginowska, A. 2001, *Proc. SPIE*, **4477**, 76
- Fruscione, A., McDowell, J. C., Allen, G. E., et al. 2006, *Proc. SPIE*, **6270**, 62701V
- Grifols, J. A., Masso, E., & Toldra, R. 1996, *PhRvL*, **77**, 2372
- Harris, D. E., Cheung, C. C., Biretta, J. A., et al. 2006, *ApJ*, **640**, 211
- Horns, D., Maccione, L., Meyer, M., et al. 2012, *PhRvD*, **86**, 075024
- Iakubovskiy, D. 2016, *AASP*, **6**, 3
- Meyer, M., Horns, D., & Raue, M. 2013, *PhRvD*, **87**, 035027
- Mirizzi, A., & Montanino, D. 2009, *JCAP*, **0912**, 004
- Morrison, R., & McCammon, D. 1983, *ApJ*, **270**, 119
- Payez, A., Cudell, J. R., & Hutsemekers, D. 2012, *JCAP*, **1207**, 041
- Payez, A., Evoli, C., Fischer, T., et al. 2015, *JCAP*, **1502**, 006
- Powell, A. J. 2015, *JCAP*, **1509**, 017
- Raffelt, G., & Stodolsky, L. 1988, *PhRvD*, **37**, 1237
- Ringwald, A. 2012, *PDU*, **1**, 116
- Schliederer, M., & Sigl, G. 2016, *JCAP*, **1601**, 038
- Sikivie, P. 1983, *PhRvL*, **51**, 1415
- Sikivie, P. 1985, *PhRvD*, **32**, 2988
- Svrcek, P., & Witten, E. 2006, *JHEP*, **06**, 051
- Tavecchio, F., Roncadelli, M., Galanti, G., & Bonnoli, G. 2012, *PhRvD*, **86**, 085036
- Taylor, G. B., Gugliucci, N. E., Fabian, A. C., et al. 2006, *MNRAS*, **368**, 1500
- Vacca, V., Murgia, M., Govoni, F., et al. 2012, *A&A*, **540**, A38
- Wouters, D., & Brun, P. 2012, *PhRvD*, **86**, 043005
- Wouters, D., & Brun, P. 2013, *ApJ*, **772**, 44
- Yamazaki, S., Fukazawa, Y., Sasada, M., et al. 2013, *PASJ*, **65**, 30