

3 **Foreground-limited XMM-Newton Constraints on Soft X-Ray Emission toward the Inner Halo of M31**

4 RUI HUANG¹

5 ¹*Affiliations and coauthors to be completed*

ABSTRACT

The soft X-ray emission toward the inner halo of M31 contains contributions from both M31 and the Milky Way, whose thermal spectra are nearly degenerate at CCD resolution. We analyze 22 XMM-Newton pointings at projected radii of approximately 10–31 kpc and treat the directly measured cool component as the full line-of-sight sum rather than assigning it to either galaxy. Our homogeneous baseline excludes the MOS 1.4–2.0 keV instrumental-line interval, uses a pointing-dependent 15 arcmin HI4PI atomic hydrogen column, and retains 14 technically valid dual-MOS fields in the primary sample. The primary spectra are consistent with a common temperature of $kT = 0.17531 \pm 0.00436$ keV, with no evidence for additional temperature scatter. The absorbed 0.4–1.25 keV surface flux is $(1.106 \pm 0.054) 10^{-15}$ erg cm⁻² s⁻¹ arcmin⁻² in the North/NW fields and $(0.894 \pm 0.203) 10^{-15}$ erg cm⁻² s⁻¹ arcmin⁻² in the South/SE fields. Their difference, $(0.212 \pm 0.210) 10^{-15}$ erg cm⁻² s⁻¹ arcmin⁻², is positive but not statistically compelling. The current data therefore establish a spatially widespread, approximately 0.18 keV line-of-sight component but do not detect a large-scale North–South asymmetry or uniquely identify an M31 contribution. Any decomposition into M31, Milky Way, or Local-Group emission requires an explicit spatial or foreground prior.

6 *Keywords:* Andromeda Galaxy (39) — Circumgalactic medium (1879) — Diffuse x-ray background
7 (384) — X-ray astronomy (1810) — Milky Way Galaxy (1054)

8 1. INTERNAL DRAFT STATUS

9 This manuscript is a complete internal scientific draft
10 based on the validated v19 independent-fit baseline. It
11 is not yet ready for submission. The numerical val-
12 ues quoted below use local LevMar covariance matri-
13 ces. The submission version must replace the descriptive
14 side means with a shared radial model plus a side off-
15 set, compute profile-likelihood or simulation-calibrated
16 intervals at non-negative boundaries, and add a con-
17 trolled spectral-systematic grid. These requirements are
18 summarized in Section 7.

9 2. INTRODUCTION

20 Hot circumgalactic gas is expected to contain a sub-
21 stantial fraction of the baryons associated with massive
22 spiral galaxies, but its low density makes direct X-ray
23 measurements difficult. Deep observations of individ-
24 ual massive spirals and stacked samples demonstrate
25 extended soft X-ray halos, while also showing that in-
26 ferred profiles depend on galaxy mass, unresolved-source
27 subtraction, and background treatment (Anderson et al.
28 2016; Li et al. 2017, 2018; Zhang et al. 2024a,b). M31
29 is close enough that its inner halo can be sampled with

30 multiple XMM-Newton fields, but its large angular ex-
31 tent creates a different limitation: emission from the
32 M31 halo is superposed on spatially structured soft emis-
33 sion from the Milky Way and potentially the Local
34 Group.

35 The central few kiloparsecs of M31 contain multiphase
36 diffuse X-ray emission associated with the bulge and disk
37 (Shirey et al. 2001; Takahashi et al. 2004; Li & Wang
38 2007). Those measurements do not determine the am-
39 plitude of cooler gas at projected radii of 10–31 kpc. Off-
40 disc XMM-Newton observations provide a closer spec-
41 tral precedent. Kavanagh et al. (2020) measured a com-
42 ponent near 0.17–0.18 keV and explicitly modeled it as
43 the combined emission of the Galaxy, M31, and the in-
44 tergalactic medium. A possible Local Hot Bridge in the
45 direction of M31 has a similar inferred temperature (Qu
46 et al. 2021). Milky Way surveys also find characteristic
47 temperatures near 0.15–0.22 keV but substantial spatial
48 variation (Henley & Shelton 2013a, 2015; Ponti et al.
49 2023; Kaaret et al. 2020; Ueda et al. 2022; Locatelli et al.
50 2024).

51 These overlapping temperature scales make distance
52 attribution from an EPIC spectrum alone unreliable.

We therefore organize the analysis around the directly observed line-of-sight component,

$$S_{\text{obs}}(E) = S_{\text{MW}}(E) + S_{\text{M31}}(E), \quad (1)$$

with an optional third Local-Group term in explicitly labeled conditional scenarios. We ask two assumption-light questions. First, what cool thermal surface flux is measured toward the North/NW and South/SE inner-halo fields? Second, how strongly do the data constrain a difference between the two sampled sides? The M31-versus-Milky-Way decomposition is discussed only after these direct measurements.

3. OBSERVATIONS AND DATA REDUCTION

3.1. XMM-Newton program and field selection

We use observations from XMM-Newton proposal 080073, “XMM-Newton Legacy Survey of M31 Halo: Searching for the Missing Accreted Hot CGM” (PI J.-T. Li), obtained with the EPIC cameras (Jansen et al. 2001; Turner et al. 2001). The selected observations were taken between 2017 June 28 and 2018 January 17. All 43 detector products used here have the Thin1 optical blocking filter. The v19 analysis is MOS-only: 21 fields have one MOS1 and one MOS2 spectrum, while 0800730101 contributes MOS2 only. EPIC-pn products are not used in this homogeneous baseline.

The 22-field parent set is the proposal-080073 subset for which the upstream project screening retained the required diffuse-field products. Observation 0800731701 is absent from the local reduction tree, 0800732401 lacks the required `4.background` products, and 0800731201, 0800731401, 0800731801, 0800732501, 0800732901, and 0800733001 were rejected upstream because of extreme 2.0–3.2 keV observed-to-QPB ratios. Observation 0800733101 was rejected because of extreme residual Soft Proton contamination and incomplete MOS coverage. These exclusions define the parent set before the v19 spectral-quality cuts described in Section 4.4; they are not selected on fitted CGM_{sum} temperature or surface flux.

Pointing centers are taken from the final PHA headers. For an adopted M31 center and distance of 785 kpc (McConnachie et al. 2005), their circular projected radii span 10.5–30.8 kpc. The pointings form two approximately minor-axis chains, labeled North/NW and South/SE, rather than a complete azimuthal annulus. In the final primary sample, the median radii are 19.28 and 14.97 kpc on the two sides, respectively, and no South/SE primary field lies at 20–25 kpc. Table 1 gives the observation-level inventory.

3.2. SAS processing provenance

The spectra were derived from the existing project reduction with the XMM-Newton Science Analysis System (SAS; Gabriel et al. 2004). We audited the headers of every v19 input rather than assigning generic “standard SAS” settings. All 43 source PHA files report `xmmsas_20211130_0941-20.0.0`, `evselect-3.71.1`; their ARFs report `arfgen-1.102.1`, and their RMFs report `rmfgen-2.8.5`. The event selections stored in every PHA require `PATTERN<=12` and `FLAG==0`. The spectra therefore include valid MOS single–quadruple events and reject flagged events.

The final PHAs contain the applied good-time intervals as embedded GTI extensions, but the local tree does not retain the upstream flare-light-curve thresholds or a command log that uniquely reconstructs those GTIs for the full sample. One partial command record survives for 0800732701, but it cannot define a homogeneous 22-field prescription. Likewise, the local `ccf.cif` files are not bound to the PHA/ARF/RMF products by a recorded identifier or hash, and their file dates do not uniquely establish which calibration index was active during each extraction. We therefore report the verified SAS task versions and final cleaned exposures, but do not invent sample-wide flare thresholds or a CCF release.

The retained tree is sufficient to reproduce the paper analysis from the 43 final spectral product groups. It is not sufficient for a bitwise ODF-to-PHA rebuild: complete ODF material is retained locally for only three selected observations, and original `mos-spectra/QPB` command files are absent for most fields. Recovery of the homogeneous flare-screening, calibration-index, and extraction-command provenance remains a submission requirement.

3.3. Source masks and extraction footprint

The analysis uses the products under `rpc/4.background/bkg_mos[12]_00500-02000_exclude_extent_source_mask`. The extraction is the usable MOS field of view after detector-coordinate chip/gap masks and point/extended-source exclusions; it is not a simple 15 arcmin circle. The authoritative mask geometry is embedded in each PHA as `REGION` extensions and in the `SLCTEXPR` selection. The embedded regions contain negative detector-coordinate ellipses for excluded sources together with the valid chip and field-of-view geometry. The reduction tree also retains the associated `edetect_stack` source-position and extended-source region products. The exact source-detection likelihood threshold and the rule that maps detections to exclusion-ellipse size are not recoverable

Table 1. XMM-Newton Observation and v19 Product Inventory

ObsID	R.A.	Decl.	Start date	Side	R_{proj}	t_{M1}	t_{M2}	N_{HI}	v19 use
	(deg)	(deg)	(UTC)		(kpc)	(ks)	(ks)	(10^{20} cm^{-2})	
0800730101	9.2301	43.2408	2017-06-28	North/NW	30.78	–	10.01	5.58	Numerical only
0800730201	8.8793	42.8879	2017-07-25	North/NW	28.79	13.92	11.49	5.78	Primary
0800730301	8.4932	42.5826	2017-07-14	North/NW	28.69	16.11	18.08	5.76	Primary
0800730401	8.7369	42.2086	2017-07-27	North/NW	23.71	14.51	14.86	5.46	Excluded
0800730501	9.0777	42.4873	2017-07-28	North/NW	23.40	11.83	14.28	5.53	Primary
0800730601	9.3863	42.8204	2017-08-12	North/NW	25.03	11.57	11.10	5.79	Primary
0800730701	9.5865	42.4375	2017-08-12	North/NW	19.54	13.61	14.35	5.73	Primary
0800730801	9.2710	42.1599	2017-08-10	North/NW	18.92	14.49	14.75	5.60	Primary
0800730901	9.0152	41.8738	2018-01-07	North/NW	19.01	18.41	18.80	5.59	Primary
0800731001	9.1456	41.5646	2017-08-10	North/NW	16.32	14.82	15.09	5.79	Excluded
0800731101	9.8256	42.0115	2018-01-07	North/NW	13.45	16.48	18.05	5.67	Primary
0800731301	10.6467	42.2062	2018-01-11	North/NW	12.85	6.14	7.06	7.64	Excluded
0800731501	9.7074	41.4971	2017-08-12	North/NW	10.52	14.73	14.98	5.85	Primary
0800731601	9.4171	41.2226	2017-08-02	North/NW	13.07	20.57	20.27	6.30	Primary
0800731901	11.6666	40.9259	2018-01-13	South/SE	11.17	17.08	16.83	6.48	Primary
0800732001	11.9666	41.3051	2018-01-13	South/SE	13.21	16.51	16.85	6.92	Primary
0800732101	12.2198	40.9845	2018-01-13	South/SE	16.31	7.99	10.33	5.90	Excluded
0800732201	12.0572	40.7280	2018-01-14	South/SE	16.01	5.92	5.88	5.58	Excluded
0800732301	11.3785	40.1663	2018-01-15	South/SE	16.73	8.39	9.74	5.29	Primary
0800732601	12.6676	40.3879	2018-01-16	South/SE	23.84	3.04	2.80	5.17	Excluded
0800732701	12.2969	39.9949	2018-01-17	South/SE	24.20	5.33	2.20	5.27	Excluded
0800732801	12.1075	39.7012	2018-01-17	South/SE	26.10	6.36	7.98	4.81	Primary

NOTE—Exposures are the cleaned PHA exposures for the individual detector products, not independent spacecraft time. The 43 MOS products sum to 533.571 ks of detector exposure. “Numerical only” denotes the MOS2-only field 0800730101. “Excluded” includes technical or numerical-quality failures; the exact reason is retained in the machine-readable observation table.

152 from the current command history and must be frozen
153 before submission.

154 We convert the PHA BACKSCAL value with the produc-
155 tion relation $\Omega_i = \text{BACKSCAL}_i (1/20/60)^2 \text{ arcmin}^2$, cor-
156 responding to the 0.05 arcsec DETX/DETY sampling.
157 The resulting detector-specific areas span 238.3–565.8
158 arcmin². MOS1 and MOS2 areas are kept separate, so
159 failed CCDs, chip gaps, bad detector regions, and source
160 holes propagate directly into each model through its own
161 area factor. Relative to the 706.9 arcmin² area of a
162 nominal 15 arcmin circle, the individual detector foot-
163 prints cover 33.7–80.0%. The one-PHA-per-field conven-
164 tion used in the earlier absorption-footprint audit covers
165 33.7–52.6% (median 45.2%). The 15 arcmin aperture is
166 used only for the baseline HI4PI column estimate, not
167 for spectral extraction.

168 3.4. Observed, QPB, and response spectra

169 For each retained detector we use four matched prod-
170 ucts: `*-obj.pi` for the observed diffuse-field spectrum,
171 `*-back.pi` for the ESAS quiescent-particle background
172 (QPB), and the corresponding `*.arf` and `*.rmf`. The
173 QPB is generated in the same detector footprint as
174 the observed spectrum. It represents the particle back-
175 ground and is not a blank-sky or astrophysical fore-
176 ground spectrum. In Sherpa, `obj.pi` is loaded as the
177 source dataset, `back.pi` is loaded as its background with
178 the stored errors, and QPB subtraction is applied before

179 grouping and fitting. The remaining spectrum therefore
180 contains celestial emission plus residual detector back-
181 ground and Soft Protons, which are modeled explicitly
182 below.

183 The cleaned detector exposures range from 2.20 to
184 20.57 ks and sum to 533.571 ks over the 43 detector
185 products. The full product paths, task-version strings,
186 exposures, areas, embedded-mask checks, and quality
187 classifications are written to the machine-readable file
188 `m31_cgmsum_v19_observation_products.csv`. Its gen-
189 erator reads the frozen production manifest and the ac-
190 tual FITS headers, and fails unless it finds exactly 22
191 pointings, 43 detector products, Thin1 in every PHA,
192 and the verified PATTERN/FLAG selections.

193 3.5. Energy filtering and response-aware grouping

194 The final MOS fitting domain is

$$195 \quad 0.4\text{--}1.4 \text{ keV} \cup 2.0\text{--}8.0 \text{ keV}. \quad (2)$$

196 The 1.4–2.0 keV interval is excluded because it contains
197 the strong MOS Al $K\alpha$ and Si $K\alpha$ instrumental features
198 and is sensitive to residual Soft Protons. Every fit-stage
199 energy selection first resets the broad band and then
200 reapplies this gap; the final Sherpa sessions preserve the
201 two disjoint intervals.

202 Spectra are grouped after QPB subtraction. Starting
203 from individual channels, a bin is closed only after both
204 (1) the QPB-subtracted signal-to-noise ratio

$$S/N = \frac{C_{\text{obj}} - C_{\text{QPB}}}{\sqrt{C_{\text{obj}} + C_{\text{QPB}}}} \quad (3)$$

reaches 6 and (2) its channel width samples the local RMF FWHM by no more than a factor of 6. Group accumulation is reset at ignored channels, and group starts are forced at 0.4, 1.4, 2.0, and 8.0 keV. Consequently, no fitted bin bridges the excluded MOS interval.

4. SPECTRAL ANALYSIS

4.1. Response construction and model topology

We fit each pointing independently in Sherpa 4.18 (Freeman et al. 2001) using XSPEC model components (Arnaud 1996). Within one pointing, the MOS1 and MOS2 celestial parameters are linked, whereas the residual instrumental-line and Soft Proton parameters remain detector specific. No parameter is linked between different pointings. If R_i denotes the detector RMF folded with its ARF, D_i denotes the same redistribution with a unit effective area, and Ω_i is the PHA solid angle, the count model for detector i is

$$M_i(E) = R_i \left\{ \Omega_i [G_i(E) + S_{\text{LB}}(E) + S_{\text{SWCX}}(E) + A(N_{\text{HI},i}) \{S_{\text{CXB}}(E) + S_{\text{CGM}_{\text{sum}}}(E) + S_{\text{hot}}(E)\}] \right\} + D_i \{ \Omega_i P_i(E) \}. \quad (4)$$

Here G_i is the residual detector-line model, P_i is the Soft Proton continuum, and the remaining terms are sky surface-brightness components. Multiplication by Ω_i converts every surface-brightness normalization to the actual detector footprint before response folding. The multiplicative sky-background scale present in the implementation is fixed to unity.

The component called `MW halo` in the fitting code is denoted `CGMsum` here. It is the cool absorbed APEC component required by the line of sight, not a spectroscopic assignment to the Milky Way. The historical target-specific `SrcApec`, `SrcPo`, and `Srcabs` components remain in the implementation for compatibility but are set to zero and frozen in this independent diffuse-field analysis. Thus no additional M31-only spectral component is fitted on top of `CGMsum`.

4.2. Absorption and atomic configuration

For each pointing, the absorption column is fixed to the uniform-area mean of the HI4PI atomic N_{HI} map within a 15 arcmin aperture (HI4PI Collaboration 2016). The adopted values span 4.81×10^{20} to $7.65 \times 10^{20} \text{ cm}^{-2}$. The map has a native angular resolution of 16.2 arcmin; a nominal aperture therefore contains only about 2.38

independent beams. The aperture is centered on the pointing coordinate, has complete valid-map coverage, and is not weighted by the MOS source mask or exposure map.

These inputs are atomic H I columns, not total equivalent X-ray absorbing columns. Molecular gas, dust-to-gas conversion, 21-cm optical-depth corrections, and the mismatch between the circular HI4PI aperture and the irregular MOS footprints are not included in the formal errors. They are retained as controlled absorption-systematics branches.

The XSPEC configuration is fixed to the `angr` solar-abundance table, `vern` photoelectric cross sections, and APEC/AtomDB 3.0.9 (Foster et al. 2012). The exact APEC line and continuum files are present in the CIAO installation and their SHA-256 hashes are stored in the production manifest.

4.3. Staged fitting procedure

All fits use the Gehrels chi-square statistic (`chi2gehrels`) and the Sherpa Levenberg–Marquardt optimizer. The staged procedure is identical for every pointing:

1. In the initial 0.2–1.4∪2.0–8.0 keV fit, set the Local Bubble to $kT = 0.1$ keV and $Z = 1$, set `CGMsum` to $kT = 0.2$ keV and $Z = 0.3$, and fix the hot component to $kT = 0.7$ keV and $Z = 1$. The Local Bubble, `CGMsum`, hot-component, and CXB normalizations are free. The per-detector Soft Proton normalization and both slopes are initially frozen with zero normalization. The three SWCX line normalizations remain zero.
2. Thaw the Local Bubble and `CGMsum` temperatures within the bounds in Table 2, keep both abundances fixed, and refit the linked MOS spectra.
3. Apply the final domain in Equation 2 and refit the thermal, CXB, and residual-line parameters.
4. For each detector, thaw the broken-power-law normalization and both photon indices while keeping the 3.2 keV break fixed, and refit all active parameters jointly.
5. Freeze the fitted high-energy Soft Proton index. Refit the low-energy index and normalization with the physical ordering $0 \leq \Gamma_{\text{low}} \leq \Gamma_{\text{high}}$; if the unconstrained intermediate fit violates the ordering, initialize $\Gamma_{\text{low}} = \Gamma_{\text{high}}/2$ before the final fit.

The production worker publishes a session only when Sherpa reports `succeeded=True`. Every published side-

Table 2. Baseline Spectral Components and Parameter Treatment

Component	XSPEC/Sherpa form	Absorption/response	MOS linking	Baseline parameter treatment
Local Bubble	<code>apec</code>	unabsorbed, R_i	linked	$kT = 0.08\text{--}0.12$ keV free; $Z = 1$ and $z = 0$ fixed; normalization free.
CGM _{sum}	<code>phabs*apec</code>	absorbed, R_i	linked	$kT = 0.14\text{--}0.40$ keV and normalization free; $Z = 0.3$ and $z = 0$ fixed.
Hot thermal	<code>phabs*apec</code>	absorbed, R_i	linked	$kT = 0.7$ keV, $Z = 1$, and $z = 0$ fixed; non-negative normalization free.
CXB	<code>phabs*pegpwlw</code>	absorbed, R_i	linked	Photon index 1.46 fixed; 0.5–2.0 keV pegged normalization free.
SWCX lines	three <code>gaussians</code>	unabsorbed, R_i	linked	O VII, O VIII, and Fe-L normalizations fixed to zero.
Residual 1.3 keV line	<code>gaussian</code>	R_i	independent	Centroid free in 1.2–1.4 keV, $\sigma = 0.01$ keV fixed, normalization free.
Al/Si detector lines	two <code>gaussians</code>	R_i	independent	Normalizations fixed to zero because 1.4–2.0 keV is excluded.
Soft Proton	<code>bknpower</code>	D_i (unit ARF)	independent	Break fixed at 3.2 keV; normalization and low-energy slope free; final high-energy slope fixed after an intermediate fit.

Note. “Linked” means shared between MOS1 and MOS2 within one pointing only. All celestial and detector-continuum normalizations are expressed per unit solid angle before multiplication by the detector-specific Ω_i .

car records the fit statistic, method, number of bins, degrees of freedom, all-stage filters, input/output hashes, and final parameter states. The current baseline uses one deterministic starting sequence and local optimization; multi-start and profile-likelihood checks are intentionally listed as submission work rather than implied here.

4.4. Quality selection and validation

All 22 fields were attempted before the quality mask was applied. The structural validator requires the frozen input detector inventory, unique file paths, exact code/input/output/APEC hashes, final fit success and message, finite fit dimensions, all-stage energy filters, grouping parameters, the adopted frozen N_{HI} , and the expected frozen Al/Si line states. A second validator restores each saved Sherpa session in CIAO and rechecks the actual dataset paths, final fit result, filters, grouping starts and gap, absorption, and forbidden-line parameters. Thus a valid JSON sidecar alone cannot promote a field into the scientific sample.

Twenty fields produced structurally valid Sherpa sessions, and all 20 passed the independent restore test. Eighteen passed the complete technical gate. Fields 0800731301 and 0800732201 have restorable sessions but non-positive final degrees of freedom. For 0800732601 and 0800732701, the final Sherpa result has `succeeded=False` with an “improper input parameters” message; the worker retained machine-readable failure records rather than publishing sessions.

A field enters the numerical sample if it passes technical validation, has positive degrees of freedom, has $Q \geq 0.01$, and has finite CGM_{sum} temperature and normalization with positive local errors. The primary sample additionally requires exactly one MOS1 and one MOS2 input. The strict sensitivity sample additionally

Table 3. Sample Accounting

Category	Number of fields
Selected pointings	22
Restorable production sessions	20
Complete technical validation	18
Numerically valid	15
Dual-MOS primary	14
Strict sensitivity sample	12

NOTE—The numerical sample includes the MOS2-only field 0800730101. The primary sample contains 10 North/NW and 4 South/SE fields.

requires at least ten degrees of freedom. The resulting counts are listed in Table 3.

Among the technically valid fields, 0800730401, 0800731001, and 0800732101 fail the $Q \geq 0.01$ numerical-quality criterion. Field 0800730101 passes the numerical gate but remains outside the primary sample because it has MOS2 only.

4.5. Derived fluxes and local uncertainties

After restoring each valid session, we extract the fitted CGM_{sum} temperature, APEC normalization, and their block of the saved Levenberg–Marquardt covariance matrix. For a band b , a unit-normalization absorbed `phabs*apec` model with the pointing-specific N_{HI} defines the conversion $C_b(kT)$. Because the fitted APEC normalization is already per arcmin², the component surface flux is

$$F_b = K_{\text{CGM}_{\text{sum}}} C_b(kT_{\text{CGM}_{\text{sum}}}). \quad (5)$$

The conversion is evaluated on a model energy grid with 0.001 keV spacing from 0.05 to 10 keV, using the same `angr`, `vern`, $Z = 0.3$, $z = 0$, and APEC 3.0.9 configuration as the fit.

We propagate the saved kT -normalization covariance with a first-order delta method,

$$\sigma_F^2 = \left(\frac{\partial F}{\partial kT}\right)^2 \sigma_{kT}^2 + \left(\frac{\partial F}{\partial K}\right)^2 \sigma_K^2 + 2\frac{\partial F}{\partial kT}\frac{\partial F}{\partial K}\text{Cov}(kT, K), \quad (6)$$

where the temperature derivative is evaluated by a centered finite difference with a step no smaller than 10^{-5} keV. These are local Gaussian statistical approximations. They do not include fixed- N_{HI} , foreground, detector-background, model-topology, or calibration uncertainty and are not profile-likelihood confidence intervals.

The primary reported quantity is the absorbed component-only 0.4–1.25 keV surface flux. This band lies wholly inside the directly fitted soft interval. We also calculate 0.5–2.0 keV for comparison with the literature, but label it secondary because 1.4–2.0 keV is supplied by model extrapolation across the excluded MOS interval.

4.6. Population summaries and side comparison

For temperature and flux separately, we summarize each side and the full primary sample with a normal random-effects model. For measurements x_i with local errors σ_i , the likelihood uses variances $\sigma_i^2 + \tau^2$; at fixed τ , the mean is weighted by $(\sigma_i^2 + \tau^2)^{-1}$. We determine the non-negative intrinsic scatter τ by maximizing the normal likelihood and quote $[\sum_i(\sigma_i^2 + \tau^2)^{-1}]^{-1/2}$ as the formal error on the mean. The reported heterogeneity diagnostic is the fixed-effect statistic

$$Q = \sum_i \frac{(x_i - \bar{x})^2}{\sigma_i^2}. \quad (7)$$

North-minus-South differences are formed from the two fitted side means, with their formal errors added in quadrature; the ratio and fractional asymmetry are propagated to first order from the same side-mean errors. For the current temperature and flux measurements, the maximum-likelihood extra scatter is zero. This means only that the available local errors do not require additional field-to-field scatter; it does not establish a physically uniform plasma.

The baseline side summaries do not include radius as a covariate. Because the two sides have different radial coverage, they are descriptive diagnostics rather than the final publication likelihood. The required shared radial term plus side offset, boundary-aware intervals, and controlled systematic branches are listed in Section 7.

5. RESULTS

5.1. A common soft spectral component

The 14 primary fields give a common temperature of

$$kT_{\text{CGM}_{\text{sum}}} = 0.17531 \pm 0.00436 \text{ keV}. \quad (8)$$

The heterogeneity statistic is $Q = 11.631$ for 13 degrees of freedom, corresponding to $p = 0.558$, and the fitted extra scatter is zero. The separate means are 0.17469 ± 0.00444 keV in the North/NW and 0.19314 ± 0.02382 keV in the South/SE. Their difference is -0.01845 ± 0.02423 keV. The data therefore support a common phenomenological temperature near 0.18 keV but do not establish that all fields contain one spatially uniform plasma.

5.2. North/NW and South/SE surface fluxes

The primary absorbed 0.4–1.25 keV surface fluxes are summarized in Table 4. The North/NW best fit is higher, but the difference is only 1.01σ . The fractional contrast is also consistent with zero. Figure 1 shows that the two sides have different radial sampling and that several individual uncertainties are large.

The numerically valid sample contains 15 fields after adding the single-MOS sensitivity case. Its North-minus-South flux difference is $0.198 \pm 0.210 10^{-15} \text{ erg cm}^{-2} \text{ s}^{-1} \text{ arcmin}^{-2}$, nearly identical to the primary result. The strict 12-field sample gives $0.147 \pm 0.274 10^{-15} \text{ erg cm}^{-2} \text{ s}^{-1} \text{ arcmin}^{-2}$. These checks do not turn the mild positive best fit into a significant asymmetry.

5.3. Effect of the corrected analysis

For the same 14 fields, the legacy v18 analysis gave a North-minus-South difference of $0.173 \pm 0.094 10^{-15} \text{ erg cm}^{-2} \text{ s}^{-1} \text{ arcmin}^{-2}$. The v19 revision changes the effective band, resets grouping across the MOS gap, removes unconstrained forbidden-band line components, and uses pointing-dependent HI4PI columns. The difference between v18 and v19 cannot be interpreted as one systematic shift because several corrections were applied together. The legacy result is retained only as a regression comparison.

6. DISCUSSION

6.1. The directly measured quantity is a line-of-sight sum

The approximately 0.18 keV component is widespread across the sampled footprint and has no detected excess

Table 4. Primary v19 Measurements with Local-Covariance Uncertainties

Quantity	North/NW	South/SE	North minus South
kT_{CGMsum} (keV)	0.17469 ± 0.00444	0.19314 ± 0.02382	-0.01845 ± 0.02423
Absorbed $F_{0.4-1.25}$ (10^{-15} erg cm $^{-2}$ s $^{-1}$ arcmin $^{-2}$)	1.106 ± 0.054	0.894 ± 0.203	0.212 ± 0.210

NOTE—The flux ratio is 1.238 ± 0.288 and the fractional asymmetry $(F_N - F_S)/[0.5(F_N + F_S)]$ is 0.212 ± 0.230 . Errors are propagated local covariance errors.




Figure 1. Primary v19 temperature (top) and absorbed 0.4–1.25 keV flux (bottom) versus signed projected radius. North/NW radii are positive and South/SE radii are negative. Filled blue and red circles are the homogeneous dual-MOS primary fields. The open triangle is the single-MOS sensitivity field. Gray crosses denote finite measurements that fail the numerical or technical selection; the two fail-closed sessions without valid measurements are absent. Error bars are local covariance approximations. The unequal radial coverage motivates the shared radial-plus-side model required for the submission analysis.

438 temperature scatter. This consistency does not deter- 447
 439 mine its distance. The Milky Way halo temperature 448
 440 distribution overlaps the current measurement (Henley 449
 441 & Shelton 2013a, 2015; Ponti et al. 2023; Locatelli et al. 450
 442 2024), and target-specific work toward M31 has explic- 451
 443 itly identified a similar combined foreground and target 452
 444 component (Kavanagh et al. 2020). At EPIC resolution, 453
 445 fitting two freely varying APEC components at nearly 454
 446 the same temperature would allow the optimizer to di-

vide one spectral shape arbitrarily. Such a decomposi-
 tion would not be a physical distance measurement.

Equation 1 is therefore the central interpretation rule. If Local-Group or bridge emission contributes, the observable becomes

$$S_{\text{obs}} = S_{\text{MW}} + S_{\text{M31}} + S_{\text{LG}}. \quad (9)$$

The additional term increases the physical ambiguity unless it is constrained by an external spatial prior. It

is therefore carried only as the explicitly named Local-Group branch of the conditional hierarchy in Section 6.3.

6.2. A mild best-fitting contrast, not an asymmetry detection

The North/NW best-fitting flux exceeds the South/SE value, but the 1.01σ difference is not compelling. A small observed contrast constrains only

$$\Delta S_{\text{obs}} = \Delta S_{\text{MW}} + \Delta S_{\text{M31}}, \quad (10)$$

with an analogous Local-Group term if required. It constrains M31 asymmetry only if the Milky Way contribution is assumed uniform over the footprint. Conversely, it constrains Milky Way variation only if the sampled M31 halo is assumed symmetric. Neither assumption follows from the current data.

The geometric imbalance is equally important. The primary sample contains 10 North/NW but only four South/SE fields, and their median radii differ by 4.31 kpc. A side-only mean can mix radial structure with azimuthal structure. We therefore treat the values in Table 4 as a descriptive empirical baseline. The next analysis must fit a common radial function and a side offset in one likelihood, with the same nuisance conventions on both sides.

6.3. Foreground context and conditional interpretations

The large Milky Way surveys provide population context but not a matched local foreground subtraction. The Henley & Shelton (2013a) continuum sample was restricted to $|b| > 30^\circ$, while M31 lies at $b = -21.57^\circ$; the nearest detected catalog field is approximately 29.7° away (Henley & Shelton 2013b). The M31-region oxygen-line measurements in the all-sky catalog are observed totals from M31 pointings and cannot be reused as a pure Milky Way foreground (Henley & Shelton 2012a,b). HaloSat and shadowing measurements further demonstrate angular structure and model dependence in the Galactic component (Henley & Shelton 2015; Henley et al. 2015; Kaaret et al. 2020).

External-galaxy stacks can instead define conditional M31 profile families. The eROSITA M31-mass stack has a projected profile of a similar order to the secondary 0.5–2.0 keV total measured here (Zhang et al. 2024a), but it is an absorption-corrected population mean with uncertain source subtraction and mass matching. Its scaling and star-forming/quiescent results further show that the inferred normalization depends on sample mass and classification (Zhang et al. 2024b, 2025). More massive spiral halos provide a useful shape precedent but not a direct M31 normalization (Anderson et al. 2016; Li et al. 2017, 2018). Any inferred M31 residual must

retain the name and uncertainty of the foreground and target-profile priors used to obtain it.

A useful conditional hierarchy is:

1. no external prior: report only the full line-of-sight sum and its spatial contrast;
2. uniform Milky Way prior: interpret ΔS_{obs} as a diluted M31 side contrast;
3. symmetric M31 prior: interpret ΔS_{obs} as Milky Way variation;
4. external Milky Way distribution: profile a non-negative M31 component and report a prior-labeled upper limit;
5. external galaxy profile: fit a prior-labeled M31 radial template plus a Milky Way residual.
6. Local-Group/bridge template: add a named non-negative S_{LG} term and report how it changes the conditional M31 and Milky Way ranges.

No conditional case should be presented as a direct decomposition.

7. ANALYSIS REQUIRED BEFORE SUBMISSION

The v19 baseline is sufficiently mature to support the manuscript structure, but five items remain submission blockers.

1. **Matched spatial likelihood.** Fit all accepted fields with a common radial dependence plus a North/NW–South/SE offset. Report the joint covariance, likelihood interval for the side offset, and sensitivity to radial parameterization. The raw side means cannot be the final headline statistic.
2. **Boundary-aware inference and optimizer checks.** Use multi-start optimization, profile the non-negative thermal normalization and side contrast, and calibrate one-sided coverage with simulations under representative nuisance conditions. Replace every local-covariance interval used for a final claim.
3. **Controlled systematic branches.** At minimum vary the mask/exposure-weighted absorption treatment, HI/dust or total-column prescription, soft-proton model, CXB normalization/cosmic variance, SWCX treatment, abundance, hot-component/model topology, and fitting band. Report branch shifts separately before defining any combined uncertainty model.

4. **Quantitative constraint product.** Report a calibrated interval on the side offset or fractional asymmetry and at least the conservative U0 upper envelope obtained by assigning the full cool line-of-sight component to M31. A named-prior U1 or smooth-foreground U2 M31 limit would substantially strengthen a full ApJ article, but must remain explicitly conditional.

5. **Submission documentation and figures.** Recover the original flare-screening thresholds, bind each extraction to its calibration index, and freeze the source-detection likelihood and mask-radius rule. Add a sky-footprint figure, representative rebinned spectra and residuals, the radial-plus-side fit, and a systematic-branch summary. Archive the current observation table, code, and machine-readable product inventory with persistent versioning.

An absolute M31-versus-Milky-Way decomposition is not a required blocker if the paper remains explicitly about the line-of-sight sum. It becomes a blocker only for claims of a detected M31 halo component or a foreground-subtracted M31 luminosity. A matched off-M31 control field or an external foreground prior would strengthen that second, conditional layer.

8. CONCLUSIONS

We have constructed a homogeneous XMM-Newton baseline for 22 sightlines through the 10–31 kpc projected inner halo of M31. The direct conclusions are:

1. Fourteen technically valid dual-MOS fields contain a common cool component with $kT = 0.17531 \pm 0.00436$ keV.
2. The absorbed 0.4–1.25 keV surface fluxes are $F_N = 1.106 \pm 0.054$ and $F_S = 0.894 \pm 0.203$, in units of 10^{-15} erg cm $^{-2}$ s $^{-1}$ arcmin $^{-2}$.
3. The North-minus-South difference is $0.212 \pm 0.210 10^{-15}$ erg cm $^{-2}$ s $^{-1}$ arcmin $^{-2}$, so the positive best fit is not a statistically compelling asymmetry detection.
4. The component is the unresolved line-of-sight sum of Milky Way and M31 emission, with a possible additional Local-Group term. EPIC spectra do not provide a unique distance decomposition.
5. Publication-level spatial constraints require a shared radial-plus-side likelihood, boundary-aware intervals, and a homogeneous systematic grid.

The stable result is therefore not a detection of M31's hot halo by itself. It is a reproducible measurement of the foreground-degenerate soft component toward M31 and a weak constraint on its large-scale spatial contrast.

DATA AND CODE AVAILABILITY

The XMM-Newton observations are available from the XMM-Newton Science Archive. The machine-readable measurement table, production manifest, fitting code, validation scripts, and figure-generation script will be archived with a persistent identifier before submission.

This work made use of XMM-Newton data and the HI4PI atomic hydrogen survey. Full acknowledgments, proposal information, and funding statements will be added after the author list is finalized.

REFERENCES

- Anderson, M. E., Churazov, E., & Bregman, J. N. 2016, *Monthly Notices of the Royal Astronomical Society*, 455, 227, doi: [10.1093/mnras/stv2314](https://doi.org/10.1093/mnras/stv2314)
- Arnaud, K. A. 1996, in *Astronomical Society of the Pacific Conference Series*, Vol. 101, *Astronomical Data Analysis Software and Systems V*, 17
- Foster, A. R., Ji, L., Smith, R. K., & Brickhouse, N. S. 2012, *The Astrophysical Journal*, 756, 128, doi: [10.1088/0004-637X/756/2/128](https://doi.org/10.1088/0004-637X/756/2/128)
- Freeman, P., Doe, S., & Siemiginowska, A. 2001, in *Proceedings of SPIE*, Vol. 4477, *Astronomical Data Analysis*, 76–87, doi: [10.1117/12.447161](https://doi.org/10.1117/12.447161)
- Gabriel, C., Denby, M., Fyfe, D. J., et al. 2004, in *Astronomical Society of the Pacific Conference Series*, Vol. 314, *Astronomical Data Analysis Software and Systems XIII*, 759
- Henley, D. B., & Shelton, R. L. 2012a, *The Astrophysical Journal Supplement Series*, 202, 14, doi: [10.1088/0067-0049/202/2/14](https://doi.org/10.1088/0067-0049/202/2/14)
- . 2012b, *VizieR Online Data Catalog: Diffuse O VII and O VIII Emission from XMM*, *VizieR Online Data Catalog J/ApJS/202/14*, doi: [10.26093/cds/vizier.22020014](https://doi.org/10.26093/cds/vizier.22020014)
- . 2013a, *The Astrophysical Journal*, 773, 92, doi: [10.1088/0004-637X/773/2/92](https://doi.org/10.1088/0004-637X/773/2/92)

- 632 —. 2013b, VizieR Online Data Catalog: XMM Survey of
633 Soft Background. III. Galactic Halo, VizieR Online Data
634 Catalog J/ApJ/773/92,
635 doi: [10.26093/cds/vizieR.17730092](https://doi.org/10.26093/cds/vizieR.17730092)
- 636 —. 2015, *The Astrophysical Journal*, 808, 22,
637 doi: [10.1088/0004-637X/808/1/22](https://doi.org/10.1088/0004-637X/808/1/22)
- 638 Henley, D. B., Shelton, R. L., Cumbee, R. S., & Stancil,
639 P. C. 2015, *The Astrophysical Journal*, 799, 117,
640 doi: [10.1088/0004-637X/799/2/117](https://doi.org/10.1088/0004-637X/799/2/117)
- 641 HI4PI Collaboration. 2016, *Astronomy & Astrophysics*,
642 594, A116, doi: [10.1051/0004-6361/201629178](https://doi.org/10.1051/0004-6361/201629178)
- 643 Jansen, F., et al. 2001, *Astronomy & Astrophysics*, 365, L1,
644 doi: [10.1051/0004-6361:20000036](https://doi.org/10.1051/0004-6361:20000036)
- 645 Kaaret, P., Koutroumpa, D., Kuntz, K. D., et al. 2020,
646 *Nature Astronomy*, 4, 1072,
647 doi: [10.1038/s41550-020-01215-w](https://doi.org/10.1038/s41550-020-01215-w)
- 648 Kavanagh, P. J., Sasaki, M., Breitschwerdt, D., et al. 2020,
649 *Astronomy & Astrophysics*, 637, A12,
650 doi: [10.1051/0004-6361/201937008](https://doi.org/10.1051/0004-6361/201937008)
- 651 Li, J.-T., Bregman, J. N., Wang, Q. D., Crain, R. A., &
652 Anderson, M. E. 2018, *The Astrophysical Journal*
653 *Letters*, 855, L24, doi: [10.3847/2041-8213/aab2af](https://doi.org/10.3847/2041-8213/aab2af)
- 654 Li, J.-T., Bregman, J. N., Wang, Q. D., et al. 2017, *The*
655 *Astrophysical Journal Supplement Series*, 233, 20,
656 doi: [10.3847/1538-4365/aa96fc](https://doi.org/10.3847/1538-4365/aa96fc)
- 657 Li, Z., & Wang, Q. D. 2007, *The Astrophysical Journal*
658 *Letters*, 668, L39, doi: [10.1086/522674](https://doi.org/10.1086/522674)
- 659 Locatelli, N., Ponti, G., Zheng, X., et al. 2024, *Astronomy*
660 *& Astrophysics*, 681, A78,
661 doi: [10.1051/0004-6361/202347061](https://doi.org/10.1051/0004-6361/202347061)
- 662 McConnachie, A. W., Irwin, M. J., Ferguson, A. M. N.,
663 et al. 2005, *Monthly Notices of the Royal Astronomical*
664 *Society*, 356, 979, doi: [10.1111/j.1365-2966.2004.08514.x](https://doi.org/10.1111/j.1365-2966.2004.08514.x)
- 665 Ponti, G., Zheng, X., Locatelli, N., et al. 2023, *Astronomy*
666 *& Astrophysics*, 674, A195,
667 doi: [10.1051/0004-6361/202243992](https://doi.org/10.1051/0004-6361/202243992)
- 668 Qu, Z., Huang, R., Bregman, J. N., & Li, J.-T. 2021, *The*
669 *Astrophysical Journal*, 907, 14,
670 doi: [10.3847/1538-4357/abc9b9](https://doi.org/10.3847/1538-4357/abc9b9)
- 671 Shirey, R., Soria, R., Borozdin, K., et al. 2001, *Astronomy*
672 *& Astrophysics*, 365, L195,
673 doi: [10.1051/0004-6361:20000243](https://doi.org/10.1051/0004-6361:20000243)
- 674 Takahashi, H., Okada, Y., Kokubun, M., & Makishima, K.
675 2004, *The Astrophysical Journal*, 615, 242,
676 doi: [10.1086/424373](https://doi.org/10.1086/424373)
- 677 Turner, M. J. L., et al. 2001, *Astronomy & Astrophysics*,
678 365, L27, doi: [10.1051/0004-6361:20000087](https://doi.org/10.1051/0004-6361:20000087)
- 679 Ueda, M., Sugiyama, H., Kobayashi, S. B., et al. 2022,
680 *Publications of the Astronomical Society of Japan*, 74,
681 1396, doi: [10.1093/pasj/psac077](https://doi.org/10.1093/pasj/psac077)
- 682 Zhang, Y., Comparat, J., Ponti, G., et al. 2024a,
683 *Astronomy & Astrophysics*, 690, A267,
684 doi: [10.1051/0004-6361/202449412](https://doi.org/10.1051/0004-6361/202449412)
- 685 —. 2024b, *Astronomy & Astrophysics*, 690, A268,
686 doi: [10.1051/0004-6361/202449413](https://doi.org/10.1051/0004-6361/202449413)
- 687 —. 2025, *Astronomy & Astrophysics*, 693, A197,
688 doi: [10.1051/0004-6361/202452273](https://doi.org/10.1051/0004-6361/202452273)