Percent-level constraints on baryonic feedback with spectral distortion measurements

Leander Thiele[®], ^{1,*} Digvijay Wadekar[®], ² J. Colin Hill[®], ^{3,4} Nicholas Battaglia, ⁵ Jens Chluba, ⁶ Francisco Villaescusa-Navarro, ^{4,7} Lars Hernquist[®], ⁸ Mark Vogelsberger, ⁹ Daniel Anglés-Alcázar[®], ^{10,4} and Federico Marinacci[®]11

¹Department of Physics, Princeton University, Princeton, New Jersey 08544, USA ²School of Natural Sciences, Institute for Advanced Study, Princeton, New Jersey 08540, USA ³Department of Physics, Columbia University, New York, New York 10027, USA ⁴Center for Computational Astrophysics, Flatiron Institute, New York, New York 10010, USA ⁵Department of Astronomy, Cornell University, Ithaca, New York 14853, USA ⁶Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, United Kingdom ⁷Department of Astrophysical Sciences, Princeton University, Princeton, New Jersey 10027, USA Center for Astrophysics, Harvard & Smithsonian, Cambridge, Massachusetts 02138, USA ⁹Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

¹⁰Department of Physics, University of Connecticut, Storrs, Connecticut 06269, USA ¹¹Department of Physics and Astronomy "Augusto Righi," University of Bologna, Bologna I-40129, Italy



(Received 19 January 2022; accepted 22 March 2022; published 11 April 2022)

High-significance measurements of the monopole thermal Sunyaev-Zel'dovich cosmic microwave background spectral distortions have the potential to tightly constrain poorly understood baryonic feedback processes. The sky-averaged Compton-y distortion and its relativistic correction are measures of the total thermal energy in electrons in the observable universe and their mean temperature. We use the CAMELS suite of hydrodynamic simulations to explore possible constraints on parameters describing the subgrid implementation of feedback from active galactic nuclei and supernovae, assuming a PIXIE-like measurement. The small 25 h^{-1} Mpc CAMELS boxes present challenges due to significant sample variance. We utilize machine learning to construct interpolators through the noisy simulation data. Using the halo model, we translate the simulation halo mass functions into correction factors to reduce sample variance where required. Our results depend on the subgrid model. In the case of IllustrisTNG, we find that the best-determined parameter combination can be measured to $\approx 2\%$ and corresponds to a product of active galactic nuclei (AGN) and supernova (SN) feedback. In the case of SIMBA, the tightest constraint is ≈0.2% on a ratio between AGN and SN feedback. A second orthogonal parameter combination can be measured to ≃8%. Our results demonstrate the significant constraining power a measurement of the late-time spectral distortion monopoles would have for baryonic feedback models.

DOI: 10.1103/PhysRevD.105.083505

I. INTRODUCTION

A robust prediction of the standard model of cosmology is the presence of subtle deviations from the blackbody spectrum in the cosmic microwave background (CMB) (e.g., [1,2]). Such monopole spectral distortions can be caused by a variety of physical processes in both the earlyand late-time Universe. Among these various distortion signals, most easily accessible to near-future experiments will be the y-type distortion. The y distortion arises from inverse Compton scattering of CMB photons with hot electrons, the thermal Sunyaev-Zel'dovich (tSZ) effect [3,4]. The y signal is mostly sourced by massive, collapsed structures at $z \lesssim 2$. Thus, the tSZ effect informs us about two aspects: First, the thermodynamic properties of the hot electron gas in massive halos; these are very sensitive to astrophysical small-scale processes. Second, the abundance of such halos, which translates into a constraint on the amount and clustering of matter in the Universe. This makes SZ cluster measurements a unique tool for cosmology and astrophysics [5,6].

To illustrate this point, consider Fig. 1. There, we plot a range of theory predictions (blue and magenta points) from the CAMELS suite of hydrodynamic simulations [7].

lthiele@princeton.edu

¹For conciseness, we will use the term "y distortion" or similar to mean both the nonrelativistic and the relativistic effects.

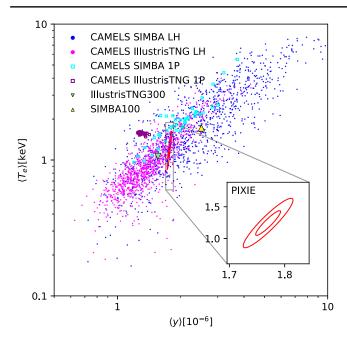


FIG. 1. Comparison between the range of CAMELS simulations with the forecast covariance matrix for the PIXIE experiment (the PIXIE ellipses are centered at the fiducial model assumed in Ref. [8]: $\langle y \rangle_{\rm fid} = 1.77 \times 10^{-6}, \ \langle T_e \rangle_{\rm fid} = 1.24 \ \rm keV$; they correspond to 68 and 95% cl). From CAMELS, we plot both the LH and the 1P set. The LH data points are contaminated by sample variance in addition to varying subgrid parameters, while the 1P points have the same initial conditions and only differ by their subgrid parameters. We have also indicated the values measured from the large boxes (triangles), as discussed in Sec. II B. All simulation data have been rescaled to the fiducial CAMELS cosmology but each data point corresponds to a different subgrid model. It is worth noting that the 1P initial conditions appear to be slightly atypical, with generally larger $\langle T_e \rangle$ than expected from the LH set.

The axes are the two distortion monopole observables considered in this work, further discussed below. Between the plotted simulations the astrophysical subgrid model differs, translating into a large spread of theory predictions. For comparison, we also show a covariance matrix for a nearfuture monopole distortion measurement (cf. Sec. III). The forecast measurement errors are tiny compared to the current theoretical uncertainty, which means that a near-future measurement would provide a substantial gain in information on astrophysics. This basic observation forms the underpinning of the calculations performed in the following: forecast constraints on simulation subgrid models from a measurement of the y-distortion monopoles, using simulations from the CAMELS suite.

Conventionally, the y distortion is separated into a nonrelativistic and a relativistic [9–13] component, each

having a distinct spectral signature that makes it possible to disentangle them observationally.³ The nonrelativistic contribution is determined by a line-of-sight integral over electron pressure,⁴

$$\langle y \rangle \equiv \langle y(\hat{n}) \rangle_{\hat{n}} = \int \frac{d\hat{n}}{4\pi} \frac{\sigma_{\rm T}}{m_e} \int P_e(\hat{n}, l) dl,$$
 (1)

where $\sigma_{\rm T}$ denotes the Thomson cross section, m_e is electron mass, \hat{n} is the line of sight, and l is physical length along the line of sight. To leading order, the relativistic component is proportional to the *y*-weighted mean electron temperature [15]

$$\langle T_e \rangle \equiv \langle T_e(\hat{n}) \rangle_{\hat{n}} = \langle y \rangle^{-1} \int \frac{d\hat{n}}{4\pi} \frac{\sigma_{\rm T}}{m_e} \int [T_e P_e](\hat{n}, l) dl. \quad (2)$$

This effective temperature is typically higher than the mass-weighted (or τ -weighted) temperature [16,17], and can be directly obtained from a moment expansion of the SZ signal (see Appendix A for additional discussion).

In this work we focus on the dominant and theoretically well-established contributions to the distortion signals. In particular, we will neglect the $\approx 10\%$ contribution of reionization to $\langle y \rangle$, which we discuss in Appendix C, as well as signals due to the Milky Way and local group (which are estimated to be another 1 and 2 orders of magnitude below the reionization signal, respectively). We also neglect other, more exotic sources of y distortions, for example from primordial magnetic field heating (e.g., [18–20]) or decaying particles (e.g., [21–26]). Conversely, if one is interested in using the y distortions to constrain or detect such processes beyond standard Λ cold dark matter (Λ CDM), astrophysical feedback must be very well understood.

There are further relativistic corrections involving higher moments of the electron temperature, as well as the kinetic Sunyaev-Zel'dovich effect sourced by coherent motion; we will neglect these complications and simply treat $\langle y \rangle$ and $\langle T_e \rangle$ as observables, as in Ref. [15].

Locally, the tSZ effect is well established observationally as a CMB temperature change correlating with the locations of clusters. However, the global distortion to the CMB spectrum has not yet been detected. Only an upper limit on the nonrelativistic y distortion exists from the COBE FIRAS experiment, which yielded $|\langle y \rangle| < 15 \times 10^{-6} (95\% \text{cl})$ [27], which is about 1 order of magnitude above the expected Λ CDM signal [15]. We thus anticipate significant detections with future CMB spectroscopy [28–31].

A monopole measurement would be complementary to the existing higher-moment tSZ analyses (e.g., [32–46]), since it features very different systematics and would also yield the relativistic component at high significance [8].

²For IllustrisTNG, the spread is dominated by sample variance [as can be seen by comparing the latin hypercube (LH) and one parameter (1P) points], a point that we will return to later.

³All these signals can be accurately modeled using SZpack [14]. ⁴We set the speed of light and Boltzmann's constant to one.

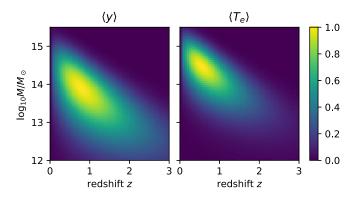


FIG. 2. Illustration of the halo mass and redshift contributions to the *y*-distortion observables. Plotted is the normalized integrand in Eq. (4), using the fitting formulas described in Sec. IV B.

Furthermore, in contrast to cluster-stacking and power spectrum [47–49] approaches, the $\langle y \rangle$ measurement is more sensitive to lower-mass objects, as illustrated in Fig. 2.

The $\langle y \rangle$ and $\langle T_e \rangle$ signals constitute unique probes of baryonic physics in galaxy clusters and groups. Since $\langle y \rangle$ probes thermal energy, it is subject to the energy conservation equation

$$\underbrace{E_e^{\text{th,tot}}}_{\text{Total}} = \underbrace{E_e^{\text{coll}}}_{\text{Collapse}} + \underbrace{E_e^{\text{inj}}}_{\text{Injected}} - \underbrace{E_e^{\text{cool}}}_{\text{Cooling}}.$$
(3)

The most uncertain term in the above equation is $E_e^{\rm inj}$ which can largely be attributed to feedback processes from massive stars, supernovae (SNe), and active galactic nuclei (AGN). These processes inject additional energy into the interstellar, intergalactic, and intracluster media (ISM, IGM, and ICM, respectively). Such feedback processes are standard ingredients in any theoretical models of galaxy formation, both semianalytic (e.g., [50,51]) and simulation-based models e.g., [52–57].

The most reliable way to explore how feedback models influence the *y* distortions is by analyzing hydrodynamical simulations with qualitatively and quantitatively different subgrid prescriptions. Such an approach is complicated by the fact that, owing to their bias towards rare high-density peaks, the distortion signals are heavily influenced by sample variance. Furthermore, the parameter space of subgrid models is vast and poorly explored, meaning that ideally we would need many large-volume hydrodynamical simulations, which is currently not feasible. We will demonstrate later that these problems can be overcome by utilizing machine learning methods as well as analytical corrections using the halo model.

Since the y distortions are predominantly sourced by galaxy groups and clusters [15], an analytical description based on the halo model [58–60] is a natural first approximation. In the halo model formalism, we assume spherically symmetric halos described only by mass M and redshift z, yielding

$$\langle y \rangle_{\text{hm}} = \frac{\sigma_{\text{T}}}{m_e} \int dz dM \frac{(1+z)^2}{4\pi H(z)} \frac{dn}{dM} \int d\vec{r} P_e(|\vec{r}|; M, z), \quad (4)$$

where dn/dM is the halo mass function and \vec{r} denotes position within a given halo, and the expression assumes a flat universe. The expression for $\langle T_e \rangle_{\rm hm}$ is analogous. Note that the halo model neglects the IGM contribution discussed in Appendix D.

In the following, it will be useful to think of the observables $x_i \equiv \{\langle y \rangle, \langle T_e \rangle\}$ in terms of the approximate factorization

$$x_i \sim f_i^c(\sigma_8, \Omega_m, \dots) f_i^b(\{A_i\}) f_i^{CV}(\delta), \tag{5}$$

where f_i^c describe the dependence on cosmological parameters, f_i^b are functions of a set of feedback parameters A_j , and $f_i^{\rm CV}$ depends on the initial conditions and thus encapsulates sample variance. Such factorizations are frequently used and good approximations to observables that are well described by the halo model. Nonetheless, our results typically only weakly depend on the validity of this approximation.

It should be noted that we group all our uncertainty on the simulation subgrid model in the feedback parameters A_j . These supernova and AGN feedback parameters, further elaborated on in Sec. II A, predominantly affect the ICM contribution to the distortion signals. There is a non-negligible IGM contribution, however, which in Appendix D we show is a $\sim 10\%$ effect with $\sim 40\%$ theoretical uncertainty.

The rest of this paper is structured as follows. In Sec. II we describe the CAMELS simulations as well as the larger reference boxes. Section III provides a short summary of the assumed experimental setup used for forecasting. In Sec. IV we describe how we interpolate through the CAMELS data. Section V contains the main results of this work, namely, dependence of the distortion signals on feedback parameters and a Fisher forecast. We conclude in Sec. VI. The Appendices contain several technical details and some new computations that did not fit in the main discussion.

II. SIMULATIONS

The y-distortion components described above can easily be measured from a hydrodynamical simulation. In fact, Eq. (1) can be rewritten as

$$\langle y \rangle = \frac{\sigma_{\rm T}}{m_e} \int dz \frac{(1+z)^2}{H(z)} \langle P_e^c(z) \rangle,$$
 (6)

where P_e^c is now in comoving units and the average is over the volume of a given simulation snapshot. An analogous expression holds for $\langle T_e \rangle$.

TABLE I. Comparison between statistics over the CAMELS CV set and the larger boxes of sizes L=205 and $100\ h^{-1}$ Mpc for IllustrisTNG and SIMBA respectively. The left half of the table is in log since log-normal is a decent approximation; for clarity the right half gives the same information in more familiar units but without error bars. The error bars in the first line are on the mean of the CV set. The error bars on the large box measurements are, in the naive Poissonian approximation $\sqrt{(L/25\ h^{-1}\ \text{Mpc})^3/27}$, smaller by factors 4.5 and 1.5 for IllustrisTNG and SIMBA respectively, and omitted in the table. All large-box measurements have been rescaled to the fiducial CAMELS cosmology, as explained in Sec. II B. The third line lists the differences between the large-box and CV measurements, in units of the error bars on the CV means; these numbers highlight significant biases in the CV set due to box-size effects. The last line gives the corresponding bias factors introduced in Eq. (7).

	$\log_{10} 10^6 \langle y \rangle$		$\log_{10}\langle T_e \rangle$ [keV]		$10^6 \langle y \rangle$		$\langle T_e \rangle$ [keV]	
	IllustrisTNG	SIMBA	IllustrisTNG	SIMBA	IllustrisTNG	SIMBA	IllustrisTNG	SIMBA
CAMELS CV	0.046 ± 0.026	0.369 ± 0.014	-0.386 ± 0.034	0.1407 ± 0.0054	1.11	2.34	0.41	1.38
Large box	0.1928	0.3995	0.02728	0.2349	1.56	2.51	1.06	1.72
$\Delta[\sigma]$	-5.6	-2.1	-12	17				
b_i					1.41	1.07	2.59	1.25

A. CAMELS

We primarily use the CAMELS suite of hydrodynamical simulations [7] that consists of several thousand 25 h^{-1} Mpc boxes, each run with 256³ dark matter particles and 256³ initial fluid elements. Each simulation is described by the following parameters: (i) the simulation code/subgrid model, (ii) two cosmological parameters (σ_8 , Ω_m), (iii) four feedback parameters (σ_{8N1} , σ_{8N2} , σ_{8N2} , σ_{8N1} , σ_{8N2} , and (iv) the random seed for the initial conditions. The remaining cosmological parameters are fixed at *Planck*-compatible flat σ_{8N1} values, σ_{8N2} = 0.049, σ_{8N2} = 0.9624 [62].

Two different simulation codes are used. The simulations labeled "IllustrisTNG" were run with the Arepo code [52,63] and the same subgrid model as the flagship IllustrisTNG simulations [64–69]. The simulations labeled "SIMBA" were run with the GIZMO code [70] and the same subgrid model as the flagship SIMBA simulations [71]. These codes differ substantially in their subgrid implementations, so having comparable simulations with both gives a good indication of the theoretical uncertainty.

The cosmological parameters σ_8 , Ω_m are varied in the intervals [0.6, 1.0] and [0.1, 0.5], respectively, the fiducial model being (0.8,0.3).

The precise definition of the four feedback parameters is given in Ref. [7]. Broadly, $A_{\mathrm{SN}i}$ parametrize the subgrid prescription for galactic winds, while $A_{\mathrm{AGN}i}$ describe the efficiency of black hole feedback. The i=1 components can be thought of as "energy" normalizations, while the i=2 components scale the speed of outflows. In detail, however, the meaning of the feedback parameters differs substantially between IllustrisTNG and SIMBA. We thus

caution against any direct comparison in terms of the feedback parameters between the two subgrid models. All the A_j are multiplicative factors relative to the fiducial efficiencies in the original IllustrisTNG and SIMBA simulations; the fiducial model is therefore unity for all feedback parameters. $A_{\rm SN1}$ and $A_{\rm AGN1}$ are varied in [0.25, 4.0], while $A_{\rm SN2}$ and $A_{\rm AGN2}$ are varied in [0.5, 2.0]. These intervals were chosen heuristically by the CAMELS team, and as we will see the corresponding variations in the y observables differ drastically between IllustrisTNG and SIMBA. Some of the simulations at the corners of parameter space are so extreme that they are certainly not realistic, for example with regard to galaxy properties.

For each of the two simulation codes, the CAMELS suite comprises the following sets of simulations, all of which will be used in this work.

- (i) LH: (latin hypercube) 1000 simulations in which cosmology and feedback parameters are varied on a latin hypercube, each run having a different random seed,
- (ii) 1P: (one parameter at a time) 10 variations for each cosmological and feedback parameter individually at fixed random seed,
- (iii) CV: (cosmic variance) 27 simulations at the fiducial model with differing random seeds.

B. Larger boxes

In addition to the CAMELS suite, we also use larger boxes, namely IllustrisTNG300-1 (205 h^{-1} Mpc) and SIMBA100 (100 h^{-1} Mpc). These simulations are useful to calibrate against the fact that $\langle f_i^{\rm CV}(\delta) \rangle$ [cf. Eq. (5)] is relatively more biased for the small 25 h^{-1} Mpc CAMELS boxes. Denoting an observable measured in one of the large boxes as $x_i^{\rm lb}$ [cf. Eq. (5) for the notation], we compute an estimator for this multiplicative bias as

$$b_i = x_i^{\text{lb}} / \langle x_i \rangle_{\text{CV}},\tag{7}$$

⁵Note that other works, e.g., Ref. [61], explicitly differentiate in their notation between the IllustrisTNG and SIMBA feedback parameters. We choose not to do so since it simplifies the notation in various places.

where the latter average is over the CAMELS CV set and we apply the correction formula

$$x_i \leftarrow x_i b_i \tag{8}$$

to the $\langle y \rangle$, $\langle T_e \rangle$ measured in the CAMELS simulations. Values for the b_i are listed in Table I.

The large boxes have the same subgrid prescription as the CAMELS fiducial model with the corresponding label, but slightly different cosmologies. We account for this fact by rescaling x_i^{lb} to the CAMELS fiducial cosmology, using power laws in h, Ω_m , Ω_b , n_s , σ_8 . These power laws were fitted using the halo model, since the scalings that could be derived from the 1P set are unreliable and do not encompass all differences in cosmology. We refer to Sec. IV B for a detailed description of the assumptions made in the halo model calculation. For reference, the power laws are listed in Appendix B.

There is, of course, a remaining bias due to the finite volume of both the larger boxes and the CV set. However, as seen in Table I, this is a relatively small source of error compared to the overall bias, since Eq. (7) is significantly different from unity for all cases except perhaps the SIMBA $\langle y \rangle$. We note that the finite error bars on the large box measurements do not affect our forecasts in the later parts of this work.

III. PIXIE EXPERIMENTAL MODEL

For deriving our forecasts later in Sec. V, we use parameters corresponding to the "extended" PIXIE experiment [28,29,72] from Ref. [8]. Their most complete model includes marginalization over a variety of foregrounds, namely galactic dust thermal emission, the cosmic infrared background, synchrotron radiation, free-free emission, spinning dust (anomalous microwave emission), and integrated CO. For each foreground component, fitting formulas were assumed for the spectral energy distribution (SED) according to *Planck* measurements [73]. There are some caveats to this approach, related to the spatial variation of the foregrounds and the relatively simplistic modeling of their SEDs (see [74] for related discussion). However, for the purposes of this work, the derived forecasts should be accurate enough. The forecast considers all non-negligible CMB spectral distortion signals (blackbody, y, relativistic y, μ) and marginalizes over them when computing the $\langle y \rangle - \langle T_e \rangle$ posterior.

While Ref. [8] used a full MCMC pipeline to arrive at their posteriors, in the $\langle y \rangle - \langle T_e \rangle$ plane, these are well approximated as Gaussian. We thus compress the marginalized posterior into a simple 2×2 covariance matrix, illustrated in the inset in Fig. 1. In general, the fiducial models used in this work differ somewhat from the fiducial model of Ref. [8]; we assume that the covariance matrix is constant regardless of the mean. This approximation is not a dominant source of systematic uncertainty.

IV. INTERPOLATING NEURAL NETWORKS

Using Eq. (6), we compute $\langle y \rangle$ and $\langle T_e \rangle$ for the CAMELS LH, CV, and 1P sets. It is our primary goal to extract the functions $f_i^b(A_{\rm SNI},A_{\rm AGNI},A_{\rm SN2},A_{\rm AGN2})$, in the language of Eq. (5). It may be argued that in the limit in which these f_i^b factorize even further into functions of the individual feedback parameters (which, as we shall see, is quite a good approximation), the 1P set should be all we need. However, the substantial sample variance in the small CAMELS boxes makes this approach unreliable. Therefore, we will instead utilize the LH set. There, σ_8 , Ω_m , and the four feedback parameters are varied by sampling from a latin hypercube. Furthermore, each data point is a noisy sample in a 6-dimensional space. Thus, we will use neural networks as smoothing interpolators through the LH set.

A. Training

We aim to learn functions

$$\langle y \rangle$$
 or $\langle T_e \rangle = F_\alpha(\sigma_8, \Omega_m, \{A_j\}),$ (9)

where $\alpha \in \{\text{IllustrisTNG}, \text{SIMBA}\}\$ and F is parametrized as a multilayer perceptron. A multilayer perceptron is a series of affine transformations $W\mathbf{x} + \mathbf{b}$, each followed by a nonlinear activation function (in our case the leaky rectifying linear unit). To mitigate overfitting, we allow dropout, i.e., probabilistic zeroing of neurons during training.

The relatively small dataset of 1,000 LH simulations makes this a somewhat nontrivial task, thus we perform automated hyperparameter optimization using the Optuna package [75] to converge at good architectures. Generally, our networks have 2-4 hidden layers with a few hundred neurons each and high dropout rates 20%-70%. A useful null-test is the following. From the 1P set we can interpolate the dependence on σ_8 , Ω_m (simple linear interpolators which will generally be biased due to sample variance). Using these interpolants, we can remove most of the dependence on the cosmological parameters from the LH data (cf. the cyan data points in Fig. 3). If the neural net is well-converged, it should matter very little whether the training data have the full dependence on cosmology or whether it is mostly factored out. The point here is that we are essentially fitting the same data, just with different transformations applied. It is not actually very important what these transformations are as a function of σ_8 and Ω_m , we simply chose ones that remove most of the dependence on these cosmological parameters.

Indeed, in the case of SIMBA, this null test is passed satisfactorily. However, IllustrisTNG is more challenging, since the feedback parameters generally have a much smaller effect on the *y* observables in this simulation.

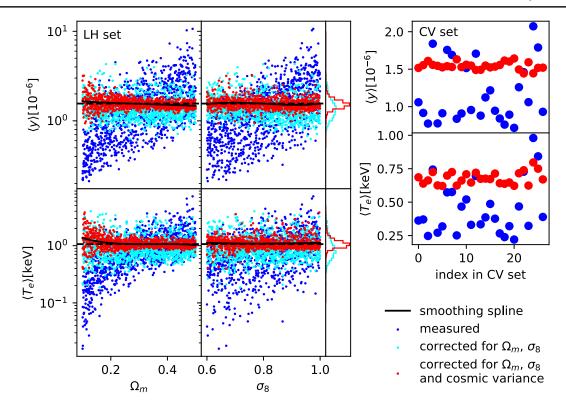


FIG. 3. Illustration of the effectiveness of analytical correction factors for the IllustrisTNG simulation outputs. Blue markers are direct measurements from the simulations. The cyan markers represent a rescaling to the fiducial cosmology ($\Omega_m = 0.3$, $\sigma_8 = 0.8$), using fits to the CAMELS 1P set. Since the CAMELS CV set is at constant cosmological and astrophysical parameters, the cosmology correction is meaningless there. For the red markers, in addition to the cosmology correction, a correction for sample variance has been applied, as described in the text. The dashed black lines represent the "ground truth" measurements from the larger boxes, while the solid black lines are smoothing splines through the red markers (these are solely to guide the eye). For the LH panels on the left, all data points have been rescaled using Eq. (7). We have also plotted histograms corresponding to the cyan and red markers.

Thus, we need to apply an additional transformation, as described in the following section.

B. Halo model correction factors for IllustrisTNG

As mentioned before, naively training the neural networks on the IllustrisTNG LH simulations does not yield robust results according to the null test described above. The reason is that, in contrast to SIMBA, the dependence on feedback parameters is obscured by the overwhelming noise due to sample variance. This issue can be seen in Fig. 1, where for IllustrisTNG the 1P data points scatter much less than the LH points. There is, however, a simple method to mitigate this problem. As we have argued in the Introduction, the halo model provides a relatively good description of sparse fields such as Compton y. Furthermore, the factorization of Eq. (5) should be a good initial guess. This motivates the following "correction" formula:

$$\langle y \rangle \leftarrow \langle y \rangle \frac{\langle y \rangle_{\text{hm}}^{\text{Tinker HMF}}}{\langle y \rangle_{\text{hm}}^{\text{measured HMF}}}.$$
 (10)

Here, "Tinker HMF" stands for the halo mass function from Ref. [76] and the expression is analogous for $\langle T_e \rangle$. On the

other hand, the mass function can also be measured in the individual CAMELS simulations. In order for the numerical integrations to be well behaved, we express it as

$$\frac{dn}{d\log M} \sim \exp\left[-\frac{1}{2}\left(\frac{\log M'}{\sigma}\right)^2\right] * \sum_{i} \delta(\log M' - \log M_i),$$
(11)

where * indicates convolution and we have dropped some prefactors. Our results are relatively insensitive to the hyperparameter σ ; in the following we will use $\sigma = 0.19$ dex. The halo masses M_i were measured using a friends-of-friends finder [77]. We interpolate the smoothed mass function in mass and redshift using a bilinear routine.

In order to perform the halo model calculation, we use the pressure profile fitting formula from Ref. [78] and assume isothermal halos with temperatures according to Ref. [79]. We apply a 20% correction to the halo masses entering the temperature fitting formula to account for hydrostatic mass bias; the same value was used in Ref. [15]. It should be noted that the fitting formula from Ref. [79] is likely not as accurate as more recent proposals (e.g., [17]),

but as we show in the following, it is sufficient for our purposes. Following Ref. [15], we assume a radial cutoff at $2.5R_{vir}$ with the virial radius definition of Ref. [80].

Of course, this procedure rests on a number of assumptions, and we should verify that it yields reasonable results. Figure 3 shows in the right panel in blue the original measurements and in red the mass function-corrected values for the 27 realizations in the CV set. Indeed, we see that the corrected values scatter much more tightly; the fact that the procedure is more successful for $\langle y \rangle$ than $\langle T_e \rangle$ is likely because the temperature fitting formula used is calibrated at larger halo masses than those that dominate the signal. In the left panel of Fig. 3, we show projections of the LH set onto the σ_8 and Ω_m axes. As expected, the blue measurement data points have a large scatter and show a mean evolution with the cosmological parameters. The cyan data points illustrate our procedure for dividing out most of the cosmology dependence, as described in the previous section. However, they still display a large scatter which is dominated by sample variance. Finally, the red data points have the mass function correction applied in addition. It is striking how much tighter they cluster compared to the cyan markers, again indicating that the mass function correction is working rather well. The scatter is relatively independent of cosmology, except for the low- Ω_m points where the halo model approach seems to start to break down. We observe that the remaining scatter in the LH panels significantly exceeds what is observed for the red markers in the CV panels, demonstrating that applying the mass function correction yields an LH set that is dominated by signal from the feedback parameters. Indeed, the neural networks trained on these denoised LH simulations pass the null test described in Sec. IVA.

It is worth noting that we tried the described mass function correction procedure for SIMBA as well. However, for these simulations it did not work well at all, indicating that halos in SIMBA are not well approximated by the fitting formulas we used in the halo model (this may be related to the very long-range effects feedback has in SIMBA [81]). Fortunately, the signal from feedback parameters is strong enough in SIMBA that it was possible to extract from the LH data with full sample variance contamination.

V. RESULTS AND DISCUSSION

In the previous section, we have constructed interpolating neural networks which return the y-distortion monopoles $\langle y \rangle$ and $\langle T_e \rangle$ as functions of astrophysical feedback parameters $A_{\rm SN1}$, $A_{\rm AGN1}$, $A_{\rm SN2}$, $A_{\rm AGN2}$ and cosmological parameters Ω_m , σ_8 . We now aim to explore the dependence on the astrophysical parameters and translate it into forecast constraints assuming the spectral distortion measurement described in Sec. III.

A. Parameter dependence

We first explore how $\langle y \rangle$ and $\langle T_e \rangle$ depend on individual feedback parameters with the remaining ones (as well as σ_8 and Ω_m) fixed. In Fig. 4, we show as blue markers the measurements from the CAMELS 1P set, rescaled such that, their value at the fiducial point, is equal to that from the corresponding large box (indicated by the red crosses). The blue lines in Fig. 4 are evaluations of the trained neural nets. Agreement between the 1P results and the neural nets is generally better for SIMBA, because the observables depend much more strongly on the feedback parameters than in IllustrisTNG. For IllustrisTNG, there are substantial differences, although in many cases the qualitative trends are similar in 1P and the neural nets. As we have argued before, we believe that generally the neural nets are more robust since they marginalize over the simulation initial conditions. However, this may not be true for some of the cases in IllustrisTNG for which the dependence is comparatively weak, because the neural net may have little incentive to accurately capture this dependence. There are, however, strong counterarguments against this hypothesis. It seems clear from the SIMBA panels that in almost all cases the SN and AGN parameter pairs each work in the same direction. This behavior is mirrored in IllustrisTNG, also in the cases where strong deviations from the 1P set are observed. We emphasize that our neural net architecture has no implicit preference for grouping the parameters in such a way. Furthermore, we have evidence that the initial conditions for the 1P set happen to be somewhat anomalous. Visually, this manifests itself by the presence of an unusually large void in the z = 0 density field. More quantitatively, this can be seen in Fig. 1, where the 1P data points (for both IllustrisTNG and SIMBA) tend to fall at higher $\langle T_e \rangle$ than the mean trend in the LH set. Finally, as mentioned in Sec. IVA, we have performed a nontrivial null test on the trained networks which is passed satisfactorily (after correcting for sample variance in the case of IllustrisTNG). Thus, we believe that in fact for all cases the neural nets are closer to reality than the 1P set.

Turning now to physical interpretation, we observe striking differences between SIMBA and IllustrisTNG. In the case of $\langle y \rangle$, IllustrisTNG has positive slopes for all four feedback parameters, while for SIMBA the SN parameters exhibit negative slopes. It seems rather natural that slopes should generally be positive because integrated electron pressure is a measure of thermal energy. However, it has also been established that stronger supernova feedback can counteract the AGN contribution, since the outflow of matter due to stellar feedback limits the black hole growth rate [54] (see also [82]). Reference [82] used the OWLS subgrid model which in terms of $\langle y \rangle$ and $\langle T_e \rangle$ is closer to SIMBA than to IllustrisTNG. This is consistent with the frequency of negative slopes with A_{SNi} in Fig. 4.

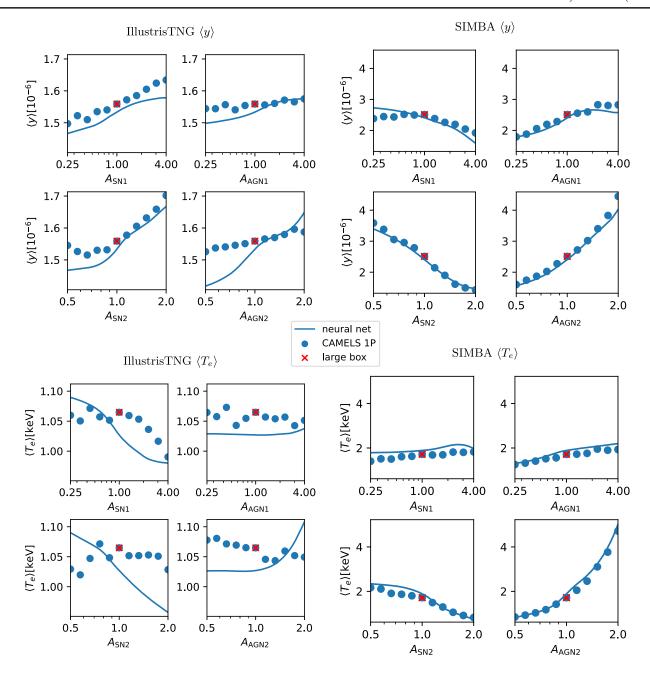


FIG. 4. Dependence of spectral distortions on individual simulation feedback parameters. In each panel, all cosmological and astrophysical parameters are held fixed apart from the one shown on the horizontal axis. The blue markers are measurements from the CAMELS 1P set, rescaled such that at the fiducial model they agree with the measurements from the larger boxes, which are indicated by the red crosses. The blue lines are evaluations of neural networks trained on the CAMELS LH set. We argue in the text that the LH-trained networks are likely more robust than the 1P data in most cases. Note that the vertical plot ranges are identical within each quadrant.

Interestingly, in some cases there are saddle points. That these are probably mostly real and not artifacts in the neural networks is supported by the fact that they are also visible in the 1P markers. Fortunately, these saddle points are generally away from the fiducial model (except for the AGN dependence of $\langle T_e \rangle$ in IllustrisTNG), so they will not contaminate the Fisher forecasts significantly.

As we show in Appendix E, the leading-order behavior of the *y*-distortion observables as a function of the four feedback parameters is well described by a factorization,

$$x_i = \prod_j g_i^{(j)}(A_j), \tag{12}$$

where the $g_i^{(j)}$ are identical to the 1-parameter functions in Fig. 4 up to normalization.

As a cross-check, we also test our results using another machine learning tool called symbolic regression. The results, presented in Appendix F, support the factorization approximation Eq. (12).

B. Fisher forecasts

We now turn to forecasting constraints on the simulation feedback parameters given our fiducial model for the PIXIE constraints on $\langle y \rangle$, $\langle T_e \rangle$ presented in Sec. III. It should be emphasized that the primary goal of this exercise is to demonstrate the substantial constraining power of a realistic y-distortion measurement on our understanding of astrophysical feedback processes. In particular, the uncertainty on the fiducial model implies that the constraints presented in the following should only be interpreted as qualitative indicators. We make the approximation that the cosmological model (specifically Ω_m and σ_8) is known perfectly; in comparison to our ignorance on the subgrid model this is a good assumption.

As we have already seen before, the feedback parameters have very different consequences in IllustrisTNG and SIMBA. We will need to assume that the chosen priors on these parameters are reasonable in terms of other observables, so that a direct comparison between the constraints (in units of the prior) makes any sense at all.

First, we construct the 4×4 Fisher matrices for IllustrisTNG and SIMBA as usual,

$$F_{ab} = \frac{\partial x^i}{\partial A^a} C_{ij}^{-1} \frac{\partial x^j}{\partial A^b},\tag{13}$$

with C the 2×2 covariance matrix, marginalized over all foreground nuisance parameters (cf. Sec. III), and the derivatives are computed from the neural network interpolators with finite difference step sizes chosen to match the scale of the posterior. Of course, these Fisher matrices are degenerate since two measurements are not sufficient to constrain four parameters.

For each simulation type, by diagonalizing the Fisher matrix we identify the best-constrained orthogonal parameter combinations as

$$\begin{split} A_{\text{IllustrisTNG}}^{(1)} &= A_{\text{SNI}}^{+0.20} A_{\text{AGN1}}^{+0.13} A_{\text{SN2}}^{+0.77} A_{\text{AGN2}}^{+0.59}, \\ A_{\text{IllustrisTNG}}^{(2)} &= A_{\text{SNI}}^{+0.52} A_{\text{AGN1}}^{-0.10} A_{\text{SN2}}^{+0.44} A_{\text{AGN2}}^{-0.73}, \\ A_{\text{SIMBA}}^{(1)} &= A_{\text{SNI}}^{-0.17} A_{\text{AGN1}}^{+0.30} A_{\text{SN2}}^{-0.72} A_{\text{AGN2}}^{+0.61}, \\ A_{\text{SIMBA}}^{(2)} &= A_{\text{SNI}}^{+0.42} A_{\text{AGN1}}^{-0.37} A_{\text{SN2}}^{+0.36} A_{\text{AGN2}}^{+0.74}. \end{split} \tag{14}$$

Note that our methodology of finding the most constrained parameter combination(s) is similar to that used in galaxy photometric survey studies for identifying the $S_8 \equiv \sigma_8 \Omega_m^{\alpha}$ parameter combination [83,84]. The induced 2×2 Fisher

matrix on the thus identified subspaces tangent to the fiducial model is not singular. Of course, these combinations are not unique since the Fisher matrix only gives linear-order information. Not surprisingly given the intuition from the upper half of Fig. 4, in $A^{(1)}$ the powers have identical signs for IllustrisTNG, while for SIMBA the SN and AGN parameters have powers of opposite signs.

We compute the 68% constraints

$$\sigma(A_{
m IllustrisTNG}^{(1)}) = 0.015, \qquad \sigma(A_{
m SIMBA}^{(1)}) = 0.0024, \ \sigma(A_{
m IllustrisTNG}^{(2)}) = 1.3, \qquad \sigma(A_{
m SIMBA}^{(2)}) = 0.075.$$

Thus, a PIXIE-like experiment could place percent-level constraints on parameter combinations that are currently only known to little better than an order of magnitude. In the case of SIMBA the measurement of the relativistic component would also allow a $\sim\!10\%$ constraint on a second parameter combination. This is not possible for IllustrisTNG, due to the extremely small variation of $\langle T_e \rangle$. Conversely, this implies that a measurement of $\langle T_e \rangle$ significantly different from the fiducial model has the potential to simply rule out the CAMELS-based implementation of the IllustrisTNG model.

1. Robustness check

We have argued before (Sec. VA) that we believe the neural network interpolators to be more robust than the 1P data, particularly for IllustrisTNG where they disagree substantially. As a robustness check, we have repeated the Fisher analysis using derivatives estimated from onedimensional interpolators through the 1P set. In the case of IllustrisTNG, the error bar on the best-constrained parameter combination inflates by a factor ~3, while the orthogonal combination remains unconstrained. The slight inflation is due to the derivatives from the 1P set being slightly shallower than those from the neural net, see Fig. 4. In the case of SIMBA, the error bars increase by about 20%. The corresponding parameter combinations differ from the ones listed in Eq. (14), but the qualitative features (relative magnitudes and signs of the exponents) are very similar. These comparisons indicate that our results for IllustrisTNG are robust within a factor of a few while those for SIMBA are very accurate.

VI. CONCLUSIONS

This work is one of the first systematic studies of the information content in the late-time Sunyaev-Zel'dovich spectral distortions. Besides a simulation-based forecast in the context of specific subgrid models, in the Appendices we have also given some novel theoretical results regarding the IGM and reionization contributions.

We have measured the nonrelativistic and relativistic mean spectral distortion amplitudes $\langle y \rangle$ and $\langle T_e \rangle$ in the CAMELS simulations suite. By training neural networks

on this data, we have constructed interpolators returning the two signals as a function of Ω_m , σ_8 , and four feedback parameters. In the case of IllustrisTNG, the observables depend only weakly on feedback, necessitating the use of halo model-derived correction factors to reduce the large sample variance due to the small CAMELS box size.

Incidentally, the described method to reduce scatter arising from small simulation boxes by using mass function-dependent scaling factors should be more generally applicable to many works concerned with fields that can be approximated with a halo model. We have also tried to reduce biases as much as possible by matching our data to the comparatively much larger size flagship IllustrisTNG and SIMBA simulations.

Using the interpolating neural networks, we have performed a Fisher forecast assuming a Gaussian posterior on $\langle y \rangle$ and $\langle T_e \rangle$, which had been computed assuming a PIXIE-like experiment and a realistic foreground model. Of course, our work is relevant for experiments other than PIXIE, e.g., ESA's Voyage 2050 large-scale proposal (see e.g., [31]), which is expected to reach even tighter constraints on distortion parameters.

We find that in the case of IllustrisTNG, only a single parameter combination can be strongly constrained, at the $\sim\!2\%$ level. On the other hand, in the case of SIMBA, the availability of two data points enables two orthogonal combinations to be measured, to $\sim\!0.2\%$ and $\sim\!8\%$. We emphasize again that although the feedback parameters are qualitatively similar between IllustrisTNG and SIMBA, their detailed meaning differs substantially so direct comparisons must be carried out with caution.

Given the limited information content from only two observables, it is interesting to ask what other measurements could be added in order to improve the constraints presented here. The tSZ and kinematic SZ effect (kSZ) profiles of halos are already being used in order to assess the viability of simulation models, and Ref. [85] demonstrates using CAMELS that next-generation CMB experiments could place constraints on the feedback parameters. Similarly, Ref. [61] shows that deviations from selfsimilarity in the integrated tSZ flux—halo mass relation $(Y_{SZ} - M)$ —can also be used to constrain feedback parameters. A combination of such constraints could substantially improve upon the error bars presented in this work. Furthermore, the CAMELS feedback parameters also affect observables beyond the SZ effects, like the properties of galaxies. Typically observations of such astrophysical nature are difficult to propagate into hard posteriors, but at least qualitatively the check for simultaneous viability of a given subgrid model should be very useful for simulators.

The main source of uncertainty in the Fisher forecast is likely due to errors in the interpolators. We have performed extensive consistency checks on the trained neural nets, but the limited number of data points in the 1,000 CAMELS

LH simulations places fundamental limits in the possible accuracy. For this reason, the given constraints have some associated uncertainty, of order unity for IllustrisTNG and at the 10% level for SIMBA. Nonetheless, our results are indicative of the transformative effect that measurements of the low-redshift spectral distortions could have for our understanding of baryonic feedback.

ACKNOWLEDGMENTS

L. T. thanks Will Coulton, Shy Genel, and David Spergel for useful discussions. D. W. gratefully acknowledges support from the Friends of the Institute for Advanced Study Membership. The Flatiron Institute is supported by the Simons Foundation. J. C. H. acknowledges support from NSF Grant No. AST-2108536. D. A.-A. was supported in part by NSF Grants No. AST-2009687 and No. AST-2108944. J. C. was supported by the Royal Society as a University Research fellow (No. URF/R/191023) and by the ERC Consolidator Grant *CMBSPEC* (No. 725456). N. B. acknowledges support from NSF Grant No. AST-1910021 and NASA Grants No. 21-ADAP21-0114 and No. 21-ATP21-0129.

APPENDIX A: ELECTRON TEMPERATURE AND THE RELATIVISTIC SZ EFFECT

A cluster's thermal SZ contribution for a given line of sight can be written as

$$\Delta I_{\nu}(\hat{n}) = I_0 \int N_e(\hat{n}, l) \sigma_{\rm T} S_{\nu}(T_e(\hat{n}, l)) dl \qquad (A1)$$

where $I_0=B_{\nu}(T_0)$ is the CMB blackbody spectrum, N_e denotes the electron number density, dl parametrizes the line of sight integration and $S_{\nu}(T)$ determines the SZ spectrum with relativistic temperature corrections. If the electron temperature is constant along the line of sight, $T_e(\hat{n},l)$ = $T_e(\hat{n})$, one can simply write $\Delta I_{\nu}(\hat{n}) = \tau(\hat{n})I_0S_{\nu}(T_e(\hat{n}))$, where the Thomson optical depth is $\tau(\hat{n}) = \int N_e(\hat{n},l)\sigma_{\rm T}dl$. However, generally the temperature varies along the line of sight, such that a moment expansion provides a simpler method for analyzing and describing the SZ signal [86].

Defining $S_{\nu}^{(k)}(T) \equiv \partial^k S_{\nu}(T)/\partial T^k$, we can perform the moment expansion of the SZ signal around the τ -weighted temperature, $T_e^{\tau}(\hat{n}) = \int N_e(\hat{n}, l) \sigma_T T_e(\hat{n}, l) dl/\tau(\hat{n})$. Up to second order in temperature, this yields

$$\Delta I_{\nu}(\hat{n}) \approx \tau I_{0} \left\{ S_{\nu}(T_{e}^{\tau}) + \frac{1}{2} S_{\nu}^{(2)}(T_{e}^{\tau}) [\langle T_{e}^{2} \rangle_{\tau} - (T_{e}^{\tau})^{2}] \right\}, \ \ (A2)$$

where we suppressed the dependence on \hat{n} and introduced the τ -weighted temperature moments

$$\langle T_e^k \rangle_{\tau} = \frac{\int N_e(\hat{n}, l) \sigma_{\rm T} T_e^k(\hat{n}, l) dl}{\tau(\hat{n})} \tag{A3}$$

such that $T_e^{\tau} \equiv \langle T_e \rangle_{\tau}$. For the standard ΛCDM cosmology the first two temperature moments from halos alone are $\langle T_e \rangle_{\tau} = 0.208 \text{ keV}$ and $\langle T_e^2 \rangle_{\tau} = 0.299 \text{ keV}^2$ [15].

Alternatively, we can perform the moment expansion around the y-weighted temperature, T_e^y . By introducing the y-weighted moments

$$\langle T_e^k \rangle_y = \frac{\int N_e(\hat{n}, l) \sigma_{\rm T} \frac{T_e(\hat{n}, l)}{m_e} T_e^k(\hat{n}, l) dl}{y(\hat{n})} \tag{A4}$$

with $y(\hat{n}) = \int N_e(\hat{n}, l) \sigma_{\rm T} \frac{T_e(\hat{n}, l)}{m_e} dl$ and $T_e^y \equiv \langle T_e \rangle_y$, again to second order in the temperature one then has

$$\Delta I_{\nu}(\hat{n}) \approx y I_0 \left\{ \tilde{S}_{\nu}(T_e^y) + \frac{1}{2} \tilde{S}_{\nu}^{(2)}(T_e^y) [\langle T_e^2 \rangle_y - (T_e^y)^2] \right\}. \ (\text{A5})$$

where $\tilde{S}_{\nu} = S_{\nu}/\Theta_{e}$, $\tilde{S}_{\nu}^{(k)} = \partial^{k}\tilde{S}_{\nu}(T)/\partial T^{k}$, and we introduced the dimensionless temperature $\Theta_{e} = \frac{T_{e}}{m_{e}}$. These two representations are essentially equivalent⁶; however, the latter is slightly more economic when it comes to capturing the relativistic tSZ effect, as we will see next.

Low temperature limit and mix of hot and cold gas: In the low temperature limit one can write (e.g., [9,11])

$$S_{\nu}(T) \approx \Theta_{e} Y_{0}(x) + \Theta_{e}^{2} Y_{1}(x)$$
 (A6)

where $x = \frac{h\nu}{T_0}$ and the functions Y_0 and Y_1 are the first two terms of the asymptotic expansion for the tSZ signal (e.g., see [11]). It is clear that in this limit, only two independent spectral parameters can be determined. This is directly evident when using the y-weighted moments with $\tilde{S}_{\nu}^{(2)}(T) \approx 0$:

$$\frac{\Delta I_{\nu}(\hat{n})}{I_0} \approx y S_{\nu}(T_e^y) = y \left[Y_0(x) + \frac{T_e^y}{m_e} Y_1(x) \right].$$
 (A7)

Since $T_e^y \equiv \langle T_e^2 \rangle_\tau / \langle T_e \rangle_\tau$, the two important parameters for the spectral analysis in the Λ CDM case are expected to be $y \approx 1.77 \times 10^{-6}$ and $T_e^y \approx 1.44$ keV [15]. However, the presence of extremely cold gas from reionization modifies the observational inference. This can be seen by adding another y-distortion contribution with no relativistic correction

$$\frac{\Delta I_{\nu}^{\text{tot}}(\hat{n})}{I_0} \approx y^{\text{re}} Y_0(x) + y \left[Y_0(x) + \frac{T_e^y}{m_e} Y_1(x) \right]. \tag{A8}$$

The reionization y parameter, y^{re} , is roughly 10% of the cluster contribution [15,87] and cannot be distinguished from the cluster contribution with distortion measurements alone. In an analysis, an effective electron temperature of $T_e^{\text{y,*}} = T_e^{\text{y}}/(1 + y^{\text{re}}/y) \approx 1.30 \text{ keV}$ would thus be

recovered. Due to small contributions from higher order temperature corrections, the recovered result for Λ CDM is $T_e^{y,*} \approx 1.24$ keV [8].

APPENDIX B: COSMOLOGY DEPENDENCE

Using the halo model we compute how the distortion monopoles $\langle y \rangle$ and $\langle T_e \rangle$ depend on the relevant Λ CDM parameters. We use the halo model fitting formulas described in Sec. IV B to compute the observables $x_i = \{\langle y \rangle, \langle T_e \rangle\}$ as functions of $\theta_j = \{h, \Omega_m, \Omega_b, n_s, \sigma_8\}$. The resulting functions are well approximated by power laws, $x_i = x_i^{(0)} (\theta_j/\theta_j^{(0)})^{\alpha_{ij}}$. The resulting one-parameter fits are given by

$$10^{6}\langle y\rangle(h) = 1.48(h/0.6711)^{1.59},$$

$$keV^{-1}\langle T_{e}\rangle(h) = 1.51(h/0.6711)^{-0.86},$$

$$10^{6}\langle y\rangle(\Omega_{m}) = 1.48(\Omega_{m}/0.3)^{0.92},$$

$$keV^{-1}\langle T_{e}\rangle(\Omega_{m}) = 1.51(\Omega_{m}/0.3)^{0.36},$$

$$10^{6}\langle y\rangle(\Omega_{b}) = 1.48(\Omega_{b}/0.0490)^{0.84},$$

$$keV^{-1}\langle T_{e}\rangle(\Omega_{b}) = 1.51(\Omega_{b}/0.0490)^{0.09},$$

$$10^{6}\langle y\rangle(n_{s}) = 1.48(n_{s}/0.9624)^{1.52},$$

$$keV^{-1}\langle T_{e}\rangle(n_{s}) = 1.51(n_{s}/0.9624)^{-0.84},$$

$$10^{6}\langle y\rangle(\sigma_{8}) = 1.48(\sigma_{8}/0.8)^{3.86},$$

$$keV^{-1}\langle T_{e}\rangle(\sigma_{8}) = 1.51(\sigma_{8}/0.8)^{1.93}.$$
(B1)

Note that the given scalings have a number of uncertainties. First, they only include the halo contribution. Second, the assumed cluster temperature model may not be optimal, as discussed in Sec. IV B. Thus, we caution against blind use of these equations. However, for small deviations from the pivot cosmology they should provide reasonable approximations for the slope, even though the prefactors are most likely not very useful.

APPENDIX C: REIONIZATION CONTRIBUTION

The nonrelativistic $\langle y \rangle$ receives a small contribution from the epoch of reionization, which is not modeled in our simulations. This Appendix describes an estimate of the magnitude and uncertainty of this effect. We calculate the $\langle y \rangle$ from reionization using maps constructed by ray tracing through the past light cone of a semianalytic realization from z=5.5 to z=20, which defines the redshift range we consider for reionization. In this semianalytic realization, reionization fields are constructed on a gravity-only simulation following the method in Ref. [88]. For each spatial cell in the *N*-body simulation we have a density and ionization state as a function of time. We set the initial temperature of the cells when they reionize as $T_0=2\times 10^4$ K, then the *i*th cell cools adiabatically according to

⁶They become indistinguishable when more temperature terms are included in the expansion.

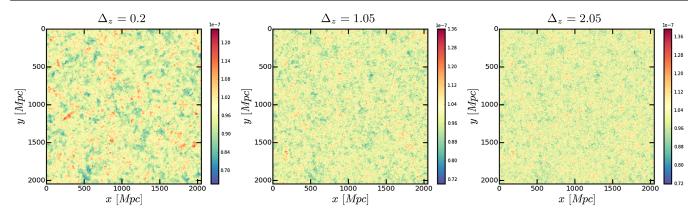


FIG. 5. Maps of the Compton-y field from reionization, generated using a semianalytic model on a gravity-only simulation. The panels assume different durations of reionization, with the short, fiducial, and long duration models from left to right. Note the slightly different color scales.

$$\frac{T_i(z)}{T_0} = \frac{1+z}{1+z_{i,RE}},$$
 (C1)

where $z_{i,RE}$ is the redshift at which the given cell reionizes. Using the parametric reionization model from Ref. [88], we change the duration and midpoint of reionization from the fiducial parameters. We show Compton-y maps from reionization for our fiducial model and the two extreme duration models, short and long duration, in Fig. 5. For the fiducial reionization model, with a median reionization redshift of $z_{\rm mid} = 10$, where 50% of the Universe has reionized by mass, and a duration parameter $\Delta_z = 1.05$, we find that $\langle y \rangle_{\rm reio} = 9.9 \times 10^{-8}$. For the short ($\Delta_z = 0.2$) and long ($\Delta_z = 2.05$) duration models we find 9.6×10^{-8} and 1.0×10^{-7} , respectively, while the dependence on $z_{\rm mid}$ (which we varied in [8, 12]) is smaller. Thus, reionization contributes less than 10% to the total signal and we estimate the uncertainty as $\sim 5 \times 10^{-9}$, significantly below the error budget for our assumed experiment.

Clearly, the relativistic distortion receives minuscule contributions from reionization as $T_e/m_e \lesssim 10^{-4}$.

APPENDIX D: IGM CONTRIBUTION

In this Appendix we compute the contribution from the IGM to the $\langle y \rangle$ signal. We define this quantity as all Compton y generated at z < 5.5 (which marks the end of reionization, cf. Appendix C) outside of any halo. Since the IGM pressure is low compared to the ICM, a reasonable approximation is

$$\langle y \rangle_{\text{IGM}} = \frac{\sigma_{\text{T}}}{m_e} \int dl \, \bar{n}_e(z) T_{\text{IGM}}(z)$$
 (D1)

where the integration is over distance up to z=5.5, $\bar{n}_e(z) \equiv x_e(z)\Omega_b\rho_{\rm crit}(z)$, and $x_e(z)$ is the free electron fraction. In our fiducial model, we assume $T_{\rm IGM,0}=2\times 10^4$ K for the IGM temperature at the end of reionization. We then assume the temperature drops adiabatically. The

function $x_e(z)$ will depend on the redshift $z_{\rm HeII}$ at which HeII reionizes. In our fiducial model we assume instantaneous HeII reionization at $z_{\rm HeII}=3.5$ such that

$$x_e(z) = \begin{cases} \frac{3X_H + 1}{4}; & z > z_{\text{HeII}}, \\ \frac{X_H + 1}{2}; & z < z_{\text{HeII}}, \end{cases}$$
 (D2)

where X_H is the primordial hydrogen mass fraction. The fiducial model yields $\langle y \rangle_{\rm reio} = 7.1 \times 10^{-8}$. Changes in the starting value of $T_{\rm IGM}$ linearly affect the Compton y. Assuming $z_{\rm HeII} = 2.5$ results in 7×10^{-8} and $z_{\rm HeII} = 4.5$ gives 7.2×10^{-8} . Thus, the IGM contribution is in magnitude comparable to the reionization signal with somewhat larger theoretical uncertainty driven by the temperature normalization.

APPENDIX E: TWO-PARAMETER DEPENDENCE

In this Appendix, we consider dependence on pairs of parameters. The primary goal of this exercise is to establish how strongly couplings between feedback parameters affect the *y* observables. In Fig. 6, in each panel we plot the quantity

$$\frac{x_i(A_j, A_k)}{\operatorname{const} \times x_i(A_i)x_i(A_k)} - 1, \tag{E1}$$

where the constant normalizes such that the ratio is one at the fiducial point and the x_i are evaluations of the neural nets with either one or two parameters varied from the fiducial point. Thus, we illustrate deviations from perfect factorization. We observe that for IllustrisTNG the factorization is a rather good approximation, with couplings of at most 5%. SIMBA exhibits stronger corrections, but the factorization is still relatively accurate to within ~40% and of course the variations are also much larger by about an order of magnitude. In units of the overall differences in the y observables the inter-parameters couplings are quite similar in IllustrisTNG and SIMBA. It must be noted that the

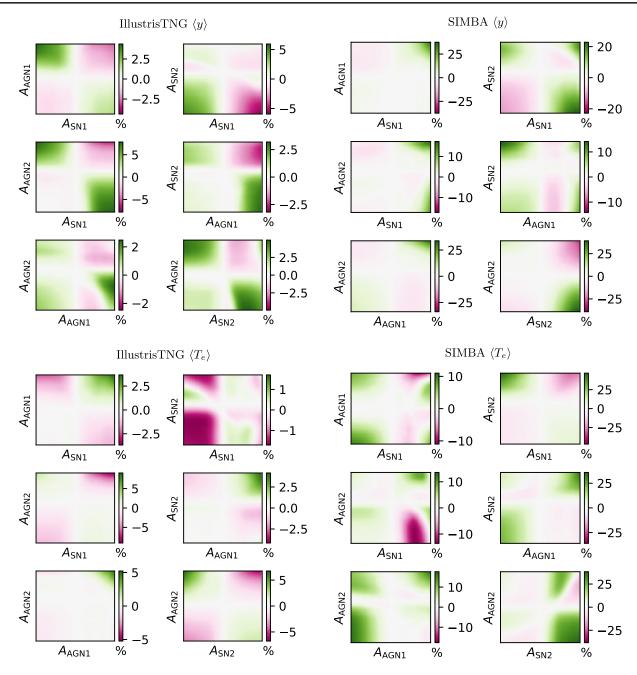


FIG. 6. Dependence of spectral distortions on pairs of simulation feedback parameters. Each panel shows evaluations of the neural networks trained on the LH set for two astrophysical parameters varied, divided by the simple model in which the dependence factorizes (using the curves from Fig. 4). As can be seen, for IllustrisTNG the factorization approximation is good to within ~5%, while for SIMBA it is somewhat worse but still reasonable. For readability, the axes ticks have been suppressed. They are exactly identical to the ones in Fig. 4, so that the fiducial model with all feedback parameters equal one is in the center of each panel.

edges of parameter space are likely not quite accurately represented by the neural nets. We should emphasize that the neural nets have no structural preference for factorized representations. In summary, most of the dependence of *y* observables on feedback parameters can be readily read off from Fig. 4.

APPENDIX F: SYMBOLIC REGRESSION

Symbolic regression identifies equations with parsimonious combinations of input parameters that have the smallest scatter with the given quantity of interest [61,89–99]. We employ it to predict $\langle y \rangle$ and $\langle T_e \rangle$ separately as a function of the feedback parameters:

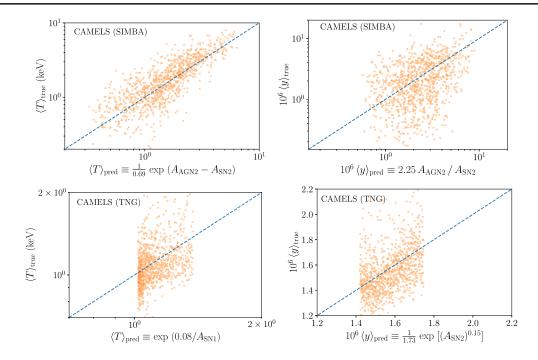


FIG. 7. As a complement to the results from the neural network, we use a machine learning tool called symbolic regression to predict $\langle y \rangle$ and $\langle T_e \rangle$ separately as an analytic function of the feedback parameters. The blue dashed line represents predicted = true; the closer the points are to the line the more accurate the prediction. The equations perform comparatively better for SIMBA. It is also interesting to see that the effects of AGN and supernovae on both $\langle y \rangle$ and $\langle T_e \rangle$ are opposite to each other for the case of SIMBA.

$$\{\langle y \rangle, \langle T_e \rangle\} = f(A_{\text{SN1}}, A_{\text{AGN1}}, A_{\text{SN2}}, A_{\text{AGN2}}) \quad (\text{F1})$$

We show the equations obtained and their performance in Fig. 7. The data points shown are the measurements from the CAMELS LH set (for the case of TNG, the data has been corrected for sample variance and the cosmology dependence has also been removed, cf. Fig. 3). We also obtained equations more complex than the ones in Fig. 7, however, as the risk of overfitting goes up as the equations get more complex, we show simple ones which have a

substantial reduction in the mean squared error. The equations perform much better for the case of SIMBA. This could be either because the dependence of $\langle y \rangle$, $\langle T_e \rangle$ on the feedback parameters is weaker for TNG, or the dependence for the case of TNG might have a very complicated functional form and it is hard to find a good approximation given the limited size of our data set. Overall, the results in Fig. 7 are consistent with the assumption in Eq. (12) that the individual parameter feedback dependence can be factorized.

R. A. Sunyaev and R. Khatri, Int. J. Mod. Phys. D 22, 1330014 (2013).

^[2] J. Chluba, Mon. Not. R. Astron. Soc. 460, 227 (2016).

^[3] Y. B. Zeldovich and R. A. Sunyaev, Astrophys. Space Sci. 4, 301 (1969).

^[4] R. A. Sunyaev and Y. B. Zeldovich, Astrophys. Space Sci. 7, 3 (1970).

^[5] J. E. Carlstrom, G. P. Holder, and E. D. Reese, Annu. Rev. Astron. Astrophys. 40, 643 (2002).

^[6] T. Mroczkowski, D. Nagai, K. Basu, J. Chluba, J. Sayers, R. Adam, E. Churazov, A. Crites, L. Di Mascolo, D. Eckert, J. Macias-Perez, F. Mayet, L. Perotto, E. Pointecouteau, C. Romero, F. Ruppin, E. Scannapieco, and J. ZuHone, Space Sci. Rev. 215, 17 (2019).

^[7] F. Villaescusa-Navarro et al., Astrophys. J. 915, 71 (2021).

^[8] M. H. Abitbol, J. Chluba, J. C. Hill, and B. R. Johnson, Mon. Not. R. Astron. Soc. 471, 1126 (2017).

^[9] S. Y. Sazonov and R. A. Sunyaev, Astrophys. J. 508, 1 (1998).

^[10] A. Challinor and A. Lasenby, Astrophys. J. 499, 1 (1998).

^[11] N. Itoh, Y. Kohyama, and S. Nozawa, Astrophys. J. **502**, 7 (1998).

^[12] J. Chluba, G. Hütsi, and R. A. Sunyaev, Astron. Astrophys. **434**, 811 (2005).

^[13] S. Nozawa, N. Itoh, Y. Suda, and Y. Ohhata, Nuovo Cimento B Ser. 121, 487 (2006).

^[14] J. Chluba, D. Nagai, S. Sazonov, and K. Nelson, Mon. Not. R. Astron. Soc. 426, 510 (2012).

- [15] J. C. Hill, N. Battaglia, J. Chluba, S. Ferraro, E. Schaan, and D. N. Spergel, Phys. Rev. Lett. 115, 261301 (2015).
- [16] S. T. Kay, L. C. Powell, A. R. Liddle, and P. A. Thomas, Mon. Not. R. Astron. Soc. 386, 2110 (2008).
- [17] E. Lee, J. Chluba, S. T. Kay, and D. J. Barnes, Mon. Not. R. Astron. Soc. 493, 3274 (2020).
- [18] K. Jedamzik, V. Katalinić, and A. V. Olinto, Phys. Rev. Lett. 85, 700 (2000).
- [19] K. E. Kunze and E. Komatsu, J. Cosmol. Astropart. Phys. 01 (2014) 009.
- [20] J. Chluba, D. Paoletti, F. Finelli, and J. A. Rubiño-Martín, Mon. Not. R. Astron. Soc. 451, 2244 (2015).
- [21] S. Sarkar and A. M. Cooper, Phys. Lett. 148B, 347 (1984).
- [22] W. Hu and J. Silk, Phys. Rev. Lett. 70, 2661 (1993).
- [23] J. Chluba, Mon. Not. R. Astron. Soc. 436, 2232 (2013).
- [24] Y. Ali-Haïmoud, J. Chluba, and M. Kamionkowski, Phys. Rev. Lett. 115, 071304 (2015).
- [25] B. Bolliet, J. Chluba, and R. Battye, Mon. Not. R. Astron. Soc. 507, 3148 (2021).
- [26] Y. Ali-Haïmoud, Phys. Rev. D 103, 043541 (2021).
- [27] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright, Astrophys. J. 473, 576 (1996).
- [28] A. Kogut, D. J. Fixsen, D. T. Chuss, J. Dotson, E. Dwek, M. Halpern, G. F. Hinshaw, S. M. Meyer, S. H. Moseley, M. D. Seiffert, D. N. Spergel, and E. J. Wollack, J. Cosmol. Astropart. Phys. 07 (2011) 025.
- [29] A. Kogut, J. Chluba, D. J. Fixsen, S. Meyer, and D. Spergel, in *SPIE Conference Series*, Proceedings SPIE, Vol. 9904 (SPIE-International Society for Optical Engineering, Bellingham, WA, 2016), p. 99040W.
- [30] B. Maffei et al., arXiv:2111.00246.
- [31] J. Chluba et al., Exp. Astron. 51, 1515 (2021).
- [32] T. Plagge, B. A. Benson *et al.*, Astrophys. J. **716**, 1118 (2010).
- [33] N. Hand *et al.* (Atacama Cosmology Telescope Collaboration), Astrophys. J. 736, 39 (2011).
- [34] Planck Collaboration et al., Astron. Astrophys. 557, A52 (2013).
- [35] Planck Collaboration et al., Astron. Astrophys. 550, A131 (2013).
- [36] J.C. Hill, B.D. Sherwin, K.M. Smith *et al.* (Atacama Cosmology Telescope Collaboration), arXiv:1411.8004.
- [37] J. P. Greco, J. C. Hill, D. N. Spergel, and N. Battaglia, Astrophys. J. **808**, 151 (2015).
- [38] Planck Collaboration *et al.*, Astron. Astrophys. **594**, A22 (2016).
- [39] V. Vikram, A. Lidz, and B. Jain, Mon. Not. R. Astron. Soc. **467**, 2315 (2017).
- [40] J. P. Dietrich, S. Bocquet, T. Schrabback, D. Applegate, H. Hoekstra, S. Grandis, J. J. Mohr *et al.*, Mon. Not. R. Astron. Soc. **483**, 2871 (2019).
- [41] M. S. Madhavacheril, J. C. Hill, S. Næss *et al.* (Atacama Cosmology Telescope Collaboration), Phys. Rev. D 102, 023534 (2020).
- [42] L. E. Bleem, T. M. Crawford *et al.*, Astrophys. J. Suppl. Ser. **258**, 36 (2022).
- [43] E. Schaan, S. Ferraro, S. Amodeo, N. Battaglia, S. Aiola et al. (Atacama Cosmology Telescope Collaboration), Phys. Rev. D 103, 063513 (2021).

- [44] S. Amodeo, N. Battaglia, E. Schaan, S. Ferraro, E. Moser, S. Aiola *et al.* (Atacama Cosmology Telescope Collaboration), Phys. Rev. D 103, 063514 (2021).
- [45] E. M. Vavagiakis, P. A. Gallardo, V. Calafut, S. Amodeo, S. Aiola *et al.* (Atacama Cosmology Telescope Collaboration), Phys. Rev. D 104, 043503 (2021).
- [46] C. T. Pratt, Z. Qu, and J. N. Bregman, Astrophys. J. 920, 104 (2021).
- [47] G. Hurier and C. Tchernin, Astron. Astrophys. 604, A94 (2017).
- [48] J. Erler, K. Basu, J. Chluba, and F. Bertoldi, Mon. Not. R. Astron. Soc. 476, 3360 (2018).
- [49] M. Remazeilles and J. Chluba, Mon. Not. R. Astron. Soc. 494, 5734 (2020).
- [50] D. J. Croton, V. Springel, S. D. M. White, G. De Lucia, C. S. Frenk, L. Gao, A. Jenkins, G. Kauffmann, J. F. Navarro, and N. Yoshida, Mon. Not. R. Astron. Soc. 365, 11 (2006).
- [51] R. S. Somerville, G. Popping, and S. C. Trager, Mon. Not. R. Astron. Soc. 453, 4337 (2015).
- [52] V. Springel, T. Di Matteo, and L. Hernquist, Mon. Not. R. Astron. Soc. 361, 776 (2005).
- [53] J. Schaye et al., Mon. Not. R. Astron. Soc. 446, 521 (2015).
- [54] D. Anglés-Alcázar, R. Davé, C.-A. Faucher-Giguère, F. Özel, and P. F. Hopkins, Mon. Not. R. Astron. Soc. 464, 2840 (2017).
- [55] S. Kaviraj, C. Laigle, T. Kimm, J. E. G. Devriendt, Y. Dubois, C. Pichon, A. Slyz, E. Chisari, and S. Peirani, Mon. Not. R. Astron. Soc. 467, 4739 (2017).
- [56] P. F. Hopkins et al., Mon. Not. R. Astron. Soc. 480, 800 (2018).
- [57] M. Vogelsberger, F. Marinacci, P. Torrey, and E. Puchwein, Nat. Rev. Phys. 2, 42 (2020).
- [58] P. J. E. Peebles, Astrophys. J. 142, 1317 (1965).
- [59] W. H. Press and P. Schechter, Astrophys. J. 187, 425 (1974).
- [60] A. Cooray and R. Sheth, Phys. Rep. 372, 1 (2002).
- [61] D. Wadekar, L. Thiele, F. Villaescusa-Navarro, J. C. Hill, D. N. Spergel, M. Cranmer, and S. Ho (to be published).
- [62] Planck Collaboration *et al.*, Astron. Astrophys. **641**, A6 (2020).
- [63] R. Weinberger, V. Springel, and R. Pakmor, Astrophys. J. Suppl. Ser. 248, 32 (2020).
- [64] V. Springel, R. Pakmor, A. Pillepich, R. Weinberger, D. Nelson, L. Hernquist, M. Vogelsberger, S. Genel, P. Torrey, F. Marinacci, and J. Naiman, Mon. Not. R. Astron. Soc. 475, 676 (2018).
- [65] J. P. Naiman, A. Pillepich, V. Springel, E. Ramirez-Ruiz, P. Torrey, M. Vogelsberger, R. Pakmor, D. Nelson, F. Marinacci, L. Hernquist, R. Weinberger, and S. Genel, Mon. Not. R. Astron. Soc. 477, 1206 (2018).
- [66] D. Nelson, A. Pillepich, V. Springel, R. Weinberger, L. Hernquist, R. Pakmor, S. Genel, P. Torrey, M. Vogelsberger, G. Kauffmann, F. Marinacci, and J. Naiman, Mon. Not. R. Astron. Soc. 475, 624 (2018).
- [67] F. Marinacci, M. Vogelsberger, R. Pakmor, P. Torrey, V. Springel, L. Hernquist, D. Nelson, R. Weinberger, A. Pillepich, J. Naiman, and S. Genel, Mon. Not. R. Astron. Soc. 480, 5113 (2018).
- [68] A. Pillepich, D. Nelson, L. Hernquist, V. Springel, R. Pakmor, P. Torrey, R. Weinberger, S. Genel, J. P. Naiman,

- F. Marinacci, and M. Vogelsberger, Mon. Not. R. Astron. Soc. 475, 648 (2018).
- [69] D. Nelson, V. Springel, A. Pillepich, V. Rodriguez-Gomez, P. Torrey, S. Genel, M. Vogelsberger, R. Pakmor, F. Marinacci, R. Weinberger, L. Kelley, M. Lovell, B. Diemer, and L. Hernquist, Comput. Astrophys. Cosmol. 6, 2 (2019).
- [70] P. F. Hopkins, Mon. Not. R. Astron. Soc. 450, 53 (2015).
- [71] R. Davé, D. Anglés-Alcázar, D. Narayanan, Q. Li, M. H. Rafieferantsoa, and S. Appleby, Mon. Not. R. Astron. Soc. 486, 2827 (2019).
- [72] A. Kogut and D. J. Fixsen, J. Cosmol. Astropart. Phys. 05 (2020) 041.
- [73] Planck Collaboration et al., Astron. Astrophys. 594, A10 (2016).
- [74] A. Rotti and J. Chluba, Mon. Not. R. Astron. Soc. 500, 976 (2021).
- [75] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, arXiv:1907.10902.
- [76] J. L. Tinker, B. E. Robertson, A. V. Kravtsov, A. Klypin, M. S. Warren, G. Yepes, and S. Gottlöber, Astrophys. J. 724, 878 (2010).
- [77] M. Davis, G. Efstathiou, C. S. Frenk, and S. D. M. White, Astrophys. J. 292, 371 (1985).
- [78] N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers, Astrophys. J. 758, 75 (2012).
- [79] M. Arnaud, E. Pointecouteau, and G. W. Pratt, Astron. Astrophys. 441, 893 (2005).
- [80] G. L. Bryan and M. L. Norman, Astrophys. J. 495, 80 (1998).
- [81] J. Borrow, D. Anglés-Alcázar, and R. Davé, Mon. Not. R. Astron. Soc. 491, 6102 (2020).
- [82] C. M. Booth and J. Schaye, Sci. Rep. 3, 1738 (2013).
- [83] A. Amon et al., Phys. Rev. D 105, 023514 (2022).

- [84] B. Jain and U. Seljak, Astrophys. J. 484, 560 (1997).
- [85] E. Moser, N. Battaglia, D. Nagai, E. Lau, L. F. Machado Poletti Valle, F. Villaescusa-Navarro, S. Amodeo, D. Angles-Alcazar, G. L. Bryan, R. Dave, L. Hernquist, and M. Vogelsberger, arXiv:2201.02708.
- [86] J. Chluba, E. Switzer, K. Nelson, and D. Nagai, Mon. Not. R. Astron. Soc. 430, 3054 (2013).
- [87] W. Hu, D. Scott, and J. Silk, Phys. Rev. D 49, 648 (1994).
- [88] N. Battaglia, H. Trac, R. Cen, and A. Loeb, Astrophys. J. **776**, 81 (2013).
- [89] D. Wadekar, L. Thiele, F. Villaescusa-Navarro, J. C. Hill, M. Cranmer, D. N. Spergel, N. Battaglia, D. Anglés-Alcázar, L. Hernquist, and S. Ho, arXiv:2201.01305.
- [90] M. D. Schmidt and H. Lipson, Science 324, 81 (2009).
- [91] D. Wadekar, F. Villaescusa-Navarro, S. Ho, and L. Perreault-Levasseur, arXiv:2012.00111.
- [92] M. Cranmer, A. Sanchez-Gonzalez, P. Battaglia, R. Xu, K. Cranmer, D. Spergel, and S. Ho, arXiv:2006.11287.
- [93] M. D. Cranmer, R. Xu, P. Battaglia, and S. Ho, arXiv: 1909.05862.
- [94] H. Shao, F. Villaescusa-Navarro, S. Genel, D. N. Spergel, D. Angles-Alcazar, L. Hernquist, R. Dave, D. Narayanan, G. Contardo, and M. Vogelsberger, Astrophys. J. 927, 85 (2022).
- [95] S.-M. Udrescu and M. Tegmark, Sci. Adv. 6, eaay2631 (2020).
- [96] T. Wu and M. Tegmark, Phys. Rev. E 100, 033311 (2019).
- [97] S. Kim, P. Y. Lu, S. Mukherjee, M. Gilbert, L. Jing, V. Čeperić, and M. Soljačić, IEEE. Trans. Neural Networks Learn. Syst. 32, 4166 (2021).
- [98] Z. Liu and M. Tegmark, Phys. Rev. Lett. 126, 180604 (2021).
- [99] C. Wilstrup and J. Kasak, arXiv:2103.15147.