The Astrophysical Journal, 984:190 (12pp), 2025 May 9 https://doi.org/10.3847/1538-4357/adc450
© 2025. The Author(s). Published by the American Astronomical Society.
OPEN ACCESS

X-Raying CAMELS: Constraining Baryonic Feedback in the Circumgalactic Medium with the CAMELS Simulations and eRASS X-Ray Observations

Erwin T. Lau[1,2], Daisuke Nagai[3,4], Ákos Bogdán[1], Isabel Medlock[4], Benjamin D. Oppenheimer[5], Nicholas Battaglia[6], Daniel Anglés-Alcázar[7], Shy Genel[8,9], Yueying Ni[1], and Francisco Villaescusa-Navarro[8]
[1] Center for Astrophysics—Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA; erwin.lau@cfa.harvard.edu
[2] University of Miami, Department of Physics, Coral Gables, FL 33124, USA
[3] Department of Physics, Yale University, New Haven, CT 06520, USA
[4] Department of Astronomy, Yale University, New Haven, CT 06520, USA
[5] Center for Astrophysics and Space Astronomy, University of Colorado, 389 UCB, Boulder, CO 80309, USA
[6] Department of Astronomy, Cornell University, 404 Space Sciences Building, Ithaca, NY 14853, USA
[7] Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269, USA
[8] Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave, New York, NY 10010, USA
[9] Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027, USA
Received 2024 December 5; revised 2025 January 28; accepted 2025 February 14; published 2025 May 9

Abstract

The circumgalactic medium (CGM) around massive galaxies plays a crucial role in regulating star formation and feedback. Using the Cosmology and Astrophysics with MachinE Learning Simulations (CAMELS) suite, we develop emulators for the X-ray surface brightness profile and the X-ray luminosity–stellar mass scaling relation, to investigate how stellar and active galactic nucleus (AGN) feedback shape the X-ray properties of the hot CGM. Our analysis shows that at CGM scales ($10^{12} \lesssim M_{\rm halo}/M_{\odot} \lesssim 10^{13}$, $10 \lesssim r\,{\rm kpc}^{-1} \lesssim 400$), stellar feedback more significantly impacts the X-ray properties than AGN feedback within the parameters studied. Comparing the emulators to recent eROSITA All Sky Survey (eRASS) observations, it is found that stronger feedback than is currently implemented in the IllustrisTNG, SIMBA, and Astrid simulations is required to match the observed CGM properties. However, adopting these enhanced feedback parameters causes deviations in the stellar mass–halo mass relations from observational constraints below the group-mass scale. This tension suggests possible unaccounted-for systematics in X-ray CGM observations or inadequacies in the feedback models of cosmological simulations.

Unified Astronomy Thesaurus concepts: Circumgalactic medium (1879); Hydrodynamical simulations (767); X-ray astronomy (1810); Galactic and extragalactic astronomy (563)

1. Introduction

The circumgalactic medium (CGM) around massive galaxies plays a fundamental role in galaxy evolution (see J. Tumlinson et al. 2017; C.-A. Faucher-Giguère & S. P. Oh 2023 for reviews). It serves as an important reservoir for star formation and a repository for metal-enriched outflows driven by feedback from supernovae (SNe) and active galactic nuclei (AGNs). In particular, the hot phase of the CGM can serve as a unique probe of baryon cycles in massive galaxies (e.g., see M. Donahue & G. M. Voit 2022 for a review). Constraining and understanding the physical properties of the CGM is thus crucial for understanding galaxy evolution and structure formation in the Universe.

Multiwavelength observations have revolutionized our understanding of CGM physical properties, particularly through X-ray observations (e.g., Á. Bogdán et al. 2013a, 2013b, 2017; M. E. Anderson et al. 2015) and the Sunyaev–Zeldovich (SZ) effects (e.g., S. Amodeo et al. 2021; J. N. Bregman et al. 2022; S. Das et al. 2023). Despite its relatively low angular resolution compared to Chandra and XMM-Newton, X-ray observations with the eROSITA/SRG telescope have advanced the state of CGM observations in X-ray through its large sky coverage. Recent stacked X-ray observations from the eROSITA Final Exposure Depth Survey (eFEDS) fields (U. Chadayammuri et al. 2022; J. Comparat et al. 2022) have inferred that the CGM X-ray luminosity seems to be lower than that of cosmological simulations at higher-stellar-mass halos. However, these results are still uncertain, given the statistical uncertainties due to cosmic variances in the stacking procedures. The recently released stacked X-ray surface brightness (XSB) profiles and the X-ray luminosities from the half-sky eROSITA All Sky Survey (eRASS) from Y. Zhang et al. (2024a, 2024b) represent the most state-of-the-art X-ray observations of the CGM to date. Compared to the 140 deg^2 eFEDS field, the half-sky coverage of eRASS contains a much larger galaxy sample for stacking, which increases the signal-to-noise ratio by a factor of a few, reducing the statistical uncertainty on the stacked signal, thus potentially allowing for better constraints on CGM properties and their physics.

The Cosmology and Astrophysics with MachinE Learning Simulations (CAMELS) suite (F. Villaescusa-Navarro et al. 2021a, 2021b; Y. Ni et al. 2023), which consists of a large number of cosmological hydrodynamical simulations with varying cosmological parameters and baryonic feedback, is uniquely suited for studying feedback physics at the CGM scale. The CAMELS suite has been used to study the impact of feedback on CGM properties in SZ (E. Moser et al. 2022; S. Pandey et al. 2023; D. Wadekar et al. 2023) and in radio through fast radio bursts (FRBs; I. Medlock et al. 2024). CAMELS also enables the development of machine learning techniques for predicting CGM properties with X-ray and H I maps (N. Gluck et al. 2024), characterizing the impact of

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

feedback on matter power clustering (A. M. Delgado et al. 2023; M. Gebhardt et al. 2024), and breaking the degeneracies between feedback and cosmological parameters via kinetic SZ and FRBs (A. Nicola et al. 2022). Zoomed-in simulations from CAMELS (M. E. Lee et al. 2024) also predict how X-ray and SZ profiles of clusters and groups can be used to constrain cosmology and feedback physics (E. Hernández-Martìnez et al. 2025).

In this paper, we explore the dependence of CGM X-ray properties on feedback models using CAMELS and use the latest eRASS observations to constrain the CAMELS feedback parameters. We study how different subgrid models of feedback impact the X-ray CGM emission and assess how the latest stacked XSB and X-ray luminosities from eRASS are able to constrain CGM feedback physics in CAMELS. To do so, we follow the approach of E. Moser et al. (2022), by constructing emulators of the XSB profiles and X-ray luminosity from CAMELS as a function of the feedback parameters. By comparing the X-ray-emulated observables to the eRASS data, we derive constraints on the feedback parameters.

2. Methodology

2.1. Simulations

The CAMELS data set (F. Villaescusa-Navarro et al. 2021a, 2021b; Y. Ni et al. 2023) consists of cosmological simulations with different ΛCDM cosmologies, hydrosolvers, and subgrid feedback physics. Each CAMELS simulation we use follows the evolution of a $(25h^{-1}$ comoving Mpc$)^3$ volume with base parameters $\Omega_b = 0.049$, $h = 0.6711$, and $n_s = 0.9624$, assuming a flat ΛCDM Universe. The other two main parameters Ω_M and σ_8 are varied in the simulations. All simulations follow the evolution of 256^3 dark matter (DM) particles and fluid elements, corresponding to a DM particle resolution of $M_{\rm DM} = 6.49 \times 10^7 (\Omega_M - \Omega_b)/0.251\, h^{-1} M_\odot$ and a gas-mass resolution of $1.25 \times 10^7 h^{-1} M_\odot$.

2.1.1. Simulation Suites

These simulations were run using different simulation codes. For this work, we focus on IllustrisTNG (D. Nelson et al. 2018; A. Pillepich et al. 2018), SIMBA (R. Davé et al. 2019), and Astrid (S. Bird et al. 2022; Y. Ni et al. 2022), which have been run using the AREPO, GIZMO, and MP-Gadget codes, respectively. We used the Latin Hypercube (LH) sets for each of the three runs: each consists of 1000 simulations with different parameters sampled using LH. These parameters include Ω_m and σ_8, two parameters corresponding to SNe feedback, and another two corresponding to AGN feedback. There are also fundamental differences between the different simulation suites. In IllustrisTNG, the hydrodynamic equations are solved on a moving mesh; in SIMBA, they are solved using meshless methods; and in Astrid, they are solved using the pressure–entropy formulation of smoothed particle hydrodynamics. Additionally, each simulation suite has different subgrid physics; for example, SIMBA tracks dust grains, while IllustrisTNG includes magnetohydrodynamics.

We focus on the four key feedback parameters explored in the CAMELS suites: two parameters for stellar feedback—SN1 and SN2; and two parameters for AGN feedback—AGN1 and AGN2. Here, we give a qualitative overview of the differences between the feedback implementations of the different simulation suites. We refer the reader to Y. Ni et al. (2023) and I. Medlock et al. (2025) for more detailed comparisons of the models. Here, we provide a qualitative overview of how the three simulation suites implement their stellar and black hole feedback and highlight the most salient differences between them. To avoid confusion between the different feedback physics in the three simulation suites, we hereafter denote the feedback parameters of the CAMELS–IllustrisTNG run as $A_{\rm SN1}$, $A_{\rm SN2}$, $A_{\rm AGN1}$, and $A_{\rm AGN2}$; of the CAMELS–SIMBA run as $B_{\rm SN1}$, $B_{\rm SN2}$, $B_{\rm AGN1}$, and $B_{\rm AGN2}$; and of the CAMELS–Astrid run as $C_{\rm SN1}$, $C_{\rm SN2}$, $C_{\rm AGN1}$, and $C_{\rm AGN2}$. For all parameters, a value of 1 represents the fiducial value that was adopted in the original simulations.

2.1.2. Stellar Feedback Models

All three suites implement stellar feedback as galactic winds. For CAMELS–IllustrisTNG, the $A_{\rm SN1}$ parameter represents a prefactor that controls the energy per unit star formation rate (SFR); for CAMELS–SIMBA, $B_{\rm SN1}$ represents the wind-mass outflow rate per unit SFR, i.e., the mass-loading factor; and for CAMELS–Astrid, $C_{\rm SN1}$ controls the energy per unit SFR, the same as for CAMELS–IllustrisTNG. In all three suites, the SN2 parameter represents the normalization factor for the galactic wind speed. Consequently, the stellar feedback model is essentially the same for CAMELS–IllustrisTNG and CAMELS–Astrid, and slightly different for CAMELS–SIMBA. Specifically, keeping SN1 constant while varying the SN2 parameter maintains a fixed wind energy output but changes the wind speed for both CAMELS–IllustrisTNG and CAMELS–Astrid; for CAMELS–SIMBA, though, while the mass-loading factor remains constant for a fixed SN1, changing SN2 varies both the wind speed and the wind energy.

2.1.3. AGN Feedback Models

The AGN feedback models differ substantially between the three simulation suites. In CAMELS–IllustrisTNG, the AGN feedback is implemented as kinetic feedback from supermassive black holes (SMBHs) seeded onto massive halos. The SMBHs in IllustrisTNG accrete assume a spherical Bondi accretion model. The AGN feedback mode switches from thermal to kinetic modes from high to low SMBH accretion rates. The $A_{\rm AGN1}$ parameter controls the energy output of the AGN feedback per unit SMBH accretion in the kinetic mode. Kinetic feedback is implemented by injecting kinetic energy around the SMBH in a random direction. The kinetic energy is injected until a minimum amount of energy is accumulated, to ensure that the energy injection events are powerful and independent of each other. This minimum amount of energy is controlled by the $A_{\rm AGN2}$ parameter. Thus, $A_{\rm AGN2}$ can be thought of as controlling the "burstiness" in the kinetic feedback of the AGN in IllustrisTNG.

On the other hand, CAMELS–SIMBA implements AGN feedback differently from IllustrisTNG. The SMBH in SIMBA has two accretion modes: Bondi accretion and gravitational torque accretion. The kinetic feedback mode in CAMELS–SIMBA switches from wind/QSO mode to jet mode from high to low SMBH accretion rates. The $B_{\rm AGN1}$ parameters control the momentum flux of the feedback per unit SMBH accretion for both jet and wind modes, and $B_{\rm AGN2}$ controls the speed of the jet.

Table 1
Parameters Associated with Different Simulations and Feedback Mechanisms

Parameter	IllustrisTNG	SIMBA	Astrid
SN1	Galactic winds: energy per unit SFR [0.25, 4.0].	Galactic winds: mass loading [0.25, 4.0].	Galactic winds: energy per unit SFR [0.25, 4.0].
SN2	Galactic winds: wind speed [0.5, 2.0].	Galactic winds: wind speed [0.5, 2.0].	Galactic winds: wind speed [0.5, 2.0].
AGN1	Kinetic mode BH feedback: energy per unit BH accretion [0.25, 4.0].	QSO-and-jet-mode BH feedback: momentum flux [0.25, 4.0].	Kinetic mode BH feedback: energy per unit BH accretion [0.25, 4.0].
AGN2	Kinetic mode BH feedback: ejection speed / burstiness [0.5, 2.0].	Jet-mode BH feedback: jet speed [0.5, 2.0].	Thermal mode BH feedback: energy per unit BH accretion [0.25, 4.0].

For CAMELS–Astrid, the accretion rate onto the SMBH is estimated via Bondi accretion. The SMBH feedback also follows a two-mode approach that switches from thermal to kinetic feedback modes from high to low SMBH accretion rates. $C_{\rm AGN1}$ and $C_{\rm AGN2}$ modulate the efficiency of the kinetic and thermal feedback separately. $C_{\rm AGN1}$ has the same physical meaning as $A_{\rm AGN1}$ in CAMELS–IllustrisTNG. Compared to CAMELS–IllustrisTNG, CAMELS–Astrid turns on the AGN kinetic feedback at a lower SMBH accretion-rate threshold with a lower upper limit of feedback efficiency, resulting in milder AGN kinetic feedback compared to CAMELS–IllustrisTNG.

In Table 1, we summarize the physical meanings of the four feedback parameters in CAMELS–IllustrisTNG, CAMELS–SIMBA, and CAMELS–Astrid, and the ranges of the values adopted in the CAMELS runs.

2.1.4. Cosmological Parameters

Finally, for the two cosmological parameters, Ω_M varies between 0.1 and 0.5, and σ_8 varies between 0.6 and 1.0. The fiducial values for the feedback parameters are 1.0 for Ω_M and 0.3 and 0.8 for σ_8, respectively.

2.2. Modeling X-Ray Observables

The XSB profile as a function of projected radius $R_p = [10, 10^3]$ kpc, binned in 20 uniform logarithmic bins, for each halo is computed from the simulation as

$$S_{\rm X}(R_p) = 2\int_{R_p}^{\infty} \frac{\epsilon(r) r}{\sqrt{r^2 - R_p^2}} dr, \tag{1}$$

where

$$\epsilon(r) = 4\pi \int n_e(r) n_{\rm H}(r) r^2 \Lambda(T(r), Z(r)) dr \tag{2}$$

is the 3D X-ray emissivity of the CGM as a function of the distance from the halo center r; n_e and $n_{\rm H}$ are the electron and hydrogen densities, respectively; and Λ is the X-ray cooling function, which is a function of the CGM temperature T and metallicity Z, computed using the APEC model (A. R. Foster et al. 2012). For each halo, we compute the densities, temperatures, and metallicity profiles (assuming spherical symmetry) for the hot-phase gas with temperature $>10^5$ K. The energy range of the emission in the observer's frame is chosen to be $E = [0.5, 2.0]$ keV. The X-ray profiles are computed using the `CGM Toolkit` code (E. T. Lau 2025a). Note that at the halo-mass scales in which we are interested ($\log_{10} M/M_\odot < 13.5$), metal lines dominate the X-ray emission (see L. Lovisari et al. 2021).

We compute the X-ray luminosity by integrating the XSB profile up to R_{500c}:

$$L_{\rm X}(<R_{500c}) = \int_0^{R_{500c}} S_{\rm X}(R_p) 2\pi R_p dR_p. \tag{3}$$

2.3. Emulation

We follow the same approach to constructing a Gaussian process emulator for the XSB profile as in E. Moser et al. (2022), where an emulator for SZ profiles was developed to capture the dependence of feedback physics. The emulator can be thought of as a multidimensional interpolator of the XSB profile.

2.3.1. XSB Profiles

We construct emulators for the XSB profiles over a total of eight parameters: the four feedback parameters; the two cosmological parameters—Ω_M and σ_8; and two halo parameters—the redshift z and stellar mass $M_\star$ of the central subhalo. For each of the 1000 LH simulations with different cosmological and feedback parameters, we compute the mean XSB profiles over the three available snapshots corresponding to $z = [0.10, 0.15, 0.18]$ that match the redshift range of the eRASS CGM data, with six uniformly spaced stellar mass bins spanning the range $\log_{10}(M_\star/M_\odot) \in [9.0, 12.0]$, covering the range of stellar masses in the eRASS CGM observations $\log_{10}(M_\star/M_\odot) \in [10.5, 11.5]$. The resulting stellar mass bins have a large enough number of halos per bin to have stable trends in profiles—changing the minimum and maximum halo mass in the binning scheme to $\log_{10}(M_\star/M_\odot) \in [10.5, 11.5]$ does not change the prediction of the emulator. The small box size ($25\,h^{-1}$ Mpc) of CAMELS limits our mass selections to objects with halo masses $M_{200c} \lesssim 10^{13} M_\odot$, and the higher-mass bins have significantly fewer objects than the lower-mass bins. The numbers of halos in the bins range from 30 to 843, depending on the specific run and snapshot. This results in 9000 (1000 runs times three redshifts times three mass bins) mean XSB profiles as inputs to the emulation. Following the SZ emulator of E. Moser et al. (2022), each input XSB profile is represented in basis vectors, and the weights of the basis vectors are decomposed using principal component analysis (PCA). The weights are interpolated using a radial basis function interpolator. Principal components are sorted by decreasing variance, and the number of PCA components is a free parameter, which impacts the amplitude of residual variance. We used 12 PCA components and also checked that adding more components does not affect the emulator results. The PCA components are then emulated over the parameter space of the LH set using the Gaussian process.

2.3.2. X-Ray Luminosity–Stellar Mass Relation

We also emulate the X-ray luminosity–stellar mass $L_X - M_\star$ relation. Similar to the XSB profile emulation, we compute the L_X values for halos in the LH simulation for the same redshift snapshots. The difference with the XSB emulation is that we emulate L_X in eight logarithmic stellar mass $\log_{10}(M_\star/M_\odot) \in [9, 12]$ bins, compared to the six stellar mass bins in XSB. This enables us to better capture the local changes in the slope of the scaling relation. We perform emulation in the same way as in the XSB profile: we decompose each relation into 12 PCA components, which are then emulated over the parameter space using the same Gaussian process as the XSB profiles. We have checked that the L_X and XSB emulators are consistent with each other, as we can recover the same L_X values by integrating the XSB-emulated profile for a given stellar mass.

Note that given the relatively small box size, for some LH runs, there are very few halos in the highest-stellar-mass bin, especially when the feedback values are extremely high or low. This could result in zero values for the emulated XSB profiles or L_X, breaking the emulator. To circumvent this, we add a small number (10^{-30}) to the zero values. We have checked that this does not affect the accuracy of the emulator.

The `CAMELS CGM Emulator` code (E. T. Lau 2025b) used in this work is publicly available on GitHub.[10]

3. Observation Data from eRASS

The eROSITA survey with its half-sky coverage overlaps with large-scale galaxy surveys such as the Sloan Digital Sky Survey (SDSS). This uniquely enables the stacking of X-ray emission on a large number ($\sim 10^5$) of optically identified galaxies, resulting in high-signal-to-noise-ratio measurements of the X-ray CGM.

We use the "CEN" sample from the eRASS CGM papers (Y. Zhang et al. 2024a, 2024b) for comparison. The sample is based on the spectroscopic galaxy sample of the SDSS Main Galaxy Sample. The sample for the stacked XSB profiles consists of three main bins based on stellar mass: "MW," "M31," and "2M31," corresponding to the median stellar mass of $\log_{10}(M_{\star,\mathrm{med}}/M_\odot) = 10.74, 11.11, 11.34$, with median redshifts $z_{\mathrm{med}} = 0.08, 0.12$, and 0.15 and the number of stacked galaxies = 30,825, 26,099, and 20,342, respectively. The X-ray luminosity–stellar mass scaling relation data include an additional lower-stellar-mass bin of $\log_{10}(M_{\star,\mathrm{med}}/M_\odot) = 10.26$, with $z_{\mathrm{med}} = 0.05$ and 7956 galaxies. We refer the reader to Y. Zhang et al. (2024a, 2024b) for details of the data reduction, sample selection, and a discussion of the systematics of the observations.

4. Results

4.1. XSB Profiles

In this section, we demonstrate how the XSB profiles from CAMELS–IllustrisTNG, CAMELS–SIMBA, and CAMELS–Astrid depend on the feedback parameters and how they compare with the eRASS observations.

[10] https://github.com/ethlau/CAMELS_emulator

4.1.1. CAMELS–IllustrisTNG

The top row in Figure 1 shows the comparisons of the eRASS CGM profiles with the emulated XSB profiles based on the CAMELS–IllustrisTNG runs in the three stellar mass bins. Among all feedback parameters, the XSB profile is most dependent on A_{SN1}, which represents the SN energy imparted to the surrounding gas per unit SFR. Increasing A_{SN1} suppresses the AGN feedback by reducing the gas available for accretion onto the SMBHs in the halo centers, as shown in M. E. Lee et al. (2024) and I. Medlock et al. (2025), thus increasing the overall gas density. We have checked the CGM temperature also increases with A_{SN1}, but not as much as gas density. As the XSB depends on the density squared, and less so on temperature, the increase in gas density leads to higher XSB. In terms of mass dependence, the XSB profile in the least massive "MW" bin is more sensitive to variations in A_{SN1} compared to the other two stellar mass bins, due to the shallower gravitational potential of MW-size halos.

The second row shows the comparisons of the eRASS CGM XSB profiles with the emulated XSB profiles with varying A_{SN2} feedback, which represents the speed of the galactic wind. Increasing A_{SN2} leads to an overall increase in XSB, particularly in the radial range [10, 300] kpc, for all stellar mass bins. Increasing the SN wind speed also suppresses the AGN feedback by decreasing the available gas for SMBH accretion, but not enough to drive the hot gas out of the halo potential; therefore, the hot gas density, and thus the XSB, increase with increasing SN wind speed. In addition to the amplitude of the XSB profile, the shape of the eRASS XSB profile matches well with the emulated profiles with higher A_{SN2}, which show flatter profiles in [10, 300] kpc. The impact of A_{SN2} is more prevalent in the lowest-stellar-mass "MW" bin with shallower gravitational potential.

The third row shows the comparisons of the eRASS CGM profiles with the emulated XSB profiles with varying A_{AGN1} feedback in CAMELS–IllustrisTNG. A_{AGN1} represents the power of the AGN feedback in the kinetic mode. The fiducial value $A_{\mathrm{AGN1}} = 1$ (while other feedback parameters are kept fixed at their respective fiducial values) is able to match the observed eRASS profiles in the MW-stellar-mass bin, but slightly overpredicts the profiles in the other two higher-stellar-mass bins. Varying A_{AGN1} has relatively little impact on the XSB profile across all mass bins, with the least massive "MW" bins being affected the least. For the highest-mass bin "2M31," increasing A_{AGN1} changes the shape of the XSB profile, making the profile flatter outside the core.

The fourth row shows the comparisons of the eRASS CGM profiles with the emulated XSB profiles with varying A_{AGN2} feedback, which represents the "burstiness" of the AGN feedback. Specifically, a higher A_{AGN2} means less frequent but more powerful AGN feedback events. A higher A_{AGN2} leads to a decrease in amplitude in the emulated XSB profiles for all stellar mass bins, as the less frequent but more powerful AGN feedback is more capable of pushing the gas out of the potential well, reducing the gas density and thus the X-ray emission. The impact of A_{AGN2} is more prevalent in the highest-stellar-mass bin "2M31."

4.1.2. CAMELS–SIMBA

Figure 2 shows the comparisons of the eRASS CGM profiles with the emulated XSB profiles based on the CAMELS–SIMBA runs.

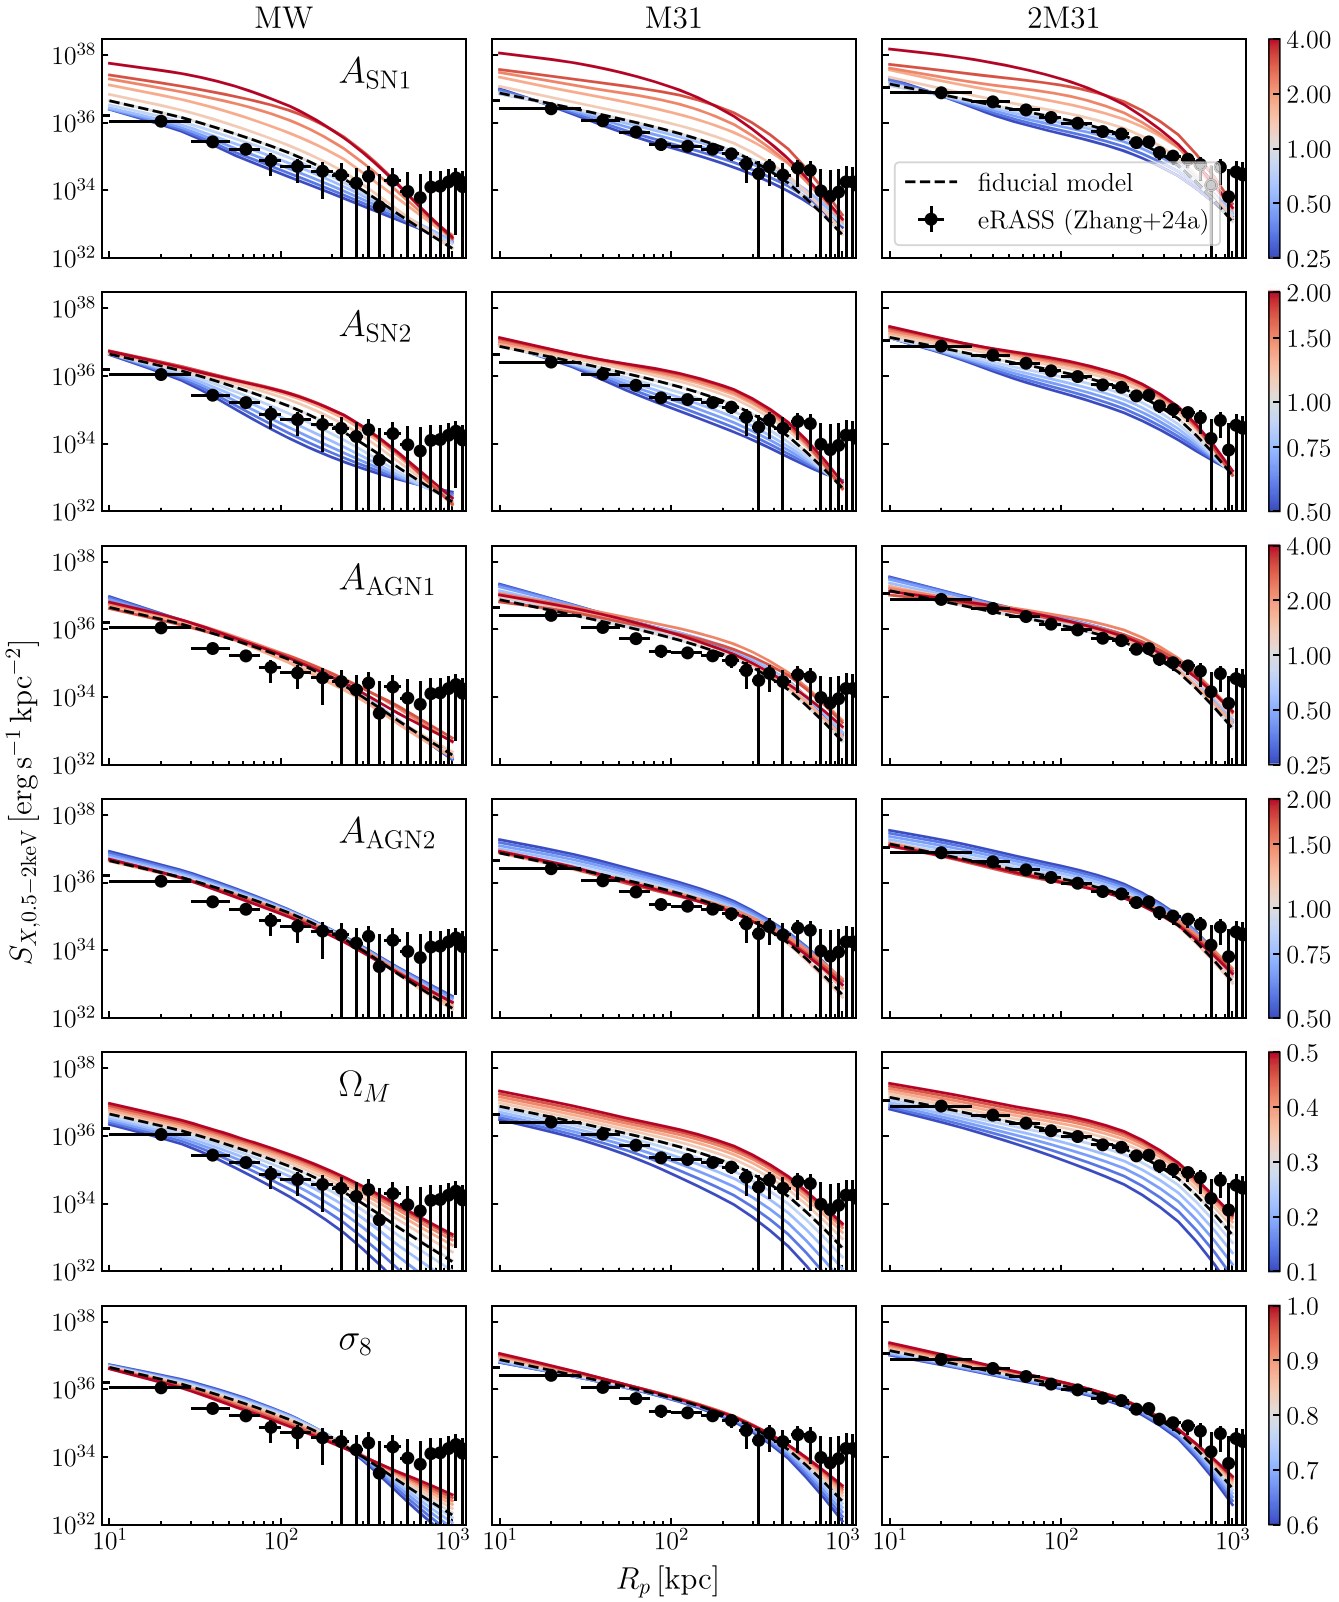

Figure 1. Comparisons of the XSB profiles from eRASS CGM in black data points with 1σ error bars (Y. Zhang et al. 2024a) and the emulated profiles for CAMELS–IllustrisTNG with varying values for the four feedback parameters—$A_{\rm SN1}$, $A_{\rm SN2}$, $A_{\rm AGN1}$, and $A_{\rm AGN2}$—are shown in the panels from the first to fourth rows. The colors indicate the strengths of the feedback parameters, which vary from [0.25, 4.0] for $A_{\rm SN1}$ and $A_{\rm AGN1}$ and [0.5, 2.0] for $A_{\rm SN2}$ and $A_{\rm AGN2}$. The bottom two rows show the dependence on Ω_M and σ_8. The black dashed lines indicate where the parameters adopt their fiducial values. The different columns show the comparisons for the three different stellar mass bins: "MW," "M31," and "2M31," in the left, middle, and right panels, respectively.

The top row shows the XSB dependence on $B_{\rm SN1}$, the energy of the SN feedback per unit stellar mass, i.e., the mass-loading factor. Similar to CAMELS–IllustrisTNG, the CAMELS–SIMBA XSB profile is dependent on the SN energy. However, increasing $B_{\rm SN1}$ leads to a decrease in XSB, mainly in the outer CGM. Increasing the mass-loading factor increases the amount

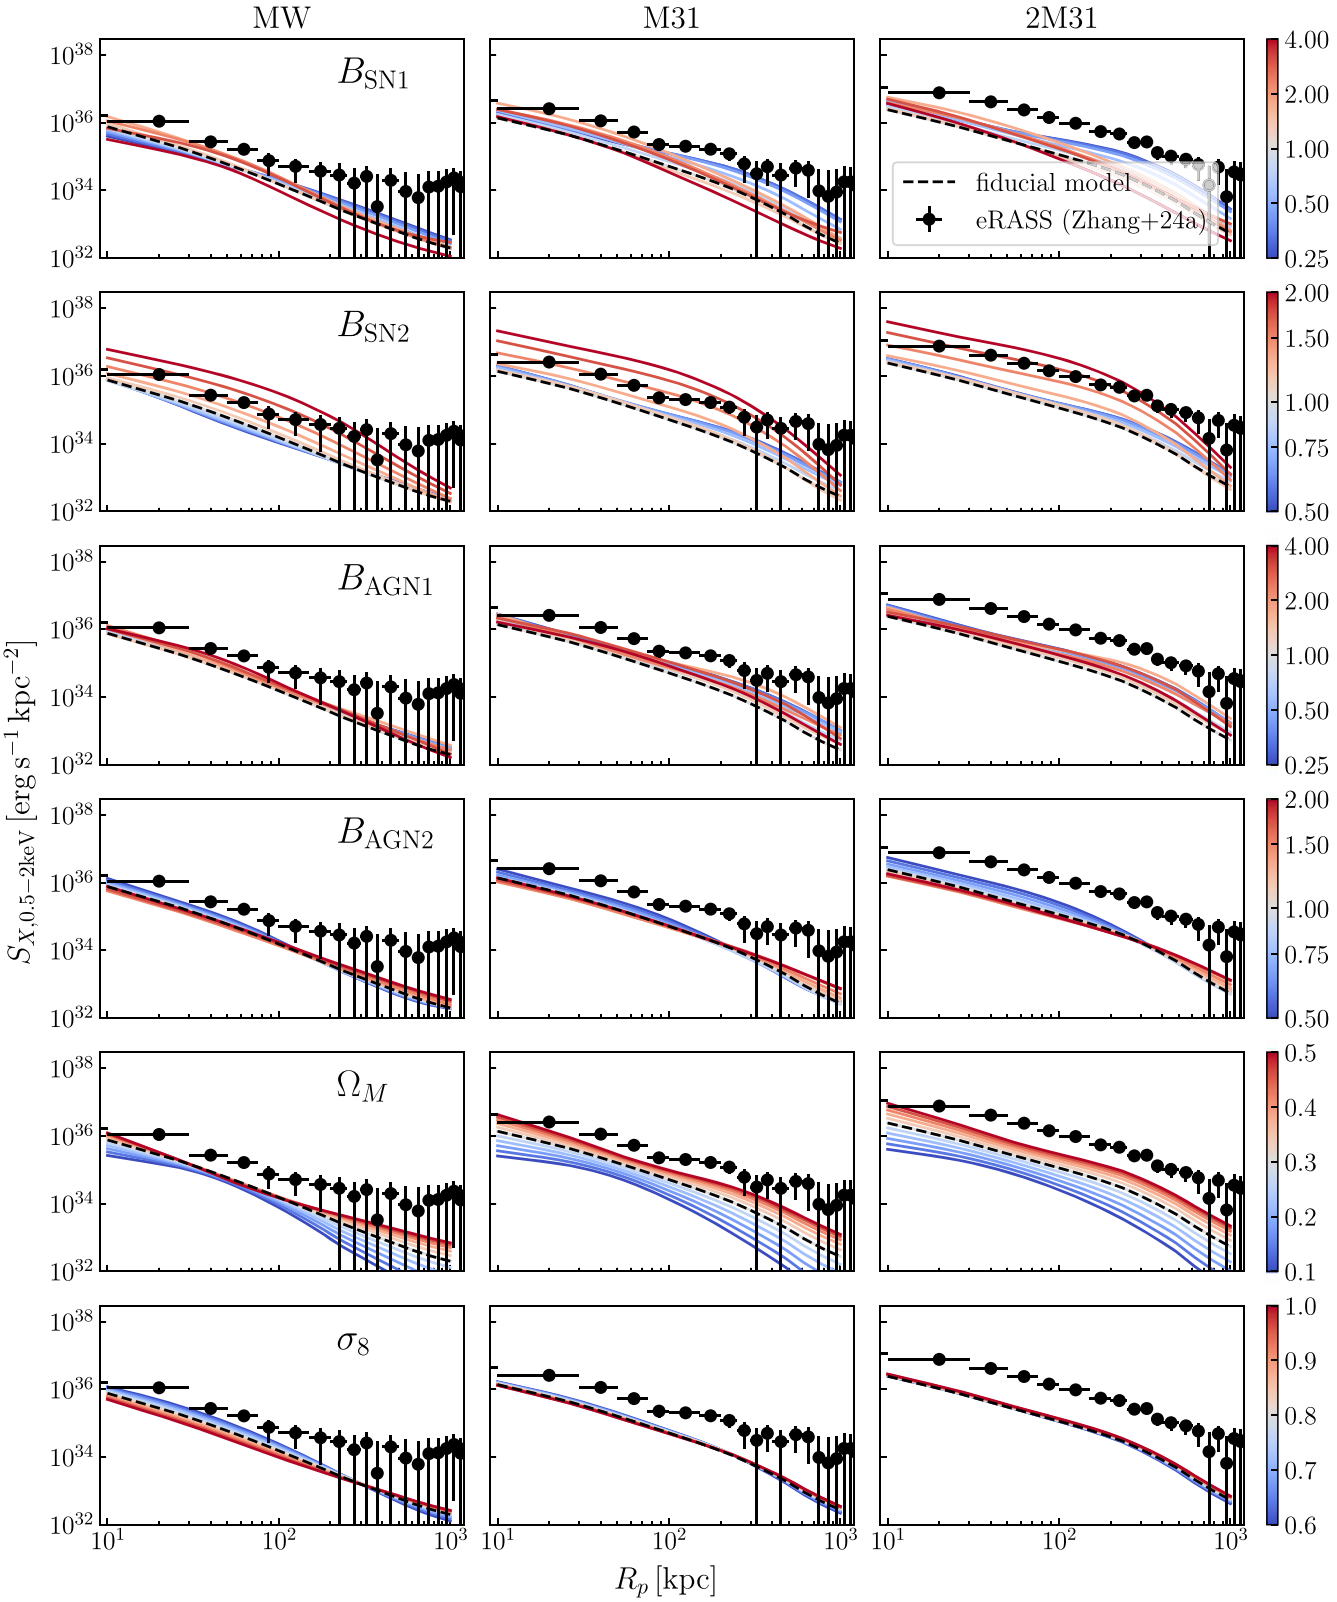

Figure 2. The same as Figure 1 but for CAMELS–SIMBA.

of gas being pushed out by the galactic wind. This is contrary to the $A_{\rm SN1}$ dependence in CAMELS–IllustrisTNG, where $A_{\rm SN1}$ represents the wind energy instead of the wind-mass outflow. With other parameters fixed at the fiducial value of 1, varying $B_{\rm SN1}$ alone systematically underpredicts the eRASS observations, especially in the lowest-stellar-mass bin.

The second row in Figure 2 shows the comparison of the eRASS CGM XSB profiles with the SIMBA profiles of varying $B_{\rm SN2}$, which represents the speed of the SN wind, similar to $A_{\rm SN2}$ in CAMELS–IllustrisTNG. The SIMBA XSB profiles are most sensitive to $B_{\rm SN2}$. In contrast to $B_{\rm SN1}$, increasing $B_{\rm SN2}$, i.e., the SN wind speed, leads to a higher XSB. This is likely

because increasing the SIMBA SN wind inhibits the growth of the SMBH, leading to suppressed AGN feedback and thus to the retention of more CGM gas (I. Medlock et al. 2025). To match the eRASS XSB profiles, the SIMBA wind speed needs to be higher than the fiducial model value.

The third row in Figure 2 shows the comparisons of the eRASS CGM profiles with the emulated XSB profiles with varying SIMBA $B_{\rm AGN1}$, the feedback energy of the AGN. Similar to CAMELS–IllustrisTNG, the XSB profile is not very sensitive to the AGN feedback energy. Overall, varying $B_{\rm AGN1}$ alone does not lead to any agreement between the SIMBA and the eRASS XSB profiles.

The fourth row of Figure 2 shows the comparison of the eRASS CGM XSB profiles with the SIMBA profiles with varying $B_{\rm AGN2}$, which represents the speed of the AGN jet. Increasing $B_{\rm AGN2}$ results in a decrease in XSB in the inner CGM and a slight increase in the outer CGM for all stellar mass bins. This means that the jet is more effective in lowering the gas density in the inner regions while transferring some gas mass to the outer regions of halos.

4.1.3. CAMELS–Astrid

Figure 3 shows the comparisons of the eRASS CGM profiles with the XSB profiles emulated based on the CAMELS–Astrid runs.

The top row shows the XSB dependence on $C_{\rm SN1}$, the energy of the SN feedback per unit SFR. The Astrid XSB profile is more sensitive to $C_{\rm SN1}$. Increasing $C_{\rm SN1}$ leads to a large increase in the normalization of XSB, qualitatively similar to CAMELS–Illustris. With other parameters fixed at the fiducial value of 1, varying $C_{\rm SN1}$ alone always overpredicts the eRASS observations for all stellar mass bins.

The second row shows the comparison of the eRASS CGM profile with the CAMELS–Astrid profile with varying $C_{\rm SN2}$, which represents the SN wind speed. Compared to CAMELS–IllustrisTNG and CAMELS–SIMBA, the XSB profiles in CAMELS–Astrid are very sensitive to SN wind speed. Increasing the $C_{\rm SN2}$ parameter increases the XSB, suggesting that increasing the SN wind speed suppresses the AGN feedback, leading to a greater retention of the CGM gas.

The third row shows the comparison of the eRASS CGM profile with the CAMELS–Astrid profile with varying $C_{\rm AGN1}$, which represents the kinetic AGN feedback energy. Compared to CAMELS–IllustrisTNG and CAMELS–SIMBA, the XSB profiles in CAMELS–Astrid are quite sensitive to the AGN feedback energy. Increasing the $C_{\rm AGN1}$ parameter decreases XSB, suggesting that the kinetic AGN feedback is effective in removing the hot gas from the halo. Similar to $C_{\rm SN1}$, varying $C_{\rm AGN1}$ alone always overpredicts the eRASS observations for all stellar mass bins.

The fourth row shows the comparison of the eRASS CGM profile with the CAMELS–Astrid profile with varying $C_{\rm AGN2}$, representing the energy of the thermal AGN feedback. Increasing $C_{\rm AGN2}$ leads to higher XSB, but the trend is not monotonic. As $C_{\rm AGN2}$ increases to about 2, the trend reverses, where increasing $C_{\rm AGN2}$ leads to a decrease in XSB normalization. This is because increasing the thermal AGN feedback increases the thermal energy of the gas, leading to adiabatic expansion of the gas, thereby reducing the overall gas density and XSB.

4.2. X-Ray Luminosity–Stellar Mass Relation

Figure 4 shows a comparison of $L_{\rm X}-M_\star$ between CAMELS–IllustrisTNG, CAMELS–SIMBA, CAMELS–Astrid, and eRASS (Y. Zhang et al. 2024b). We see a similar picture of the dependence on the four feedback parameters as in the XSB profiles. For CAMELS–IllustrisTNG, the relation is most sensitive to $A_{\rm SN1}$, followed by $A_{\rm SN2}$, $A_{\rm AGN2}$, and $A_{\rm AGN1}$. The fiducial TNG model underpredicts the eRASS data at low stellar mass but overpredicts at high stellar mass, but all within 2σ of the observations. Changing $A_{\rm SN1}$ leads to a larger change in the slope of the scaling relation, since the X-ray luminosities in low-stellar-mass halos are more sensitive to SN feedback than higher-stellar-mass halos.

For CAMELS–SIMBA, the $L_{\rm X}-M_\star$ relation is most sensitive to $B_{\rm SN1}$ and $B_{\rm SN2}$, but their trends are reversed. Higher $B_{\rm SN1}$ leads to lower $L_{\rm X}$, but higher $B_{\rm SN2}$ leads to higher $L_{\rm X}$. The impact of changing $B_{\rm SN1}$ is stronger at high stellar mass, thus varying $B_{\rm SN1}$ also changes the slope in the relation. Changing $B_{\rm SN2}$ leads to almost uniform change in $L_{\rm X}$ across the stellar mass bin, except at the lowest stellar mass, where $L_{\rm X}$ is not particularly sensitive to $B_{\rm SN2}$. The fiducial CAMELS–SIMBA model underpredicts the eRASS data at all stellar masses.

CAMELS–Astrid shows a stronger sensitivity to feedback parameters in the $L_{\rm X}-M_\star$ relation, compared to CAMELS-TNG and CAMELS–SIMBA. The relation is most sensitive to $C_{\rm SN2}$, followed by $C_{\rm SN1}$, $C_{\rm AGN2}$, and $C_{\rm AGN1}$. The fiducial CAMELS–Astrid model overpredicts the eRASS $L_{\rm X}$ over all stellar masses. Increasing the $C_{\rm SN1}$, $C_{\rm SN2}$, and $C_{\rm AGN1}$ feedback parameters leads to a slight decrease in the slope of the scaling relation, as the increases in feedback increase the X-ray luminosities in lower stellar mass halos more, due to their shallower potential wells. While increasing $C_{\rm AGN2}$ leads to higher normalization in $L_{\rm X}-M_\star$, for $C_{\rm AGN2}\geqslant 1.5$, the normalization decreases, as strong thermal AGN feedback decreases the gas density and thus the XSB in the inner CGM (see the fourth row in Figure 3).

4.3. Dependence on $\Omega_{\rm M}$ and σ_8

For all three simulation suites, we see similar dependencies of the XSB profile and the $L_{\rm X}-M_\star$ relation on Ω_M and σ_8. Increasing Ω_M leads to higher XSB and $L_{\rm X}$ values, due to the increase in the virial temperature of the CGM, as the increase in Ω_M leads to higher halo masses for a given stellar mass. Here, the Ω_M dependence is degenerate with the feedback parameters, especially with the SN parameters across the suites. In contrast, both XSB and $L_{\rm X}$ are not sensitive to σ_8.

4.4. Constraints on Feedback Parameters

We apply the emulator to the eRASS data set to obtain constraints on the feedback parameters. We follow a conventional Markov Chain Monte Carlo (MCMC) to infer the posterior distributions of the parameters given the eRASS data. Inference is performed by maximizing a Gaussian likelihood function. Specifically, the likelihood function is

$$\ln \mathcal{L}(L_{\rm X}|\theta) \propto -\sum_i \frac{(L_{{\rm X},i}^{\rm mod}(M_\star, z; \theta) - L_{{\rm X},i}(M_\star, z)^2}{2\sigma_i^2} \quad (4)$$

for the $L_{\rm X}-M_\star$ relation, where $L_{\rm X}^{\rm mod}$ is the emulated X-ray luminosity, with parameters $\theta = (A_{\rm SN1}, A_{\rm SN2}, A_{\rm AGN1}, A_{\rm AGN2})$,

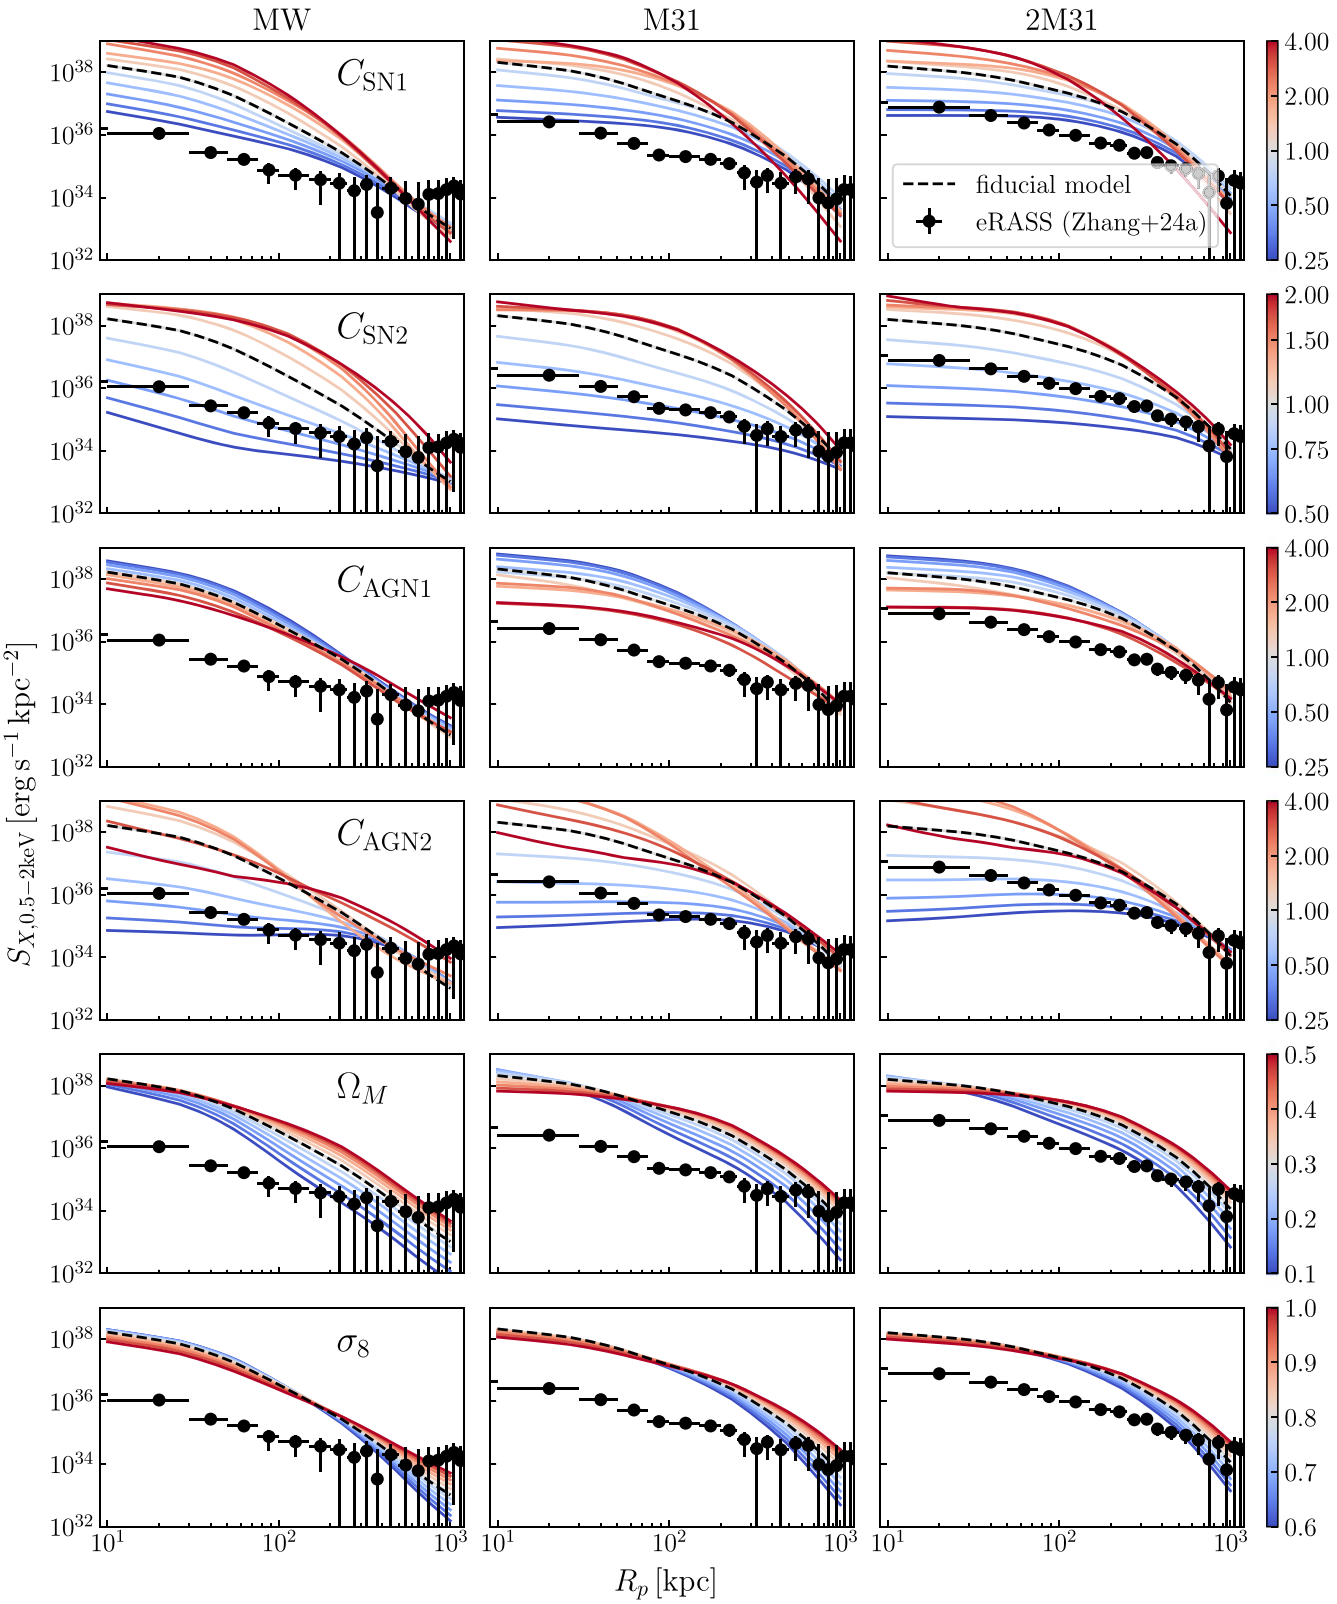

Figure 3. The same as Figure 1 but for CAMELS–Astrid.

and $L_{X,i}$ is the eRASS data, with error σ_i at the ith stellar mass bin. Because there are only four data points in the eRASS data, we do not fit for Ω_M and σ_8. Including them would result in six independent parameters, requiring two more data points. We fix $\Omega_M = 0.3089$ and $\sigma_8 = 0.8102$ (Planck Collaboration et al. 2020). The uniform prior for each feedback parameter is chosen to match the range explored in the simulations. We run MCMC for 10^5 steps. For each chain, we discard the first 5×10^4 steps, to make sure that the posteriors are not affected by the initial values of the parameters. We determine that the

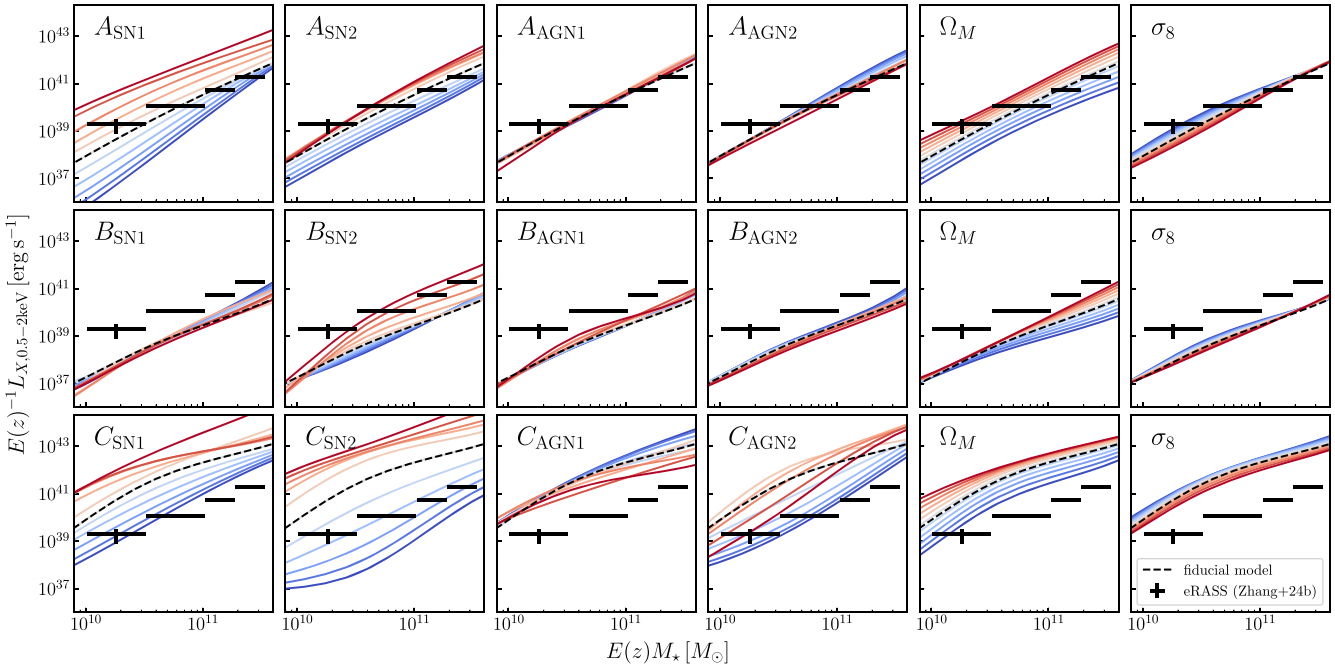

Figure 4. Comparison of the $L_{\rm X}-M_\star$ relation. The black data points are the eRASS data points with 1σ error bars (Y. Zhang et al. 2024b). The colored lines are the relations for CAMELS–IllustrisTNG (top), CAMELS–SIMBA (middle), and CAMELS–Astrid (bottom), with varying values for the four feedback parameters and Ω_M and σ_8, shown in the panels from left to right. The colors indicate the strengths of the feedback parameters as in Figures 1, 2, and 3, with red (blue) representing the maximum (minimum) feedback values. The black dashed line indicates the fiducial relation with all feedback strengths set to unity.

chains converge by making sure that the means and variances of the parameter values in the last 5000 steps remain essentially unchanged.

Figure 5 shows the posterior distributions of the feedback parameters for CAMELS–IllustrisTNG, CAMELS–SIMBA, and CAMELS–Astrid, respectively. The eRASS data provide reasonably good constraints on the feedback parameters. The reduced χ^2 values for the best-fit models are 0.83, 4.00, and 3.01 for CAMELS–IllustrisTNG, CAMELS–SIMBA, and CAMELS–Astrid, respectively, compared to those for the fiducial models: 4.17, 313, and 666. Figure 6 shows the eRASS best-fit scaling relations for CAMELS–IllustrisTNG, CAMELS–SIMBA, and CAMELS–Astrid, compared to the eRASS measurements from Y. Zhang et al. (2024b), and the best-fit values for the feedback parameters and their uncertainties are summarized in Table 2.

For CAMELS–IllustrisTNG, the eRASS1 data prefer slightly stronger SN1 (SN energy per SFR) and slightly weaker SN2 (stellar wind speed) feedback than the fiducial IllustrisTNG model. There are no constraints on AGN1 (kinetic AGN feedback energy) and AGN2 (AGN feedback frequency), which are almost the same as the fiducial values. For CAMELS–SIMBA, the model requires stronger SN and AGN feedback than the fiducial run to match the eRASS X-ray observations. On the other hand, for Astrid, the eRASS data prefer higher SN1 and AGN1 feedback, i.e., higher stellar feedback energy, and lower kinetic AGN feedback energy, respectively, but prefer lower values for SN2 and AGN2, representing the SN wind speed and thermal AGN feedback energy, respectively.

4.5. Implications for the Stellar Mass–Halo Mass Relation

The right panel of Figure 6 shows the stellar mass–halo mass relations for CAMELS–IllustrisTNG, CAMELS–SIMBA, and

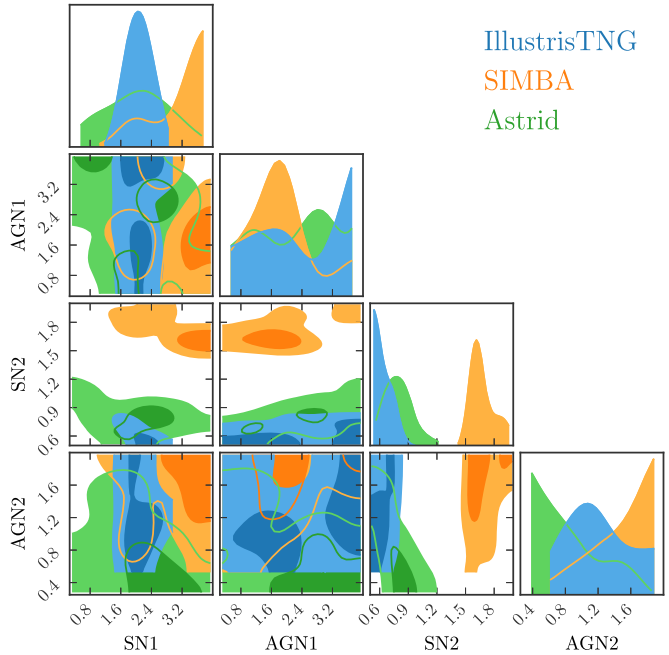

Figure 5. Posteriors of the best-fit parameters from MCMC for CAMELS–IllustrisTNG (blue), CAMELS–SIMBA (orange), and CAMELS–Astrid (green). The contours represent 1σ and 2σ of the uncertainties in the parameters.

CAMELS–Astrid, using their respective best-fit feedback parameters from the MCMC analysis performed on the eRASS data in Section 4.4 (solid curves). We compare them to the empirical fit to the observed stellar mass–halo mass relation (G. Girelli et al. 2020). There are differences between the best-fit CAMELS–eRASS relations with the observations. For CAMELS–IllustrisTNG, the eRASS best-fit stellar mass–halo mass relation matches with the fiducial one for halo masses

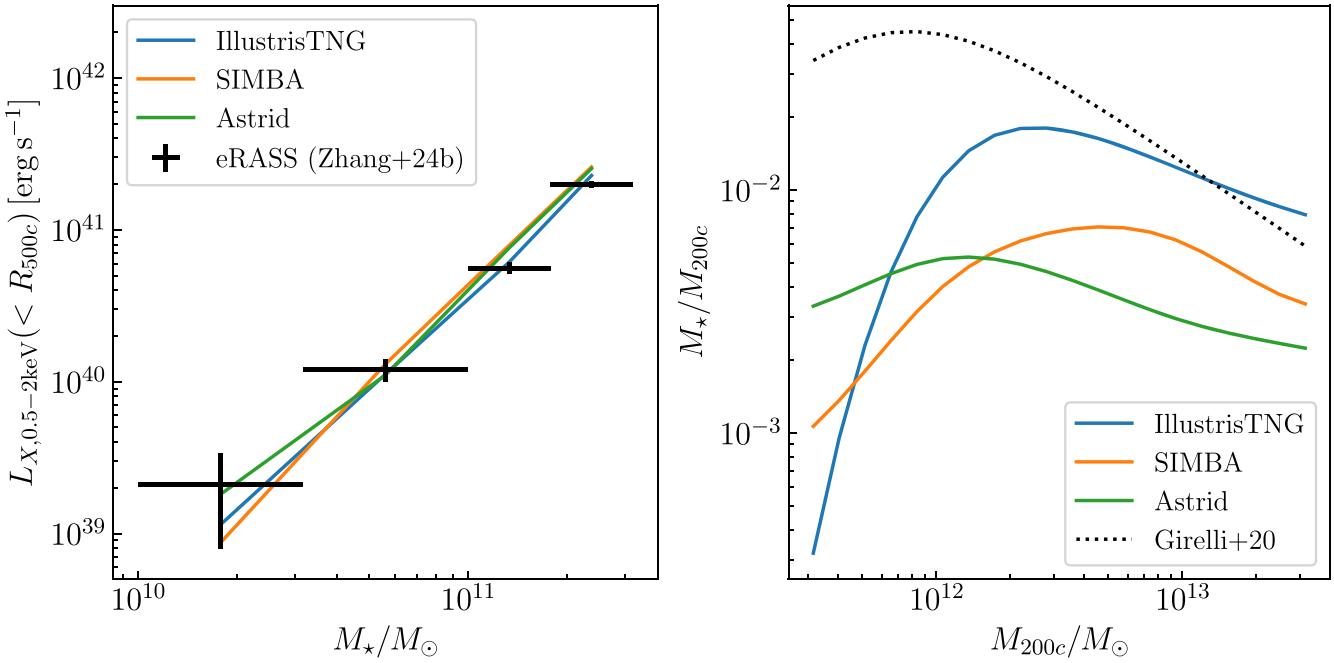

Figure 6. Left panel: the best-fit L_X–$M_\star$ scaling relations for CAMELS–IllustrisTNG (blue), CAMELS–SIMBA (orange), and CAMELS–Astrid (green), obtained by fitting against the eRASS data (black data points) from Y. Zhang et al. (2024b). Right panel: the corresponding stellar mass fraction as a function of halo mass for CAMELS–IllustrisTNG (blue), CAMELS–SIMBA (orange), and CAMELS–Astrid (green), compared to the empirical fit to the observed stellar mass fraction from G. Girelli et al. (2020; dashed line).

Table 2
Best-fit Feedback Parameters

Parameter	IllustrisTNG	SIMBA	Astrid
SN1	$2.07^{+0.35}_{-0.34}$	$3.45^{+0.35}_{-0.38}$	$2.07^{+0.82}_{-1.00}$
SN2	$0.6^{+0.13}_{-0.07}$	$1.63^{+0.20}_{-0.05}$	$0.78^{+0.13}_{-0.13}$
AGN1	$2.18^{+1.56}_{-1.14}$	$1.79^{+0.55}_{-0.80}$	$2.42^{+0.84}_{-1.39}$
AGN2	$1.15^{+0.53}_{-0.54}$	$1.62^{+0.27}_{-0.53}$	$0.67^{0.65}_{-0.30}$

$\log_{10}(M_{200c}/M_\odot) > 13$; below this mass scale, the predicted stellar mass fraction becomes smaller than the observed value. For both CAMELS–SIMBA and CAMELS–Astrid, the predicted stellar mass fraction is smaller than the observations.

The disagreement between the eRASS best-fit stellar mass fraction and the observed stellar mass–halo mass relation suggests that the feedback models in the CAMELS simulations are not able to simultaneously match both the eRASS X-ray CGM and observed stellar properties. There are also disagreements in the properties of the CGM between cosmological simulations and recent thermal and kinetic SZ observations of the CGM with the Atacama Cosmology Telescope (ACT; S. Amodeo et al. 2021), which prefers higher thermal pressure gas density and thus stronger feedback (E. Moser et al. 2022)—although for ACT, the differences showed up at larger halocentric radii ($r \gtrsim 1$ Mpc) and at higher halo masses ($M_{200c} \gtrsim 10^{13} M_\odot$).

This apparent tension between simulations and observations can be due to: (1) systematics in the stacked eRASS observations; or (2) the inadequacies of the feedback models in the cosmological simulations.

For (1), the eRASS X-ray CGM emission can potentially be overestimated. Specifically, X-ray binaries in the interstellar medium can contaminate the CGM due to the relatively wide point-spread function ($\sim 27''$) of eROSITA/SRG. The incomplete removal of background X-ray emissions from nearby galaxy groups can also boost the CGM signal. Since the observed eRASS XSB profiles and X-ray luminosities are derived from the stacking of spectroscopically selected galaxies, they are also susceptible to uncertainties in the stacking procedure. The stacked values can be biased by a few outliers in the stacking sample; the central galaxies selected for stacking can also be misclassified. For example, the X-ray properties of the CGM from the eROSITA eFEDS measurements have been shown to depend on the SFR and the morphology of the host galaxy (U. Chadayammuri et al. 2022; J. Comparat et al. 2022). Specifically, star-forming galaxies are shown to be more X-ray-bright than quiescent galaxies. Preferential selection of star-forming galaxies in the stacking procedure can lead to overestimates in the X-ray emissions. However, recent eRASS measurements (Y. Zhang et al. 2025) show that for Milky Way–mass systems, star-forming and quiescent galaxies have similar X-ray luminosities, while for more massive galaxies ($M_\star > 10^{11} M_\odot$), the X-ray luminosities of quiescent galaxies surpass those of star-forming systems. These potential systematics can be addressed with observations of the CGM of individual galaxies with XMM-Newton and Chandra that have higher angular resolution. SZ observations, which do not suffer from contamination by X-ray binaries, will also be useful for assessing these systematics. If the X-ray CGM is indeed overestimated, accounting for the contaminants will alleviate the need for stronger feedback, thus reducing the tension in the stellar mass fractions between simulations and the one inferred from eRASS.

For (2), as the CAMELS simulations have lower resolution compared to their corresponding flagship simulations (TNG100—A. Pillepich et al. 2018; SIMBA—R. Davé et al. 2019; Astrid—Y. Ni et al. 2022), the CAMELS versions show slightly lower X-ray luminosity than their flagship counterparts (compare the dashed lines in Figure 4 with Figure 2 in

Y. Zhang et al. 2024b), thus the level of feedback may not be required to be as high as for the CAMELS version for these flagship simulations, potentially reducing the tension with eRASS. In addition, the X-ray CGM properties can also depend on other parameters in the subgrid models that have not been explored in the current paper. We will address this issue with CAMELS simulations with more varied parameters, such as the 1P28 (I. Medlock et al. 2025) and the SB28 sets (Y. Ni et al. 2023).

5. Conclusions

In this paper, we use the CAMELS simulation suite to construct emulators of the XSB profile and the X-ray luminosity–stellar mass (L_X–$M_\star$) relation at the massive galaxy scale (with $\log_{10} M_\star/M_\odot \in [10, 11.5]$), to model the dependence of the CGM on feedback due to SNe and AGNs in subgrid models from three modern cosmological hydrodynamical simulations: IllustrisTNG, SIMBA, and Astrid. We compare the emulated XSB profiles and the L_X–$M_\star$ relation to the recent stacked measurements from eRASS (Y. Zhang et al. 2024a, 2024b). Here are our findings:

1. The XSB and L_X of all three models explored depend more sensitively on SN than AGN feedback. In particular, the changes in the efficiency of the AGN feedback have a lower impact on the X-ray observables.
2. The variations in the predicted X-ray properties among the three fiducial models are significant. For CAMELS–IllustrisTNG, the X-ray properties of the CGM match the eRASS observations best at higher stellar mass (M31-type galaxies and more massive ones). CAMELS–SIMBA underestimates the eRASS XSB profiles and L_X at all stellar masses by an order of magnitude, and CAMELS–Astrid overpredicts the eRASS results by an order of magnitude.
3. We derive the best-fit values of the feedback parameters of the three models using the L_X–$M_\star$ measurements from eRASS. While the eRASS data prefer stronger stellar feedback than the fiducial models in the CAMELS simulations, the stellar mass–halo mass relations using the best-fit values do not match with the observed stellar mass–halo mass relation.

Although it seems counterintuitive that the X-ray emissions are more dependent on SN and AGN feedback for CAMELS–IllustrisTNG and CAMELS–Astrid, in detail, the SN feedback is mainly responsible for regulating the AGN feedback that quenches star formation, as the SN feedback limits the amount of gas available for accreting onto the AGN (M. E. Lee et al. 2024; I. Medlock et al. 2025). This effect turns out to be more important than varying the amount of AGN feedback energy. To a lesser extent, more SN feedback also increases the metallicity of the CGM, thereby increasing the X-ray emission. In addition, the small box size of the CAMELS simulation also limits the number of more massive halos ($\sim 10^{13} M_\odot$) that are more susceptible to AGN feedback. Upcoming CAMELS simulations with larger box sizes, or zoom-in simulations (e.g., M. E. Lee et al. 2024), have larger numbers of more massive galaxies, groups, and clusters and will provide tighter constraints on AGN feedback physics.

Although modern cosmological hydrodynamical simulations often calibrate their galaxy formation models using observation data of galaxies, such as the stellar mass function and the stellar mass–halo mass relation, the equally important CGM properties have not been sufficiently constrained nor compared to observations. Our work demonstrates that X-ray observables are powerful in inferring and constraining the baryon cycles in massive galaxies. As next steps, in addition to the XSB profiles and X-ray luminosities studied in this work, temperature profiles, metallicity profiles, and spatial and kinematic structures will provide even more powerful constraints on baryonic feedback (G. Schellenberger et al. 2024; J. A. ZuHone et al. 2024). SZ and radio observations will also provide complementary probes of the CGM (e.g., I. Medlock et al. 2024; P. Singh et al. 2024). Forward modeling with light-cone simulations (S. Shreeram et al. 2024; E. T. Lau et al. 2025c) will become useful for accounting for observational systematics in the multiwavelength CGM observations. Combined with optical properties and SZ measurements, these X-ray constraints will serve as important inputs for refining feedback models in cosmological simulations and improving our understanding of galaxy formation.

Acknowledgments

We thank Y. Zhang, Johan Comparat, and the rest of the eROSITA Collaboration for providing us with the eRASS data and for stimulating discussions. D.N. and B.D.O. acknowledge the support of the NSF grant AST-2206055. Á.B. acknowledges support from the Smithsonian Institution and the Chandra Project through NASA contract NAS8-03060. D.A.A. acknowledges support from NSF grant AST-2108944, NASA grant ATP23-0156, STScI grants JWST-GO-01712.009-A and JWST-AR-04357.001-A, the Simons Foundation Award CCA-1018464, and the Cottrell Scholar Award CS-CSA-2023-028 from the Research Corporation for Science Advancement. The Flatiron Institute is supported by the Simons Foundation.

Software: emcee (D. Foreman-Mackey et al. 2013).

ORCID iDs

Erwin T. Lau https://orcid.org/0000-0001-8914-8885
Daisuke Nagai https://orcid.org/0000-0002-6766-5942
Ákos Bogdán https://orcid.org/0000-0003-0573-7733
Isabel Medlock https://orcid.org/0009-0000-7415-8239
Benjamin D. Oppenheimer https://orcid.org/0000-0002-3391-2116
Nicholas Battaglia https://orcid.org/0000-0001-5846-0411
Daniel Anglés-Alcázar https://orcid.org/0000-0001-5769-4945
Shy Genel https://orcid.org/0000-0002-3185-1540
Yueying Ni https://orcid.org/0000-0001-7899-7195
Francisco Villaescusa-Navarro https://orcid.org/0000-0002-4816-0455

References

Amodeo, S., Battaglia, N., Schaan, E., et al. 2021, PhRvD, 103, 063514
Anderson, M. E., Gaspari, M., White, S. D. M., Wang, W., & Dai, X. 2015, MNRAS, 449, 3806
Bird, S., Ni, Y., Di Matteo, T., et al. 2022, MNRAS, 512, 3703
Bogdán, Á., Bourdin, H., Forman, W. R., et al. 2017, ApJ, 850, 98
Bogdán, Á., Forman, W. R., Kraft, R. P., & Jones, C. 2013a, ApJ, 772, 98
Bogdán, Á., Forman, W. R., Vogelsberger, M., et al. 2013b, ApJ, 772, 97
Bregman, J. N., Hodges-Kluck, E., Qu, Z., et al. 2022, ApJ, 928, 14
Chadayammuri, U., Bogdán, Á., Oppenheimer, B. D., et al. 2022, ApJL, 936, L15
Comparat, J., Truong, N., Merloni, A., et al. 2022, A&A, 666, A156
Das, S., Chiang, Y.-K., & Mathur, S. 2023, ApJ, 951, 125
Davé, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, MNRAS, 486, 2827

Delgado, A. M., Anglés-Alcázar, D., Thiele, L., et al. 2023, MNRAS, 526, 5306
Donahue, M., & Voit, G. M. 2022, PhR, 973, 1
Faucher-Giguère, C.-A., & Oh, S. P. 2023, ARA&A, 61, 131
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306
Foster, A. R., Ji, L., Smith, R. K., & Brickhouse, N. S. 2012, ApJ, 756, 128
Gebhardt, M., Anglés-Alcázar, D., Borrow, J., et al. 2024, MNRAS, 529, 4896
Girelli, G., Pozzetti, L., Bolzonella, M., et al. 2020, A&A, 634, A135
Gluck, N., Oppenheimer, B. D., Nagai, D., Villaescusa-Navarro, F., & Anglés-Alcázar, D. 2024, MNRAS, 527, 10038
Hernández-Martìnez, E., Genel, S., Villaescusa-Navarro, F., et al. 2025, ApJ, 981, 170
Lau, E. T. 2025a, CGM Toolkit, v1.0.0, Zenodo, doi:10.5281/zenodo.14714182
Lau, E. T. 2025b, CAMELS CGM Emulator, v1.0.0, Zenodo, doi:10.5281/zenodo.14714176
Lau, E. T., Nagai, D., Farahi, A., et al. 2025c, ApJ, 980, 122
Lee, M. E., Genel, S., Wandelt, B. D., et al. 2024, ApJ, 968, 11
Lovisari, L., Ettori, S., Gaspari, M., & Giles, P. A. 2021, Univ, 7, 139
Medlock, I., Nagai, D., Singh, P., et al. 2024, ApJ, 967, 32
Medlock, I., Neufeld, C., Nagai, D., et al. 2025, ApJ, 980, 61
Moser, E., Battaglia, N., Nagai, D., et al. 2022, ApJ, 933, 133
Nelson, D., Pillepich, A., Springel, V., et al. 2018, MNRAS, 475, 624
Ni, Y., Di Matteo, T., Bird, S., et al. 2022, MNRAS, 513, 670
Ni, Y., Genel, S., Anglés-Alcázar, D., et al. 2023, ApJ, 959, 136
Nicola, A., Villaescusa-Navarro, F., Spergel, D. N., et al. 2022, JCAP, 2022, 046
Pandey, S., Lehman, K., Baxter, E. J., et al. 2023, MNRAS, 525, 1779
Pillepich, A., Springel, V., Nelson, D., et al. 2018, MNRAS, 473, 4077
Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, A&A, 641, A6
Schellenberger, G., Bogdán, Á., ZuHone, J. A., et al. 2024, ApJ, 969, 85
Shreeram, S., Comparat, J., Merloni, A., et al. 2024, arXiv:2409.10397
Singh, P., Lau, E. T., Faerman, Y., Stern, J., & Nagai, D. 2024, MNRAS, 532, 3222
Tumlinson, J., Peeples, M. S., & Werk, J. K. 2017, ARA&A, 55, 389
Villaescusa-Navarro, F., Anglés-Alcázar, D., & Genel, S. 2021a, arXiv:2109.09747
Villaescusa-Navarro, F., Anglés-Alcázar, D., & Genel, S. 2021b, ApJ, 915, 71
Wadekar, D., Thiele, L., Hill, J. C., et al. 2023, MNRAS, 522, 2628
Zhang, Y., Comparat, J., Ponti, G., et al. 2024a, A&A, 690, A267
Zhang, Y., Comparat, J., Ponti, G., et al. 2024b, A&A, 690, A268
Zhang, Y., Comparat, J., Ponti, G., et al. 2025, A&A, 693, A197
ZuHone, J. A., Schellenberger, G., Ogorzałek, A., et al. 2024, ApJ, 967, 49