



Calibrating Cosmological Simulations with Implicit Likelihood Inference Using Galaxy Growth Observables

Yongseok Jo¹ , Shy Genel^{2,3} , Benjamin Wandelt^{2,4} , Rachel S. Somerville^{2,5}, Francisco Villaescusa-Navarro^{2,6} , Greg L. Bryan^{2,7} , Daniel Anglés-Alcázar^{2,8} , Daniel Foreman-Mackey² , Dylan Nelson⁹ , and Ji-hoon Kim¹

¹Center for Theoretical Physics, Department of Physics and Astronomy, Seoul National University, Seoul 08826, Republic of Korea; kerex@snu.ac.kr

²Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY, 10010, USA

³Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY, 10027, USA

⁴Sorbonne Université, CNRS, UMR 7095, Institut d’Astrophysique de Paris, 98 bis boulevard Arago, F-75014 Paris, France

⁵Department of Physics and Astronomy, Rutgers University, 136 Frelinghuysen Road, Piscataway, NJ 08854, USA

⁶Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544 USA

⁷Department of Astronomy, Columbia University, 550 West 120th Street, New York, NY 10027, USA

⁸Department of Physics, University of Connecticut, 196 Auditorium Road, U-3046, Storrs, CT 06269-3046, USA

⁹Universität Heidelberg, Zentrum für Astronomie, Institut für Theoretische Astrophysik, Albert-Ueberle-Str. 2, D-69120 Heidelberg, Germany

Received 2022 July 18; revised 2022 November 21; accepted 2022 December 4; published 2023 February 13

Abstract

In a novel approach employing implicit likelihood inference (ILI), also known as likelihood-free inference, we calibrate the parameters of cosmological hydrodynamic simulations against observations, which has previously been unfeasible due to the high computational cost of these simulations. For computational efficiency, we train neural networks as emulators on ~ 1000 cosmological simulations from the CAMELS project to estimate simulated observables, taking as input the cosmological and astrophysical parameters, and use these emulators as surrogates for the cosmological simulations. Using the cosmic star formation rate density (SFRD) and, separately, the stellar mass functions (SMFs) at different redshifts, we perform ILI on selected cosmological and astrophysical parameters (Ω_m , σ_8 , stellar wind feedback, and kinetic black hole feedback) and obtain full six-dimensional posterior distributions. In the performance test, the ILI from the emulated SFRD (SMFs) can recover the target observables with a relative error of 0.17% (0.4%). We find that degeneracies exist between the parameters inferred from the emulated SFRD, confirmed with new full cosmological simulations. We also find that the SMFs can break the degeneracy in the SFRD, which indicates that the SMFs provide complementary constraints for the parameters. Further, we find that a parameter combination inferred from an observationally inferred SFRD reproduces the target observed SFRD very well, whereas, in the case of the SMFs, the inferred and observed SMFs show significant discrepancies that indicate potential limitations of the current galaxy formation modeling and calibration framework, and/or systematic differences and inconsistencies between observations of the SMFs.

Unified Astronomy Thesaurus concepts: Nonparametric inference (1903); Likelihood ratio test (1942); Magnetohydrodynamical simulations (1966); Cosmological parameters (339); Neural networks (1933); Bayes’ Theorem (1924); Galaxy formation (595)

1. Introduction

The significant progress of cosmological simulations and observations has greatly improved our understanding of a wide variety of phenomena, such as the formation and evolution of the large-scale structure. N -body simulations have successfully simulated the formation and evolution of the large-scale structure of the Λ CDM universe (Springel et al. 2005; Boylan-Kolchin et al. 2009; Klypin et al. 2011, 2016). Furthermore, (magneto) hydrodynamic simulations that include comprehensive subgrid models such as star formation, stellar winds, and active galactic nucleus (AGN) feedback have been performed in a cosmological context and have made significant strides toward reproducing a realistic galaxy population across a range of cosmic epochs (see Vogelsberger et al. 2020 for a review). These include ILLUSTRIS (Genel et al. 2014; Vogelsberger et al. 2014a, 2014b; Nelson et al. 2015), cosmo-OWLS (Le Brun et al. 2014), MAGNETICUM (Hirschmann et al. 2014; Ragagnin et al. 2017), HORIZON-AGN (Dubois et al. 2014), MASSIVEBLACK-II (Khandai et al. 2015),

EAGLE (Schaye et al. 2015), BLUETIDES (Feng et al. 2016), MUFASA (Dave et al. 2016), ROMULUS (Tremmel et al. 2017), BAHAMAS (McCarthy et al. 2017), SIMBA (Dave et al. 2019), ILLUSTRIS-TNG (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018a; Springel et al. 2018), HORIZON RUN 5 (Lee et al. 2021), and ASTRID (Bird et al. 2022; Ni et al. 2022).

Meanwhile, wide-field and deep surveys have identified samples of many thousands of nearby and distant galaxies, respectively (Aihara et al. 2018; Venhola et al. 2018; Kuijken et al. 2019; Brunner et al. 2022). In addition, high-resolution imaging and spectroscopy have enabled investigations into the structure and kinematics of galaxies (Dey et al. 2019; Wang et al. 2021). These observational breakthroughs have enabled not only access to a plethora of galaxies from which to construct global distributions of galaxy properties, such as cosmic star formation history (Madau & Dickinson 2014) and galaxy stellar mass functions (SMFs) at different redshifts (Baldry et al. 2012; Leja et al. 2020; McLeod et al. 2021), but also the exploration of many dimensions of galaxy properties, leading to scaling relationships such as the (baryonic) Tully–Fisher relation (Tully & Fisher 1977; McGaugh et al. 2000, 2021), the mass–metallicity relation (Lequeux et al. 1979; Gallazzi et al. 2005; Fontanot et al. 2021), the star-forming



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

sequence (Noeske et al. 2007; Speagle et al. 2014; Leja et al. 2022), the size–mass relation (Shen et al. 2003; van der Wel et al. 2014; Mowla et al. 2019), and relations between galaxy properties and the mass of the central massive black hole (Kormendy & McClure 1993; Kormendy & Gebhardt 2001; Merritt & Ferrarese 2001).

The remarkable progress of simulations and observations has provided considerable insights into physical processes for galaxy formation and evolution and has played a crucial role in constraining theoretical models. However, simulations and observations have not been fully reconciled. Contributing factors include observational uncertainties, modeling uncertainties in the simulations, and ad hoc comparisons. For instance, cosmological simulations—specifically, subgrid models such as stellar winds and black hole feedback—have generally been calibrated against only a handful of observables through by-eye comparisons between simulations and observations, along with educated guesses or simple parameter-space search algorithms (Schaye et al. 2015; Pillepich et al. 2018a; Oh et al. 2020). The limits of this conventional calibration approach are as follows: (1) The dimensions of the subgrid parameter space that one can cover are significantly limited when using by-eye comparisons. In addition, as the number of parameters of interest increases, it becomes harder to provide educated guesses due to complex and intertwined relations between physical models and observables. (2) It is challenging to calibrate against numerous observables simultaneously. (3) The accuracy of the calibrated parameters is hard to determine due to the objective nature of the comparison process. (4) The simulation uncertainty, due to such sources as cosmic variance, is generally not taken into consideration. Cosmological simulations suffer from uncertainty that comes from various sources of randomness such as initial conditions (Genel et al. 2014, 2019; Keller et al. 2019). This can lead to appreciable bias in the calibration.

Similarly, semianalytical models (SAMs), which estimate properties of the galaxy population using parameterized physical models that include a number of free parameters, have rigorously tuned those parameters to reproduce certain observational properties of the galaxy population using Bayesian inference together with the Markov Chain Monte Carlo method (MCMC) (Lu et al. 2012, 2014; Benson 2014). Bayesian inference is a widely used method of statistical inference that updates one’s knowledge—or belief in the Bayesian sense—of the parameters by making new observations. The majority of the problems of the conventional calibration method for cosmological simulations can be alleviated by the use of Bayesian inference. For instance, Bayesian inference is conducted through the likelihood function, which can mathematically guarantee the precision of calibration and enables inference from numerous observations simultaneously. Also, the probabilistic nature of Bayesian inference captures the uncertainty of parameters through Bayes’s theorem. However, Bayesian inference usually entails MCMC, which is computationally expensive, for determining the posterior distribution. Moreover, MCMC sometimes fails to retrieve a posterior distribution, usually when the target posterior is complex and high-dimensional. In contrast to SAMs, hydrodynamical cosmological simulations with subgrid models are too computationally costly to be used to perform hundreds of thousands of sequential simulations for MCMC, rendering such an approach impossible in practice. This is the primary reason why, despite all the merits, calibration of cosmological simulations in the Bayesian framework has not been conducted thus far. In addition, conventional Bayesian inference

has its limitations in terms of the need for an explicit (analytic) likelihood. Since the likelihood should be explicitly formulated, commonly used analytic likelihoods such as Gaussians are only an approximation to the true unknown one.

The implicit likelihood inference approach (ILI)—also known as likelihood-free inference or simulation-based inference—provides a framework for performing rigorous Bayesian inference in a computationally efficient way, especially for inferences on computationally expensive simulations (Mohamed & Lakshminarayanan 2016; Alsing et al. 2018; Papamakarios et al. 2018; Cranmer et al. 2020; Durkan et al. 2020). In contrast to conventional Bayesian inference, which requires an explicit (analytic) formulation for the likelihood, ILI learns the likelihood—the conditional distribution of observables given the parameters—directly from simulated parameter–observable pairs using a neural density estimator (NDE). NDEs are a flexible representation of the likelihood placing only mild assumptions on the likelihood form.

On the computational cost side, the likelihood can be evaluated through the trained NDE without performing further simulations. The number of emulations required for the inference is equivalent to the number of emulations for training the NDEs, which is generally thousands of simulations (Alsing et al. 2019). In conventional inference, although it generally depends on the complexity of the problem, the convergent MCMC typically requires at least 10^5 samples in cosmological applications (Feroz & Hobson 2008; Trotta et al. 2011), which entails significantly more simulations than the NDE requires. ILI has already been vigorously exploited for inference and estimation of physical quantities in astrophysics, for example for inference of the Hubble constant from binary neutron star mergers (Gerardi et al. 2021), constraints on the cosmological parameters from weak lensing (Tam et al. 2022), mass estimations of the Milky Way and M31 (Lemos et al. 2021; Villanueva-Domingo et al. 2021), inference of strong gravitational lensing parameters (Legin et al. 2021), dynamical mass estimation of galaxy clusters (Kodi Ramanah et al. 2021), and inference of reionization parameters from the 21 cm power spectrum and light cones (Zhao et al. 2022a, 2022b).

In this work, we adopt the `sbi` package (Tejero-Cantero et al. 2020), the successor of the `delfi` package, which is equipped with various NDEs for ILI, to calibrate cosmological simulations against observations. We also exploit the suite of cosmological simulations of the Cosmology and Astrophysics with Machine Learning Simulations (CAMELS) project (Villaescusa-Navarro et al. 2021), which includes the largest data set designed to train machine-learning models and provides more than a thousand simulations for exploring the cosmological and astrophysical parameter space. Despite the large number of simulations in the CAMELS project, it is still not enough for directly employing our ILI technique, so we build an emulator that is trained on the CAMELS simulations to estimate the target observable taking the cosmological and astrophysical parameters as input. This provides much flexibility and reduces computational cost during inference, which is in line with Elliott et al. (2021), who used an emulator to calibrate a SAM. Using the emulator as a surrogate for the actual cosmological simulations, we perform ILI using the observed cosmic star formation history (Leja et al. 2022) and the observed SMFs (Leja et al. 2020) to infer the parameters that are varied in the CAMELS suite—two cosmological parameters (Ω_m and σ_8) and four astrophysical parameters

(related to stellar wind feedback and kinetic black hole feedback) (Figure 1).

The structure of this paper is as follows. In Section 2, we describe the CAMELS simulations that we use to train the emulators and the details of the target observables. In Section 3, we give a brief review of the ILI method including the NDE and our emulator design. In Section 4.1, we investigate the performance and convergence of the posterior distributions of the cosmological and astrophysical parameters inferred from an emulated star formation rate density (SFRD) as the target observable. In Section 4.2, we perform the inference from the observed SFRD and study how the inferred SFRD matches the observed one. In Section 5.1, we investigate the performance and convergence of the posterior distributions from an emulated SMF as the target observable and in Section 5.2, we perform the inference from observed SMFs and study the discrepancies between those and the inferred ones. In Section 6, we discuss the properties of the inferred posteriors and any mismatch between inferences and observations with respect to the correlation between the parameter–observable pairs and physical analysis of the cosmological simulations. In Section 7, we present a summary of the results and findings.

2. Cosmological Simulations: The CAMELS Project

2.1. Overview

CAMELS¹⁰ is a suite of 4233 cosmological simulations: 2184 (magneto)hydrodynamic simulations with the AREPO and GIZMO codes, and 2049 *N*-body simulations (Villaescusa-Navarro et al. 2021). Each simulation contains 256^3 dark matter particles of mass $6.49 \times 10^7 (\Omega_m - \Omega_b)/0.251 h^{-1} M_\odot$ and 256^3 gas cells with an initial mass of $1.27 \times 10^7 h^{-1} M_\odot$ in a periodic box of a comoving volume of $(25 h^{-1} \text{ Mpc})^3$, which results in a resolution comparable to but slightly lower than that of the TNG300 simulation of the ILLUSTRISTNG project (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018a; Springel et al. 2018). The CAMELS project has been exploring a wide cosmological and astrophysical parameter space for applications of machine learning in astrophysics. The cosmological and astrophysical parameters of interest are Ω_m , σ_8 , $A_{\text{SN}1}$, $A_{\text{SN}2}$, $A_{\text{AGN}1}$, and $A_{\text{AGN}2}$ (refer to Equations (1)–(7) for details). The suite of (magneto)hydrodynamic CAMELS simulations comprises three different sets for each of the AREPO and GIZMO codes, as follows: (1) the LH set consists of 1000 simulations with different initial conditions varying all parameters sampled from a Latin hypercube, (2) the 1P set consists of 61 simulations with the same initial conditions varying only one parameter at a time, and (3) the CV set consists of 27 simulations with fixed cosmology and astrophysics that sample cosmic variance using different initial conditions.

The simulations run with the AREPO and GIZMO codes use the TNG and SIMBA models, respectively, and we refer to these suites as the TNG and SIMBA runs. Throughout this work, we adopt the TNG suites of the CAMELS simulations unless specified otherwise. The LH set of the TNG suites is exploited to train the emulator (Section 3.3). The CV set of the TNG suites is used to model simulation uncertainty (Appendix B). The 1000 simulations of the LH set are run with $\Omega_m \in [0.1, 0.5]$, $\sigma_8 \in [0.6, 1.0]$, $A_{\text{SN}1} \in [0.25, 4.0]$, $A_{\text{SN}2} \in [0.5, 2.0]$,

$A_{\text{AGN}1} \in [0.25, 4.0]$, and $A_{\text{AGN}2} \in [0.5, 2.0]$ arranged in a Latin hypercube. The 27 simulations of the CV set are run with $\Omega_m = 0.3$, $\sigma_8 = 0.8$, and $A_{\text{SN}1} = A_{\text{SN}2} = A_{\text{AGN}1} = A_{\text{AGN}2} = 1$ but with different initial conditions. Meanwhile, the following cosmological parameters are fixed across all simulations: $\Omega_b = 0.049$, $h = 0.6711$, $n_s = 0.9624$, $M_\nu = 0.0 \text{ eV}$, $w = -1$, and $\Omega_K = 0$. The TNG suite of the CAMELS simulations implements the subgrid physics models of ILLUSTRISTNG (Weinberger et al. 2017; Pillepich et al. 2018a). These simulations employ the AREPO code¹¹ (Springel 2010; Weinberger et al. 2020) to solve gravity (TreePM) and magnetohydrodynamics using a Voronoi moving-mesh approach. The ILLUSTRISTNG physics includes various subgrid models: radiative cooling and heating, star formation, stellar evolution, feedback from galactic winds, the formation and growth of supermassive black holes (SMBHs), and feedback from AGN.

2.2. Physics of Astrophysical Parameters

We have four astrophysical parameters that control the strengths of star formation–driven galactic winds and SMBH feedback. The star formation–driven galactic wind is isotropically injected in a kinetic form (see Pillepich et al. 2018b for details). The winds are characterized by a mass loading factor η_w , which is defined by

$$\eta_w \equiv \frac{\dot{M}_w}{\dot{M}_{\text{SFR}}} = \frac{2}{v_w^2} e_w (1 - \tau_w), \quad (1)$$

where \dot{M}_w and \dot{M}_{SFR} are the rate of gas mass to be converted into wind particles and the instantaneous, local star formation rate (SFR), respectively. With a fixed thermal energy fraction τ_w , the mass loading factor is determined by the total energy injection rate per unit star formation e_w and the wind speed v_w , which involve $A_{\text{SN}1}$ and $A_{\text{SN}2}$, respectively, as follows:

$$e_w = A_{\text{SN}1} \times \bar{e}_w \left[f_{w,Z} + \frac{1 - f_{w,Z}}{1 + (Z/Z_{w,\text{ref}})^{\gamma_{w,Z}}} \right] \times N_{\text{SN}1} E_{\text{SN}1,51} 10^{51} \text{ erg } M_\odot^{-1} \quad (2)$$

and

$$v_w = A_{\text{SN}2} \times \max \left[\kappa_w \sigma_{\text{DM}} \left(\frac{H_0}{H(z)} \right)^{1/3}, v_{w,\text{min}} \right]. \quad (3)$$

Details on the parameters \bar{e}_w , $f_{w,Z}$, $Z_{w,\text{ref}}$, Z , $\gamma_{w,Z}$, $N_{\text{SN}1}$, $E_{\text{SN}1,51}$, κ_w , σ_8 , and $v_{w,\text{min}}$ can be found in Pillepich et al. (2018a, Table 1).

An SMBH particle with mass $M_{\text{seed}} = 8 \times 10^5 h^{-1} M_\odot$ is seeded on the fly at the center of any halo with mass $M_{\text{FoF}} > 5 \times 10^{10} h^{-1} M_\odot$ that does not yet contain an SMBH. To prevent the SMBH particles from artificially wandering around the galaxy, they are kept close to the potential minimum of their host dark matter halos using an ad hoc prescription.

ILLUSTRISTNG adopts Bondi–Hoyle–Lyttleton accretion (Hoyle & Lyttleton 1939; Bondi & Hoyle 1944; Bondi 1952) with the Eddington cap for the growth of SMBHs. The states of the SMBHs are distinguished into high accretion (a classical thin disk) and low accretion (hot accretion flow) based on a

¹⁰ <https://www.camel-simulations.org>

¹¹ <https://arepo-code.org>

threshold of

$$\chi = \min \left[0.002 \left(\frac{M_{\text{BH}}}{10^8 M_{\odot}} \right)^2, 0.1 \right] \quad (4)$$

in units of the Eddington accretion limit. According to the state of accretion, the feedback mode is determined. For the high-accretion state, the feedback energy is injected as pure thermal energy into the vicinity of the SMBH (thermal mode). For the low-accretion state, feedback energy is released kinetically in a random direction (kinetic mode) as

$$\dot{E}_{\text{low}} = A_{\text{AGN}1} \epsilon_{\text{f,kin}} \dot{M}_{\text{BH}} c^2, \quad (5)$$

where

$$\epsilon_{\text{f,kin}} = \min \left[\frac{\rho}{0.05 \rho_{\text{SFthresh}}}, 0.2 \right]. \quad (6)$$

Here, ρ and ρ_{SFthresh} are the gas density around the SMBH and the density threshold for star formation. The injection of kinetic feedback occurs every time the accumulated energy reaches the energy threshold from the last feedback. The energy threshold for the kinetic feedback is parameterized as

$$E_{\text{inj,min}} = A_{\text{AGN}2} \times f_{\text{re}} \frac{1}{2} \sigma_{\text{DM}}^2 m_{\text{enc}}, \quad (7)$$

where σ_{DM}^2 is the one-dimensional dark matter velocity dispersion around the SMBH, m_{enc} is the enclosed gas mass within the feedback sphere, and f_{re} is a free parameter that is set to 20 for the fiducial TNG model. $A_{\text{AGN}2}$ controls the frequency and speed of the SMBH feedback. The details of the prescription for the SMBH physics in ILLUSTRISTNG are described in Weinberger et al. (2017).

2.3. Target Observables: Cosmic SFRD and SMF

In this work, the observables from which the cosmological and astrophysical parameters are inferred are the cosmic SFRD and the SMFs. For a fair comparison between observations and simulations, we take into consideration the consistency between SFRD and SMFs. For instance, cosmological simulations, by nature, can guarantee consistency between the evolution of SFR and that of stellar mass. That is, the SFRD $f(z) = dM_{\star}(z)/dz + \dot{M}_{\text{return}}$, where $M_{\star}(z) = \int_M \phi(M, z)$, $\phi(M, z)$ is the galaxy SMF, and \dot{M}_{return} is the rate of mass return from evolving stellar populations. On the observational side, the consistency depends on, e.g., the modeling for each observable, and it is not guaranteed that the SFRD and the SMFs at different redshifts are consistent with each other (for more details, refer to Section 5 of Leja et al. 2019).

To circumvent this, we adopt the SFRD of Leja et al. (2022) ranging from $z = 3$ to 0.5 and the five SMFs of Leja et al. (2020) at $z = 0.5, 1, 1.5, 2.0$, and 2.5, both of which are inferred with Prospector- α (Leja et al. 2019) using galaxies in the 3D-HST (Skelton et al. 2014) and COSMOS-2015 (Laigle et al. 2016) catalogs in such a way that the SFRD and SMFs are consistent with each other (Leja et al. 2019). The galaxies are selected above the stellar mass completeness limit taken from Tal et al. (2014) for 3D-HST and that from Laigle et al. (2016) for COSMOS-2015. In the case of the SFR estimation, the mass completeness limits are adjusted upward by $\mathcal{O}(0.1)$ dex since red galaxies, which can have an impact on

SFR, are more likely to be excluded. However, Leja et al. (2022) found that the resultant SFRD is not sensitive to these adjustments. In addition, neither SFRD nor SMF has error bars since the statistical uncertainties are negligible, and the true uncertainties are systematic in nature.

On the simulation side, the SFRD is constructed from the global SFR per unit comoving volume for 21 snapshots (21 redshifts), matching the redshifts between simulations and observations. The SMFs are obtained from the stellar mass of the galaxy catalog, binned into 13 bins in the range $[10^{8.9}, 10^{11.4}] M_{\odot}$, which Leja et al. (2020) aimed at.

2.4. Resolution Effects: Convergence and Rescaling

The resolution convergence and effects in the TNG simulations have been extensively studied in Weinberger et al. (2017, Appendix B), Pillepich et al. (2018a, Appendix A), and Pillepich et al. (2018b, Section 3.3). In general, observables such as the SMF at different resolutions are not converged. Figure 18 shows that there is a shift of the order of tens of percent in both SMFs and SFRDs between TNG100-1 and TNG100-2. Note that the fiducial parameters are calibrated against observations for the TNG100 simulation of the ILLUSTRISTNG project, and the CAMELS parameter variations of the LH set are chosen around the fiducial values of the TNG100 simulation. Hence, based on the approach of Pillepich et al. (2018b, Appendix A), we rescale the SFRD and SMFs with a mass-modulated rescaling factor. Since the resolution of CAMELS simulations is comparable to that of TNG100-2, we construct a rescaling factor using the SFR–halo mass relation and the stellar mass–halo mass relation from TNG100-1 and TNG100-2 (details in Appendix A). The rescaled CAMELS simulations in Figure 18 demonstrate that this procedure reduces the resolution effects at some level. One remark is that we find that the rescaling depends on the cosmological and astrophysical parameters. However, in this work, we ignore this effect. Further details will be discussed in Section 6.5.1.

2.5. Uncertainties in Simulation: Cosmic Variance and Butterfly Effects

In this section, we focus on *simulation uncertainty* that is modeled as mock uncertainty and added to emulators in Section 3.3. The *simulation uncertainty* is the intrinsic uncertainty of cosmological simulations. In cosmological simulations, randomness in the initial conditions that corresponds to the density fluctuations of the early universe leads to *cosmic variance*. The minute position differences of the initial conditions owing to random seeds manifest as differences in the large-scale structure that directly impact the galaxy populations. On the observational side, *cosmic variance* can be attributed to the limited volume of the surveys. Meanwhile, the *butterfly effect* stems from the chaotic behaviors of cosmological simulations. In dynamical systems, we can quantify chaotic or stochastic behaviors using a quantity called the Lyapunov exponent. If the Lyapunov exponent is positive, the minute perturbations evolve exponentially with the Lyapunov timescale and manifest as macroscopic differences. The chaotic behavior of the galactic dynamical systems amplifies minute fluctuations, seeded by randomness in numerical computations such as stochasticity in subgrid models and floating errors, into appreciable differences in later times (Genel et al. 2019). Also, Keller et al. (2019) studied the

stochasticity of galaxy properties using the particle-based code GASOLINE and the grid-based code RAMSES. For the details of the quantification of simulation uncertainty, refer to Appendix B.

3. ILI with Emulator

ILI aims at identifying the regions of parameter space of forward models to match observations (Cranmer et al. 2020; Tejero-Cantero et al. 2020). In Bayesian terminology, ILI retrieves the posterior distribution over the parameters given an observation. In contrast to conventional Bayesian inference, ILI requires no assumption or ansatz for likelihood $p(\mathbf{x}|\boldsymbol{\theta})$ so that it can also be applicable when analytical formulation for the likelihood is not accessible. Here, \mathbf{x} and $\boldsymbol{\theta}$ generally stand for observable and relevant parameters, respectively.

In this work, we adopt `sbi`¹², a package designed to retrieve the posterior distribution $p(\boldsymbol{\theta}|\mathbf{x})$ (Tejero-Cantero et al. 2020). Instead of analytical probability distributions such as the Gaussian distribution, `sbi` employs a neural network to output a probability distribution called an NDE. An NDE is a neural network that takes data points \mathbf{x} as input and returns a conditional probability distribution $\tilde{p}_\phi(\boldsymbol{\theta}|\mathbf{x})$ over the simulation parameters such that $\int \tilde{p}_\phi(\boldsymbol{\theta}|\mathbf{x}) d\boldsymbol{\theta} = 1$, where ϕ represents the neural network parameters.

The `sbi` package also aims at training the NDE with the lowest number of evaluations of forward models such as cosmological simulations, which are in general computationally expensive, by concentrating on the relatively small parameter space around parameter values of interest. To this end, we must choose the proposal distribution $q(\boldsymbol{\theta})$ from which we draw the sets of parameters for the new simulations in the next iteration. In general, the choice of the optimal proposal for a specific problem is an open question. Since we are interested in the highly probable regions of parameter space, it might be natural to adopt the current approximate posterior density as a proposal density (Papamakarios et al. 2018; Durkan et al. 2020). On the other hand, Alsing et al. (2018) adopted the geometric mean of the prior and the current approximate posterior density in the context of sequential approximate Bayesian computation. It might increase the probability of exploring the parameter space more broadly beyond the posterior density.

Lastly, the `sbi` package provides three different ways of computing an NDE: (1) sequential neural posterior estimation (SNPE) trains an NDE to directly compute the posterior distribution; (2) sequential neural likelihood estimation (SNLE) trains an NDE to estimate the likelihood, which allows one to evaluate the posterior distribution with the prior; and (3) sequential ratio estimation (SRE) trains a binary classifier as an NDE to estimate density ratios, which can be used to estimate the posterior distribution (details will be discussed in Section 3.2). Of the three, we adopt the ratio estimation method because it generally requires a smaller amount of computational resources for training an NDE than the others do, without losing much accuracy (Durkan et al. 2020). We have compared the performance of the three different methods in our setup based on posteriors inferred from one of our observables using all three methods, but any noticeable differences cannot be found.

¹² <https://github.com/mackelab/sbi>

3.1. NDE: Learning the Density Ratio

The NDE estimates the posterior density $p(\boldsymbol{\theta}|\mathbf{x})$ by computing the density ratio $r(\boldsymbol{\theta}, \mathbf{x}) = \frac{p(\mathbf{x}|\boldsymbol{\theta})}{p(\mathbf{x})} = \frac{p(\boldsymbol{\theta}|\mathbf{x})}{p(\boldsymbol{\theta})}$, where a simulator defines a valid probability density function $p(\mathbf{x}|\boldsymbol{\theta})$ over observations \mathbf{x} (Sugiyama et al. 2012; Mohamed & Lakshminarayanan 2016; Hermans et al. 2019).

Consider a binary random variable $Y: \Omega \rightarrow \mathbb{R}$ where $\Omega \in \{y = 0, y = 1\}$ and each outcome is equally likely a priori (i.e., $p(y = 0) = p(y = 1)$). Then,

$$\begin{aligned} p(y = 1|\mathbf{x}, \boldsymbol{\theta}) &= \frac{p(\mathbf{x}, \boldsymbol{\theta}|y = 1)}{p(\mathbf{x}, \boldsymbol{\theta}|y = 1) + p(\mathbf{x}, \boldsymbol{\theta}|y = 0)} \\ &= \frac{r(\mathbf{x}, \boldsymbol{\theta})}{1 + r(\mathbf{x}, \boldsymbol{\theta})}, \end{aligned} \quad (8)$$

where the density ratio $r(\mathbf{x}, \boldsymbol{\theta})$ is defined by $r(\mathbf{x}, \boldsymbol{\theta}) \equiv \frac{p(\mathbf{x}, \boldsymbol{\theta}|y = 1)}{p(\mathbf{x}, \boldsymbol{\theta}|y = 0)}$ (refer to the derivation¹³). That is, the binary classifier that computes $p(y = 1|\mathbf{x}, \boldsymbol{\theta})$ or $p(y = 0|\mathbf{x}, \boldsymbol{\theta})$ can be exploited to compute the ratio density $r(\mathbf{x}, \boldsymbol{\theta})$.

In the case of the conditional probability density, the $y = 1$ class represents $(\mathbf{x}, \boldsymbol{\theta}) \sim p(\mathbf{x}, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is drawn from the given prior $p(\boldsymbol{\theta})$ and \mathbf{x} is obtained subsequently by the simulations with the sampled parameters. The $y = 0$ class consists of pairs $(\mathbf{x}, \boldsymbol{\theta}) \sim p(\mathbf{x})p(\boldsymbol{\theta})$ with parameters and observations sampled independently. By training the neural classifier $\tilde{p}_\phi(\mathbf{x}, \boldsymbol{\theta})$ to take $(\mathbf{x}, \boldsymbol{\theta})$ as input and output the probability $p(y = 1|\mathbf{x}, \boldsymbol{\theta})$ or $p(y = 0|\mathbf{x}, \boldsymbol{\theta})$, where ϕ stands for the hyperparameters (e.g., weights and biases) of the neural network, we can compute the density ratio

$$r(\mathbf{x}, \boldsymbol{\theta}) = \frac{p(\mathbf{x}, \boldsymbol{\theta})}{p(\mathbf{x})p(\boldsymbol{\theta})} = \frac{p(\mathbf{x}|\boldsymbol{\theta})}{p(\mathbf{x})} = \frac{p(\boldsymbol{\theta}|\mathbf{x})}{p(\boldsymbol{\theta})}. \quad (9)$$

Therefore, if the prior $p(\boldsymbol{\theta})$ is known and can be evaluated, the posterior density $p(\boldsymbol{\theta}|\mathbf{x})$ can be obtained as $p(\boldsymbol{\theta}|\mathbf{x}) = r(\mathbf{x}, \boldsymbol{\theta})p(\boldsymbol{\theta})$.

3.2. SRE

We employ and modify the SRE (Durkan et al. 2020) implemented in `sbi`. The SRE retrieves a posterior density using the NDE described in Section 3.1. The workflow of the SRE is as follows (see also Algorithm 1): (1) The SRE samples M sets of parameters $\boldsymbol{\theta}_m$ from a prior $p(\boldsymbol{\theta})$ or a proposal density $q^{(n)}(\boldsymbol{\theta}|\mathbf{x}_0)$ and generates M observables \mathbf{x}_m with the parameters $\boldsymbol{\theta}_m$, where $m = 1, \dots, M$. (2) The SRE constructs a training set \mathcal{T} that consists of two classes of pairs: $(\mathbf{x}_m, \boldsymbol{\theta}_m)$ for $y = 1$ and $(\mathbf{x}_k, \boldsymbol{\theta}_m)$ for $y = 0$ by permutations where $k \neq m$, where y simply represents a binary class (see Section 3.1). (3) The NDE $\tilde{p}_\phi(\mathbf{x}, \boldsymbol{\theta})$ is trained on \mathcal{T} until it converges. (4) The structure of the trained NDE is saved as an external file for future use. (5) The posterior and proposal densities are updated as $p^{(n)}(\boldsymbol{\theta}|\mathbf{x}_0) \propto \tilde{p}_\phi(\mathbf{x}_0|\boldsymbol{\theta})p(\boldsymbol{\theta})$ and $q^{(n)}(\boldsymbol{\theta}|\mathbf{x}_0) \propto p^{(n)}(\boldsymbol{\theta}|\mathbf{x}_0)p(\boldsymbol{\theta})$, respectively. (6) The SRE repeats from (1) with the newly updated proposal. During the SRE, we use `emcee`¹⁴, an MCMC package, for sampling parameters from the proposal (Foreman-Mackey et al. 2013). However, for the final production plots in this paper, parameters are drawn from the

¹³ By definition of conditional probability, $p(y = 1|\mathbf{x}) = p(\mathbf{x}|y = 1)p(y = 1)/p(\mathbf{x})$. Using the law of total probability, we can rewrite it as $p(y = 1|\mathbf{x}) = p(\mathbf{x}|y = 1)p(y = 1)/(p(\mathbf{x}|y = 0)p(y = 0) + p(\mathbf{x}|y = 1)p(y = 1))$. Since each class is equally likely a priori, $p(y = 1|\mathbf{x}) = p(\mathbf{x}|y = 1)/(p(\mathbf{x}|y = 0) + p(\mathbf{x}|y = 1))$.

¹⁴ <https://github.com/dfm/emcee>

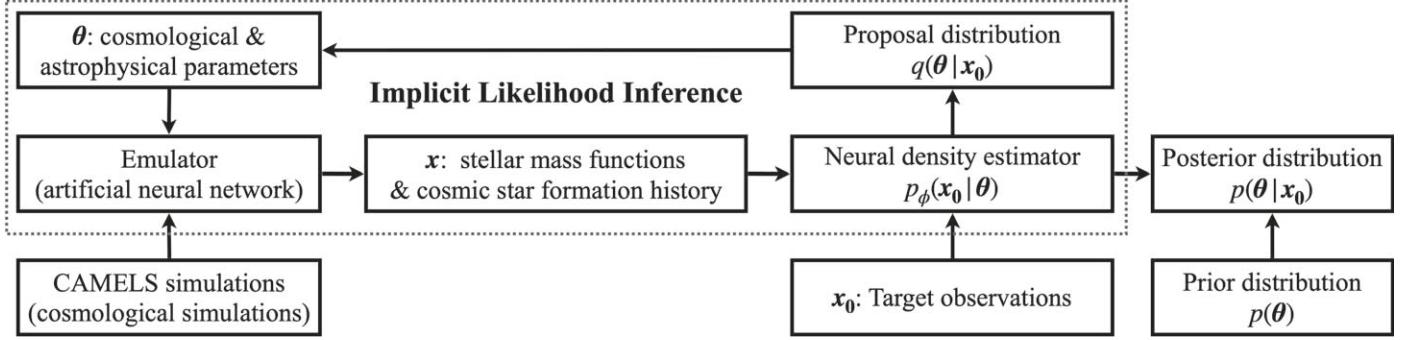


Figure 1. Diagram for the pipeline of this work, which aims at calibrating cosmological simulations against observations. We use the CAMELS simulations (Section 2) to train emulators (Section 3.3) that estimate SMFs and cosmic star formation history, \mathbf{x} , taking cosmological and astrophysical parameters, $\boldsymbol{\theta}$, as input. Using the emulators as a surrogate for cosmological simulations, we perform ILI (Section 3) from observations, \mathbf{x}_0 (Section 2.3), to retrieve the posterior distributions, $p(\boldsymbol{\theta}|\mathbf{x}_0)$, of the cosmological and astrophysical parameters given the observations.

externally saved NDE using `zeus` (Karamanis & Beutler 2020; Karamanis et al. 2021). The choice of an optimal proposal density is a crucial element for the SRE. However, this remains an open question. The posterior inferred in the previous round is the most common for the proposal density for the next round (Papamakarios et al. 2018; Durkan et al. 2020). In this work, we adopt the geometric mean of the prior and the posterior density of the current epoch, inspired by Alsing et al. (2018).

The SRE has two hyperparameters subject to optimization: the number of simulations per iteration M and the number of iterations N . Since M determines the amount of information that the training set can carry, it can strongly affect the accuracy of the posterior density. For instance, if M is biased, the resultant posterior density can also be biased. M should be set to a sufficiently large value that the sampled training set can carry enough information about the trained posterior. We set M empirically via many trials as follows: new 100 or 200 SFRDs that are generated by the emulator, with parameters sampled from the proposal density, are added to the training data set. N is directly related to the convergence and stability of the inference. We perform the SRE without imposing a definite N and stop it whenever there is no improvement in training for 20 epochs—e.g., it has converged.

On the other hand, the structure of the NDE can be more decisive and has critical hyperparameters in terms of accuracy of the inferred posterior density. The residual network (ResNet) is adopted as the baseline neural network for the NDE. We determine the complexity of the NDE depending on correlations between the parameters and observables. Since the SMF–parameter pairs have more strongly correlated behavior than the SFRD–parameter pairs (Section 6.1), we adopt a deeper network for the SFRD than for the SMFs as follows: the NDE, $\tilde{p}_\phi(\mathbf{x}, \boldsymbol{\theta})$, that we adopt for the SFRD (SMF) is a feedforward residual network that consists of two residual blocks of 250 (100) hidden units (for the detailed structure of the NDE, refer to Appendix C).

Algorithm 1. Sequential Ratio Estimation

Input: Simulator $\tilde{p}(\mathbf{x}|\boldsymbol{\theta})$, prior $p(\boldsymbol{\theta})$, target observation \mathbf{x}_0 , neural density estimator $\tilde{p}_\phi(\mathbf{x}, \boldsymbol{\theta})$, iterations N , simulations per iteration M
Initialize: Proposal $q^{(1)}(\boldsymbol{\theta}|\mathbf{x}_0) = p(\boldsymbol{\theta})$, training set $\mathcal{T} = \{\}$
for $n = 1$ to N **do**
 Draw $\boldsymbol{\theta}_m \sim q^{(n)}(\boldsymbol{\theta}|\mathbf{x}_0)$, $m = 1, \dots, M$;
 Simulate $\mathbf{x}_m \sim \tilde{p}(\mathbf{x}|\boldsymbol{\theta}_m)$, $m = 1, \dots, M$;

(Continued)

```

Construct  $\mathcal{T} = \mathcal{T} \cup \{(\mathbf{x}_m, \boldsymbol{\theta}_m)\} | m = 1, \dots, M\}$ ;  

while  $\tilde{p}_\phi$  not converged do  

  Sample mini-batch  $\{(\mathbf{x}_b, \boldsymbol{\theta}_b)\} \sim \mathcal{T}$ ;  

  Optimize a neural density estimator  $\tilde{p}_\phi(\mathbf{x}, \boldsymbol{\theta})$  using stochastic gradient descent;  

end  

  Save the neural density estimator  $\tilde{p}_\phi(\mathbf{x}, \boldsymbol{\theta})$   

  Update posterior  $p^{(n)}(\boldsymbol{\theta}|\mathbf{x}_0) \propto \tilde{p}_\phi(\mathbf{x}_0|\boldsymbol{\theta})p(\boldsymbol{\theta})$   

  Update proposal  $q^{(n)}(\boldsymbol{\theta}|\mathbf{x}_0) \propto \sqrt{p^{(n)}(\boldsymbol{\theta}|\mathbf{x}_0)p(\boldsymbol{\theta})}$   

end  

Output: Posterior  $p^{(N)}(\boldsymbol{\theta}|\mathbf{x}_0) \propto \tilde{p}_\phi(\mathbf{x}_0|\boldsymbol{\theta})p(\boldsymbol{\theta})$ 

```

3.3. Emulator: Surrogate for Cosmological Hydrodynamic Simulations

The number of simulations required to retrieve the posterior density is highly correlated with the dimensions and complexity of a problem. ILI generally requires more than $\mathcal{O}(1000)$ simulations (Hermans et al. 2019; Durkan et al. 2020; Dalmasso et al. 2021; Huppenkothen & Bachetti 2022; Figures 3 and 9 in this work), which exceeds the total number of CAMELS simulations. We thus circumvent this issue by adopting an emulator as a surrogate simulation. The emulator is constructed upon a fully connected neural network that is faster than hydrodynamic simulations by several orders of magnitude. We split the LH set of the CAMELS simulations into *training* (750), *test* (150), and *validation* (150) sets. Six independent neural *emulators* are trained on the training sets to estimate the SFRD and the five SMFs at five different redshifts $z = 0.5, 1.0, 1.5, 2.0, 2.5$ as a function of six cosmological and astrophysical parameters: $(\Omega_m, \sigma_8, A_{\text{SN1}}, A_{\text{SN2}}, A_{\text{AGN1}}, A_{\text{AGN2}})$ (further details of the six emulators are discussed in Appendix C). We use Optuna (Akiba et al. 2019), an automatic hyperparameter optimization tool, to train and optimize the emulators. The hyperparameters subject to optimization include the learning rate, weight decay, the number of layers, and the number of neurons. During training, both inputs (cosmological and astrophysical parameters) and outputs (the SFRD and the SMFs) are normalized, using linear scaling¹⁵ and the z -score¹⁶, respectively. We measure the accuracy of the emulators with

¹⁵ $x' = (x - x_{\min}) / (x_{\max} - x_{\min})$, where x' is the normalized input, and x_{\min} and x_{\max} are the minimum and maximum of the inputs, namely the edge values of the parameter ranges (refer to Section 2).

¹⁶ $x' = (x - \mu) / \sigma$, where x' is the normalized output and μ and σ are the mean and the standard deviation of the outputs, respectively.

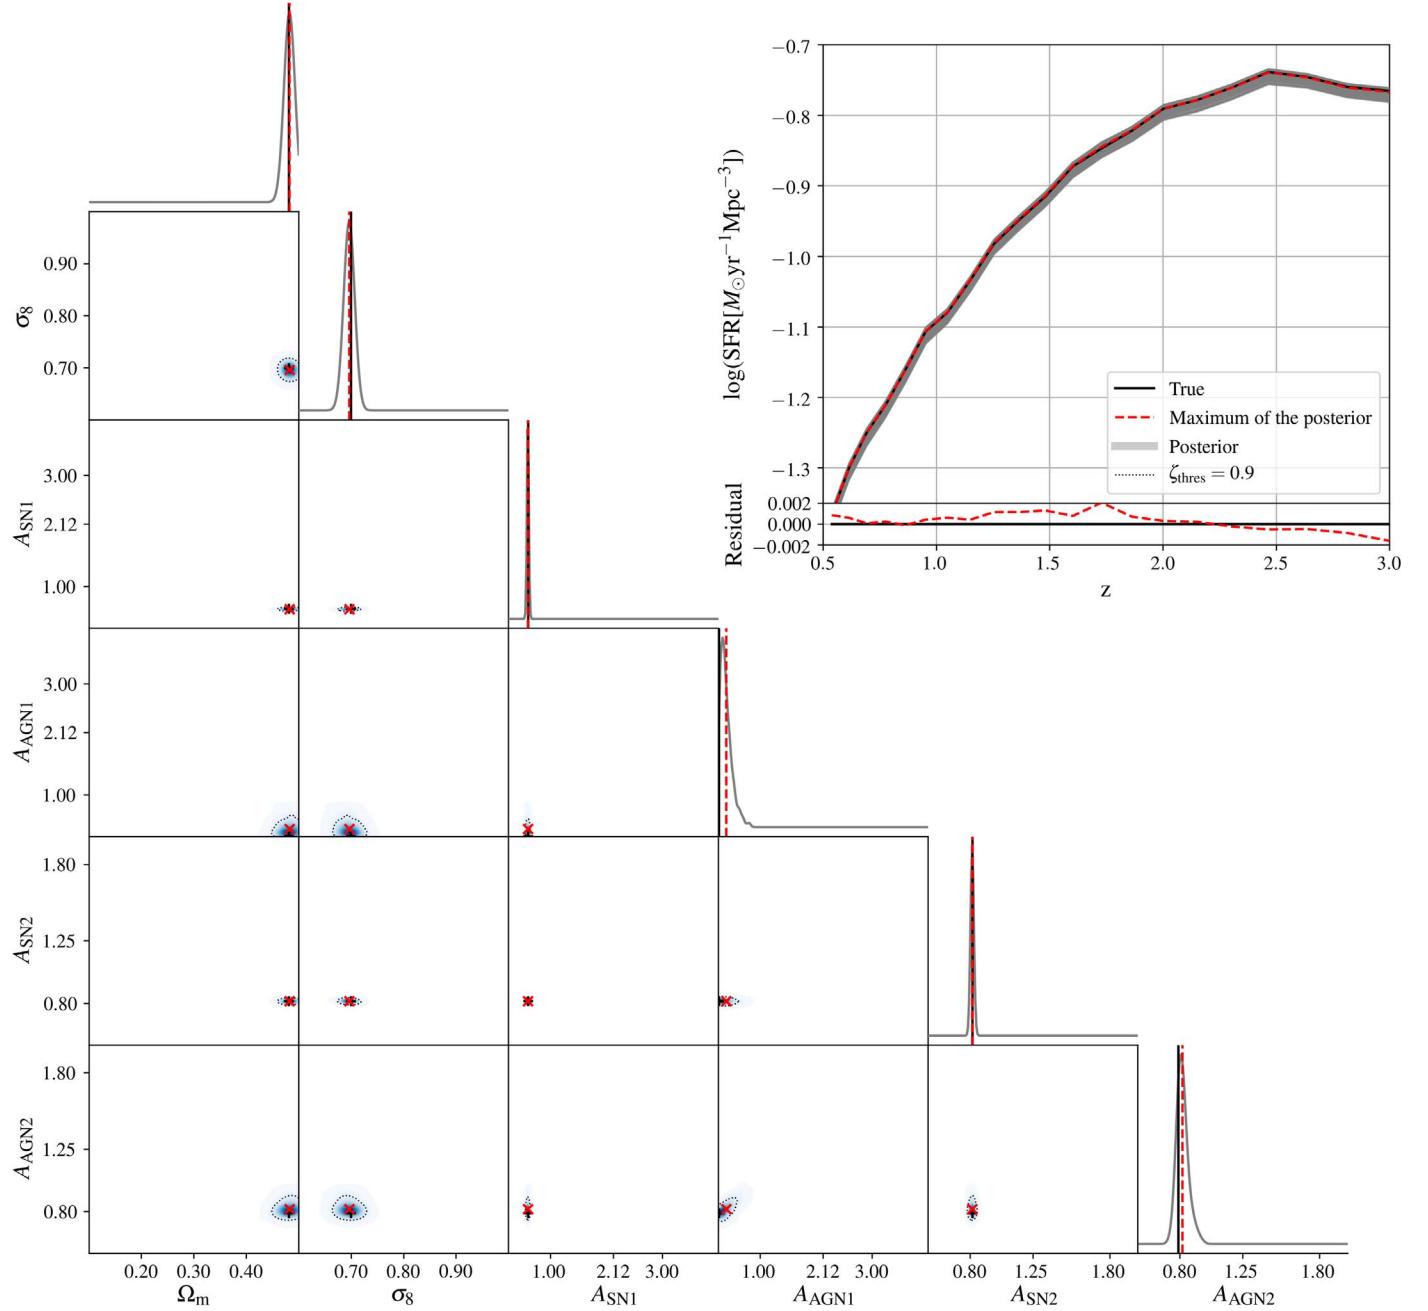


Figure 2. Top right: Emulated cosmic SFRDs from the inference without the mock uncertainty. Shown are the inferred posterior (gray), the maximum of the posterior (red dashed), and the emulator-based target SFRD (black solid). The emulator-based target SFRD is generated by the emulator taking one of the parameter combinations from the LH set (details in Section 4.1.1). Bottom left: Two-dimensional distribution of the inferred posterior. The black and red crosshairs represent the values of the target and maximum of the posterior, respectively. The marginal distributions are obtained by kernel density estimation. The black solid and red dashed vertical lines indicate the true values and maximum values of the inferred posteriors, respectively. The inferred parameters (red dashed) and their true values (black solid) are nearly on top of each other with small errors and variances (see Table 1 for details). The inferred SFRD and the target SFRD are also in good agreement with a relative error of 0.17% (see Section 4.1.1).

the mean square error (MSE) and the Pearson correlation coefficient. The MSEs for the emulators are 0.0007 dex (SFRD) and 0.0011 dex (SMF). The Pearson correlation coefficients are 0.98 (SFRD) and 0.94 (SMF).

3.3.1. Connection between Emulator and Cosmological Hydrodynamic Simulation

The emulator can be the best option for a surrogate simulation with reasonable accuracy in terms of computational cost. However, uncertainties of cosmological simulations such as cosmic variance

and butterfly effects are missing in our emulators since they are based on a simple fully connected neural network that outputs predictions in a deterministic way without having any randomness (refer to Section 2.5 and Appendix B for details of the *simulation uncertainty*). The *simulation uncertainty* can play a significant role in a probabilistic inference such as ILI, especially in quantifying uncertainty in the inferred parameters, which depends on the uncertainty of observables. In this section, we focus on (1) how emulators marginalize the *simulation uncertainty* and (2) the mock uncertainty that is implemented in emulators to mimic the *simulation uncertainty*.

Table 1
Average Relative Errors $\bar{\delta}$, Standard Deviations σ , and Coefficients of Variation c_v for Each Inference in Section 4.1

	$\bar{\delta}$ (%)				σ				c_v^a (%)							
	Figure 2		Figure 3		Figure 4		Figure 6		Figure 2		Figure 3		Figure 4		Figure 6	
	Cnvg.	U.C.	S.C.	Bmd.	Uncrt.	Cnvg.	U.C.	S.C.	Bmd.	Uncrt.	Cnvg.	U.C.	S.C.	Bmd.	Uncrt.	
Ω_m	0.02	7.9	0.9	0.45	...	0.004	0.016	0.004	0.015	0.073	0.79	3.2	0.8	3.1	16.0	
σ_8	0.49	6.1	1.0	0.36	...	0.003	0.014	0.003	0.014	0.100	0.42	2.0	0.4	2.0	14.7	
A_{SN1}	1.36	18.8	6.3	0.55	...	0.014	0.043	0.015	0.042	0.353	2.33	7.1	2.5	7.0	60.7	
A_{AGN1}	49.8	52.1	52.1	30.8	...	0.105	0.138	0.072	0.090	1.158	40.5	53.0	27.5	26.5	450.6	
A_{SN2}	0.25	8.5	1.6	0.41	...	0.007	0.026	0.009	0.032	0.145	0.76	3.2	1.1	3.9	17.8	
A_{AGN2}	3.65	17.8	8.2	5.7	...	0.040	0.117	0.036	0.087	0.358	5.01	14.8	4.6	10.5	47.0	
				...					0.148					8.1		

Notes. This table includes the values of the well-converged posterior (Cnvg.; see Figure 2), tests of convergence and stability of the ILI showing unstably convergent and stably convergent stages (U.C. and S.C.; see Figure 3), the bimodal posterior (Bmd.; see Figure 4), and the posterior inferred with the mock uncertainty (Uncrt.; see Figure 6). Since the bimodal posterior has two peaks, the corresponding columns (Bmd.) include two values.

^a See footnote 18.

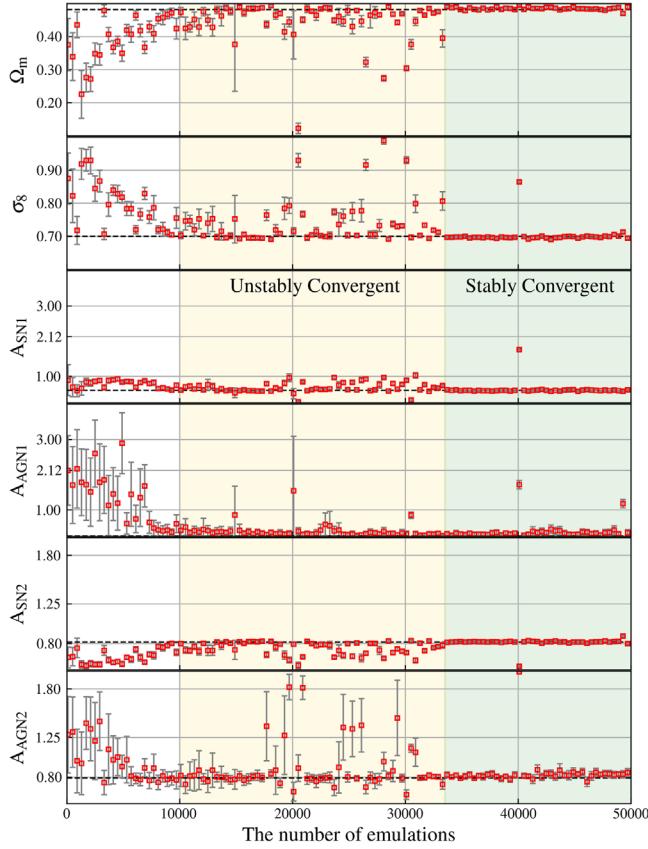


Figure 3. The convergence of each parameter as a function of the number of emulations used in the ILI on the emulator-based target SFRD used in Figure 2. The red squares and gray error bars present the maximum and the standard deviation of the posterior density. The black dashed lines show the true values. The convergence of ILI from the SFRD is divided into an unstably convergent stage (yellow region) and a stably convergent stage (green region) (see Section 4.1.1). In the unstably convergent stage, the maxima of the inferred posteriors jump around from the true values occasionally while the deviation of the corresponding SFRDs from the target is relatively small, which leads to degeneracy in Section 4.1.2. As the number of emulations for ILI exceeds $\sim 30,000$, the inferred posteriors reach the stably convergent stage.

First, to investigate how much the emulators marginalize the *simulation uncertainty*, we use the deviation of the emulator $f(\theta)$ from the uncertainty-marginalized ideal simulation $\bar{g}(\theta)$, which denotes the ideal, uncertainty-free, infinite-volume simulation, defining the deviation $\Delta(\theta) \equiv f(\theta) - \bar{g}(\theta)$. We approximate the deviation by taking an average over the ensembles of the thousand simulations in the LH set (note that the uncertainty-marginalized ideal simulation—e.g., the infinite-volume simulation—is unobtainable). Here, two assumptions are made: (1) The mean of the infinite number of simulation ensembles converges to the uncertainty-marginalized ideal simulation. (2) The *simulation uncertainty* is constant across the cosmological and astrophysical parameter space. The details of the assumptions are further described in Appendix D.

Using the deviation $\Delta(\theta)$, we estimate the mean of the deviations over the LH set that the emulators are trained on (i.e., the bias $b_{\text{LH}} = \langle \Delta(\theta) \rangle_{\text{LH}}$) and the standard deviation over the LH set ($\hat{\sigma}_{\text{LH}} = \sqrt{\langle (\Delta(\theta))^2 \rangle_{\text{LH}}}$), where $\langle \cdot \rangle_{\text{LH}}$ averages over all the parameters θ in the LH set. If $b_{\text{LH}} = 0$ and $\hat{\sigma}_{\text{LH}} = 0$, the emulator perfectly marginalizes the *simulation uncertainty* and the emulated prediction follows the uncertainty-marginalized ideal simulation. In this work, the emulators have a bias of $b_{\text{LH}} = 0.003$ dex, which can indicate that the emulators predict

the observables with a relatively high accuracy *on average*. However, given that the standard deviations from the emulators ($\hat{\sigma}_{\text{LH}} = 0.066$ dex) and simulations ($\sigma_{\text{sim,sfr}} = 0.057$ dex; refer to Appendix D for details) are comparable, it is unlikely for the emulators to marginalize the *simulation uncertainty*. In other words, the emulator predictions are in agreement with the uncertainty-marginalized ideal simulations *on average* but each emulator prediction of a point in parameter space has a similar variance to the *simulation uncertainty* with respect to the uncertainty-marginalized ideal simulations as the actual cosmological simulations have.

Without proper marginalization, the implementation of uncertainty in the emulator leads to a greater uncertainty in the parameters as well as in the observables. In addition, this weakens the connection between the simulations and emulators to a large extent in terms of physical interpretation of the parameters. Hence, in this work, we treat an emulator as ground truth or the mean of simulation uncertainty and implement a mock *simulation uncertainty* on top of the emulator. We model the *simulation uncertainty* using multivariate Gaussian noise with minor modifications. Then, the *mock simulation uncertainty*—hereafter the *mock uncertainty*—is added to the emulators manually (refer to Appendix B for details of the implementation).

4. Inference from the History of the Cosmic SFRD

4.1. Inference from Emulated Histories of the Cosmic SFRD

In this section, we study the performance and properties of ILI on the cosmic SFRD using an *emulator-based* SFRD rather than an observationally driven SFRD. To this end, we use the emulator to predict SFRDs as functions of the cosmological and astrophysical parameters and adopt an *emulated* SFRD as a target observation from which the cosmological and astrophysical parameters can be inferred. As discussed in Section 3.3, the emulators are adopted as surrogates for hydrodynamic simulations and are considered to be the ground truth throughout this work unless specified otherwise.

In contrast to deterministic approaches, probabilistic inference such as Bayesian inference and ILI has significant flexibility in that the inferred posterior distributions can take on versatile structures of probability distributions depending on the nature of the problems. Probabilistic distributions have three main beneficial aspects in our application: (i) the variance of the posteriors can be interpreted as the error bars of the inferred parameters (Section 4.1.1), (ii) the posterior can have an appreciable volume of parameter space that reproduces the same target observable within some accuracy (i.e., degeneracy; refer to Section 4.1.2), and (iii) we can measure the confidence interval of each parameter under the presence of uncertainty in the observations (Section 4.1.3). Lastly, in Section 4.2, we apply our ILI machinery to a measurement of an observationally derived estimate of SFRD and study how well the inferred posterior distribution can match the observations.

4.1.1. Performance of Inference

We first investigate the accuracy of our ILI and how stably the inferred posterior density converges without the mock uncertainty. The performance of ILI with the mock uncertainty will be discussed in Section 4.1.3. Using the suite of CAMELS simulations, we train an emulator that takes as input six cosmological and astrophysical parameters θ and predicts the SFRD x (for details, refer to Section 3.3). Then, we perform ILI to

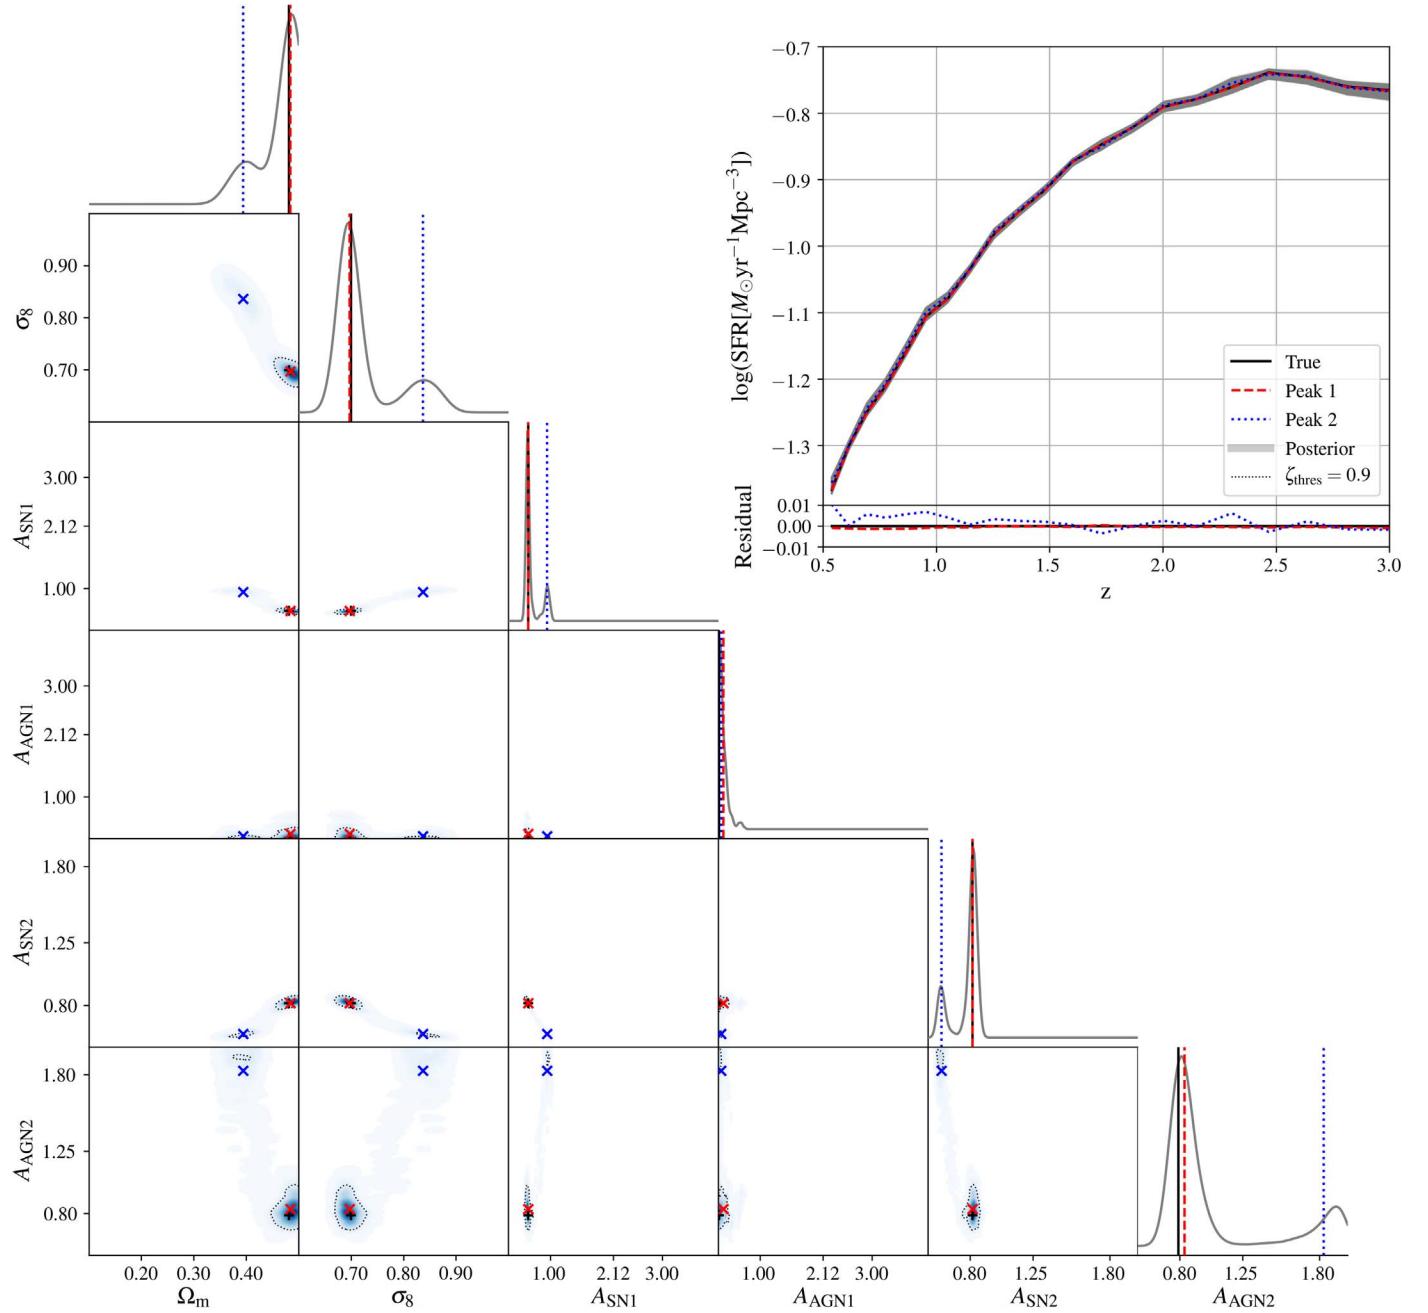


Figure 4. Top right: *Emulated SFRDs* from the inferred posterior (gray), two peaks of the posterior (red dashed and blue dotted), and the emulator-based target SFRD used in Figure 2 (black solid). Bottom left: Two-dimensional contour projections of the inferred posterior. The inferred posterior contains two strong peaks, both of which reproduce the SFRD well, with relative errors of 0.35% (red dashed) and 0.98% (blue dotted). This indicates a degeneracy in the SFRD (see Section 4.1.2).

retrieve the posterior density over the six cosmological and astrophysical parameters θ given an *emulator-based* target SFRD x_0 . Here, the *emulator-based* target SFRD x_0 is generated by the emulator with θ_0 that is one of the data points in the LH set.

Shown in Figure 2 are two-dimensional projections of the inferred posterior (bottom left) and an SFRD plot (top right) that includes the corresponding SFRDs drawn from the posterior as well as the emulator-based target SFRD. In this example, a total of 40,000 emulations are used to retrieve the posterior density $p(\theta|x_0)$ (gray) given the SFRD x_0 (black solid). The location of the maximum of the posterior density¹⁷

(red dashed line or red crosshairs) matches the true value (black solid line or black crosshairs) almost perfectly for all six parameters with small relative errors (see Table 1 for the values of the relative errors, standard deviations, and coefficients of variation¹⁸). The values of the standard deviations imply that the inferred posterior density has very small variances, in that compared to the volume of parameter space covered by the prior of $\mathcal{O}(10)$, the volume of the posterior density approximates to $\mathcal{O}(10^{-12})$. In the top right panel, the SFRD from the

¹⁷ The set of parameters giving the maximum value of the posterior density is drawn from the MCMC samples, i.e., $\theta_{\max} = \arg \max_{\theta} p(\theta|x_0)$.

¹⁸ The coefficient of variation, also known as the relative standard deviation, is a standardized measure of dispersion of a probability distribution or frequency distribution. It is usually defined as $c_v = \frac{\sigma}{\mu} \times 100$, but in this work we adopt a definition using the maximum of the posterior θ_{\max} instead of the mean μ , namely $c_v = \frac{\mu}{\theta_{\max}} \times 100$.

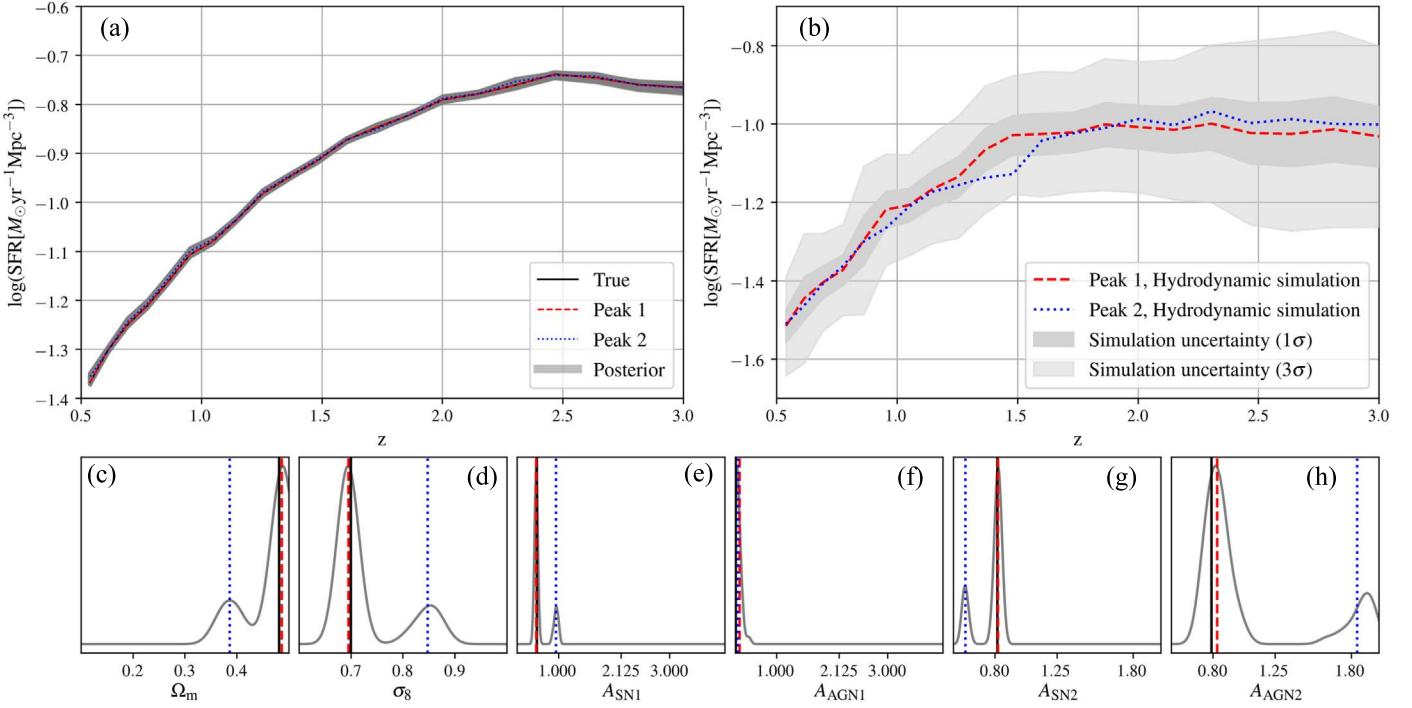


Figure 5. (a) Cosmic SFRD from the emulator, representing the posterior and its two peaks. (b) Cosmic SFRD from cosmological simulations using the points in parameter space that are the peaks of the posterior. (c)–(h) One-dimensional projections of the inferred posterior (gray) based on Figure 4. The two peaks are drawn with the red dashed and blue dotted lines consistent between Figure 4 and panels (c)–(h). However, the marginals of the *degenerate* posterior distribution in panels (c)–(h) are reconstructed with a modified probability density, namely that which is limited to only the region within $\zeta_{\text{degen}} = 0.9$ in Figure 4. The gray and light gray regions in panel (b) indicate the σ and 3σ confidence regions of the simulation uncertainty. Given that both simulated SFRDs from the two peaks lie within the 3σ region of the simulation uncertainty, we conclude that these cosmological simulations themselves are degenerate in terms of the SFRD (see Section 4.1.2).

maximum of the posterior (red dashed line) coincides with the target SFRD (black line) with a relative error of 0.17%. The SFRDs from the full posterior (gray region) have an exceedingly narrow distribution with a standard deviation of 0.003 dex.

Figure 3 illustrates the convergence of the ILI on the SFRD and its stability. Each panel shows the maximum of the posterior (red squares) and standard deviations (gray error bars) of each parameter as a function of the (cumulative) number of emulations used for training so far. The panel for A_{SN1} (third row) shows that the A_{SN1} parameter converges to the truth almost right after the beginning. The rest of the parameters (Ω_m , σ_8 , A_{AGN1} , A_{SN2} , A_{AGN2}) can seemingly come to convergence after 10,000 emulations. However, the convergence can be divided into two different stages: the “unstably convergent stage” in [10,000, 33,500] emulations (yellow region) and the “stably convergent stage” in [33,500, 50,000] emulations (green region).

In the unstably convergent stage (yellow region), the inferred parameters jump around occasionally without a particular period (unstably) but rapidly return to the truth in the next iteration (convergent). The average relative errors¹⁹, standard deviations, and coefficients of variation over [10,000, 33,500] emulations can be found in Table 1 (first, third, and fifth rows). All three measures above are an order of magnitude greater than those of the converged posterior in Figure 2. However, the mean of the relative errors of all the SFRDs,

including not only the maximum of the posteriors but also the posteriors themselves, at the unstably convergent stage approximates to 0.95%, which is clearly greater than that of a single convergent SFRD (0.17%) but the value itself seems acceptable. The most interesting feature of the stage is that although the inferred parameters are jumping around, the corresponding SFRDs are relatively well converging to the truth. This is attributed to the weak correlation between the parameters and the SFRD (see Section 6.1 for a discussion). More importantly, this implicitly indicates the possibility of a multimodal distribution that can reproduce the same observable from different sets of the parameters, as further discussed in Section 4.1.2.

On the other hand, the posterior distributions are stably convergent with relatively small variances after $\sim 30,000$ emulations (green region). All six maxima (red squares) for the six parameters stably converge to the truth (black dashed line) with the average relative errors (see footnote 10) and their standard deviations over [33,500, 50,000] emulations. The variances of the posterior density (gray error bars) are convergent as well (refer to the S.C. columns of Table 1). These values are an order of magnitude lower than the values at the unstably convergent stage, which can be a clear sign of transition. Also, in comparison to the values for the posterior in Figure 2, the average relative errors of the stably convergent posteriors are somewhat greater but the standard deviations and coefficients of variation are comparable. That is, the peaks or the maxima of the posteriors have appreciable scatter as compared to the truth at the stably convergent stage, whereas the widths of the posteriors are consistent. As for the variances, the convergence properties not only depend on the number of

¹⁹ The average relative error of a parameter θ is defined as $\bar{\theta}_\theta \equiv \frac{1}{N} \sum_n (\theta_{\text{truth}} - \arg \max_\theta p^{(n)}(\theta | \mathbf{x}_0)) / \theta_{\text{truth}} \times 100$, where n enumerates over a range of emulations in the ILI process (e.g., Figure 3) and $N = \sum_n 1$. \mathbf{x}_0 is the target observation.

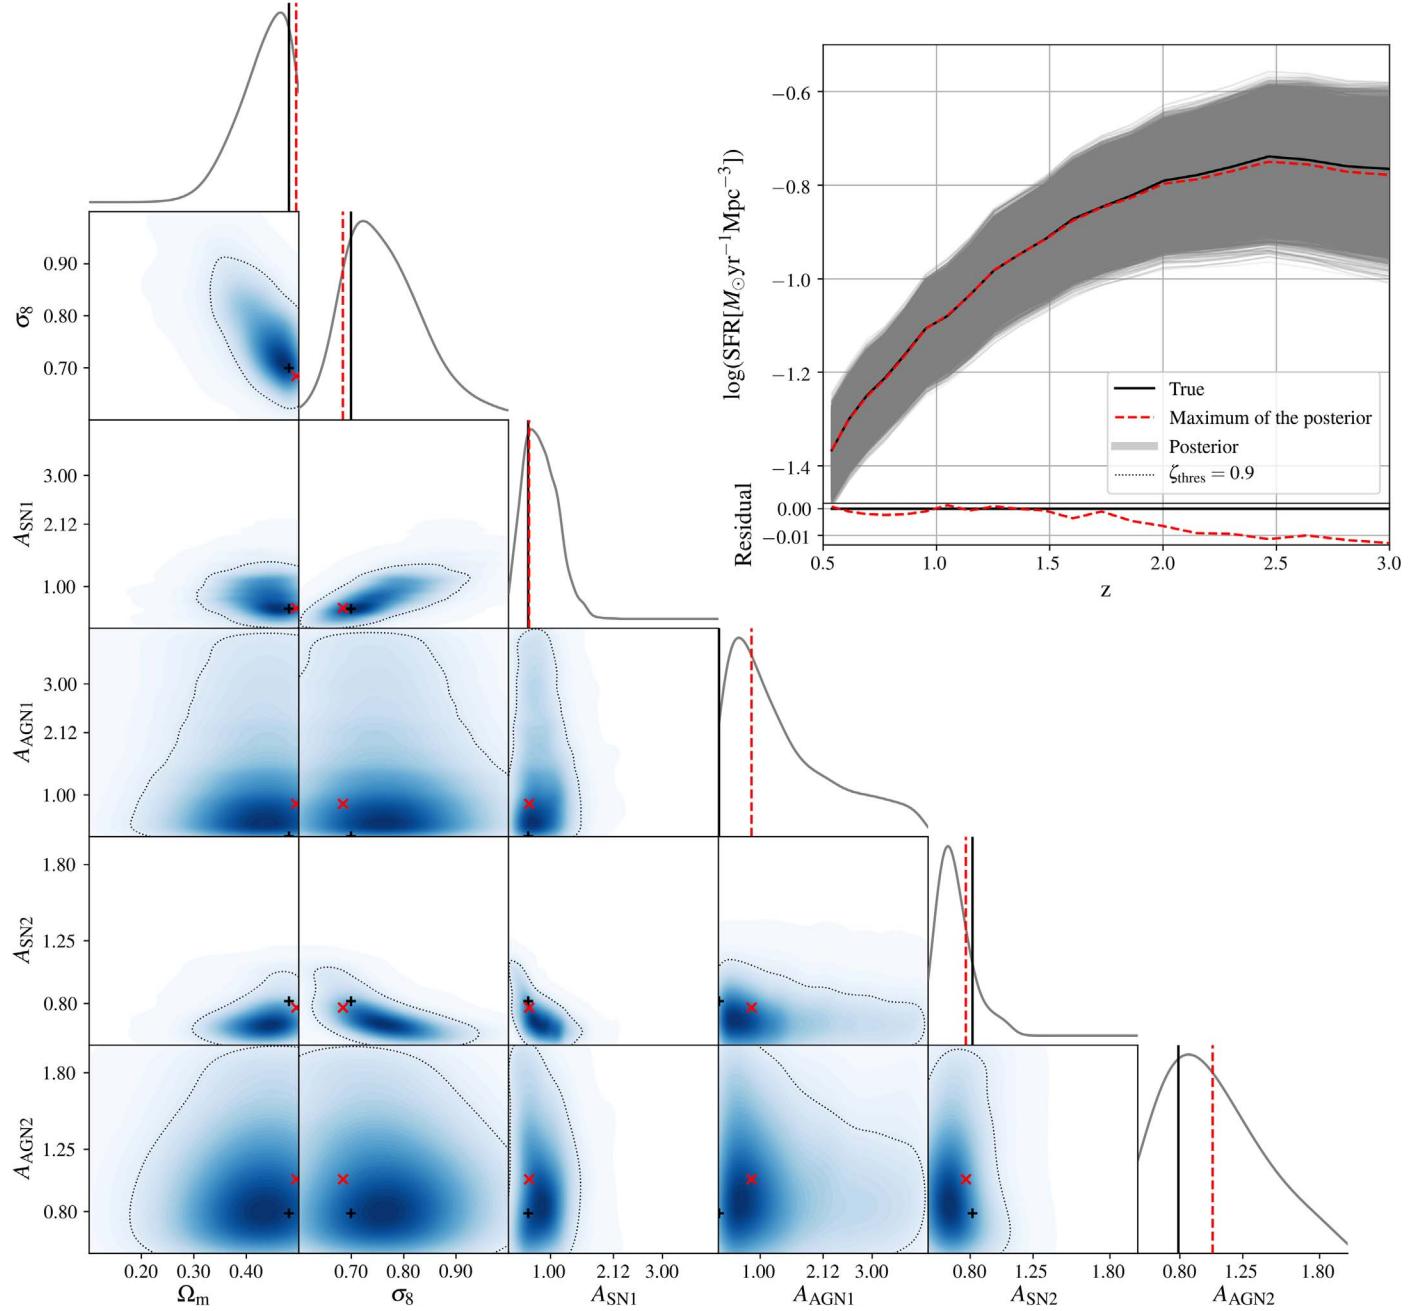


Figure 6. Bottom left: Two-dimensional projections of the posteriors (blue) inferred from the emulator-based target SFRD used in Figure 2 with the mock uncertainty. Top right: Cosmic SFRD from the inferred parameters (gray) and the truth (black solid). The red vertical lines indicate the maximum values of the inferred posteriors. The mock uncertainty in ILI leads to large variances in the inferred marginal distributions compared to the sharp marginal distributions of Figure 2 inferred without the mock uncertainty (see Section 4.1.3 and for a quantitative comparison see Table 1).

emulations but also differ by parameter. In the convergence of $A_{\text{AGN}1}$, there is no evident transition from unstably convergent to stably convergent in terms of the relative errors, standard deviations, or coefficients of variation. On the other hand, the stellar feedback parameters ($A_{\text{SN}1}$ and $A_{\text{SN}2}$) converge rapidly as soon as the training begins, leading to a seamlessly smooth transition to the unstably convergent stage. However, the transition to the stably convergent stage takes place drastically both visually and quantitatively. Meanwhile, the AGN feedback parameters ($A_{\text{AGN}1}$ and $A_{\text{AGN}2}$) have notable scatter even in the stably convergent stage (green region) as compared to the other four parameters. Hence, we can conclude that the strengths of the correlations between each

parameter and the SFRD are in the following order: $A_{\text{SN}1} \simeq A_{\text{SN}2} > \Omega_m \simeq \sigma_8 > A_{\text{AGN}1} \simeq A_{\text{AGN}2}$. An indication for such correlations can also be found in Villaescusa-Navarro et al. (2021, Figures 11 and 12).

Thus far, we have not implemented any uncertainties in the inferences. Therefore, assuming that the emulators are injective (one-to-one functions), a negligible variance is expected for all six parameters, i.e., $p(\mathbf{x}, \boldsymbol{\theta}) = \delta(\mathbf{x} - f(\boldsymbol{\theta}))$ and $\boldsymbol{\theta}$ is unique. However, a tiny amount of variance exists in both the posterior and the inferred SFRDs (gray in Figure 2). This can be attributed to (1) physical degeneracy, which is discussed in Section 4.1.2, and (2) inaccuracy of the NDE. In both Figures 2 and 3, Ω_m , σ_8 , $A_{\text{SN}1}$, and $A_{\text{SN}2}$ show high convergence and

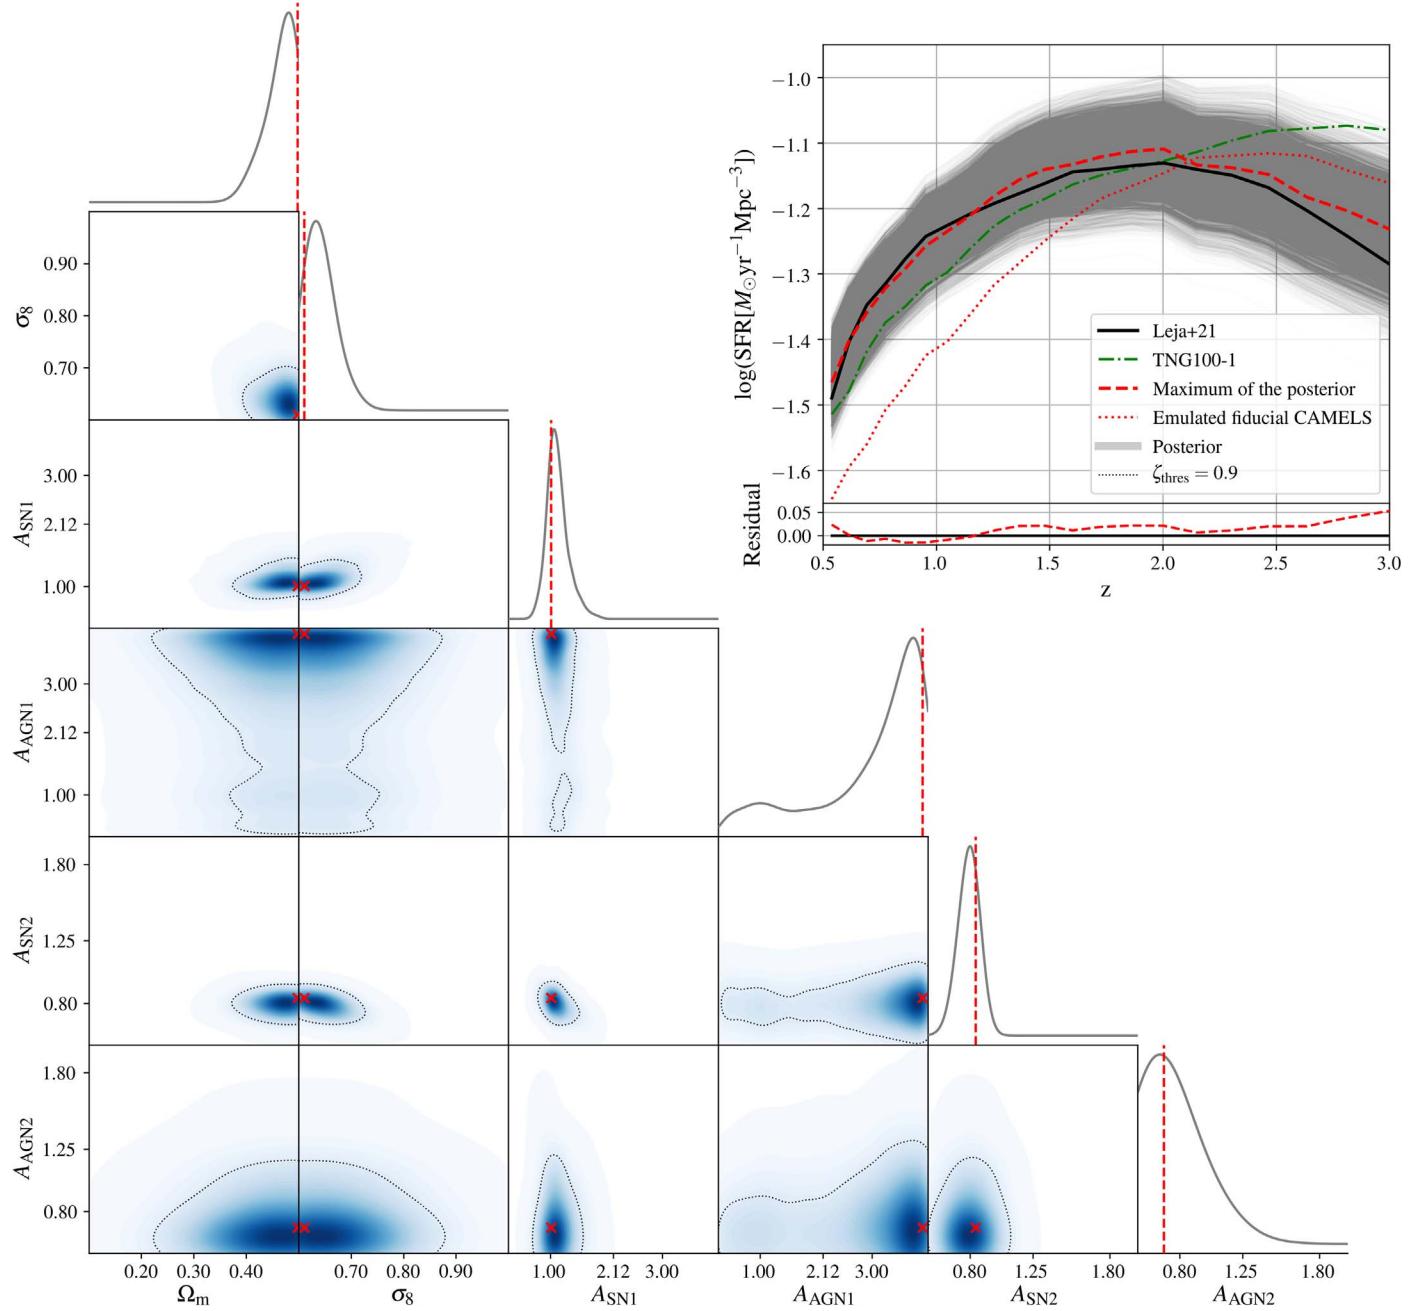


Figure 7. Top right: The inferred cosmic SFRDs (red dashed) and the observationally driven target SFRD (Leja et al. 2022) (black solid). Bottom left: Two-dimensional projections of the inferred posteriors (gray) from the observationally driven target SFRD. The red dashed vertical lines indicate the maximum values of the inferred posteriors in the six-dimensional parameter space. For comparison, the SFRD of the TNG100-1 simulation and the emulated fiducial SFRD of the CAMELS simulations are drawn in the top right panel (green dotted-dashed and red dotted, respectively). The inferred SFRD matches the observed SFRD with a relative error of 4.1% and a deviation of 0.04 dex, whereas the deviation of the fiducial CAMELS SFRD from the observed SFRD is ~ 0.15 dex (see Section 4.2). The inferred cosmological parameters are quite off from common sense due to, e.g., resolution effects and degeneracies with the astrophysical parameters (see Section 6.4).

precision for both maxima and variances whereas $A_{\text{AGN}1}$ and $A_{\text{AGN}2}$ have larger variances on average. The magnitude of the variance indicates how strongly the observable can constrain each parameter, or how intimately each parameter and observable correlates with the other. The greater the variance, the weaker the correlation. Due to the weak correlations, $A_{\text{AGN}1}$ and $A_{\text{AGN}2}$ require more simulations to converge stably and tend to have a larger variance than the other parameters. Nevertheless, the relative errors of both parameters and SFRDs are less than 1% on average with a total of 34,000 emulations.

4.1.2. Bimodality in the Posterior Distribution

In this section, we present one of the bimodal posterior distributions that can be found in the unstably convergent stage (yellow region) in Figure 3. Figure 4 illustrates two-dimensional projections of the inferred posterior (bottom left) as well as the corresponding SFRDs from the posterior and the (emulator-based) target SFRD (top right). The most intriguing feature in Figure 4 is the *bimodal peaks* (red and blue) in the posterior density. Not only do the two peaks exist in the

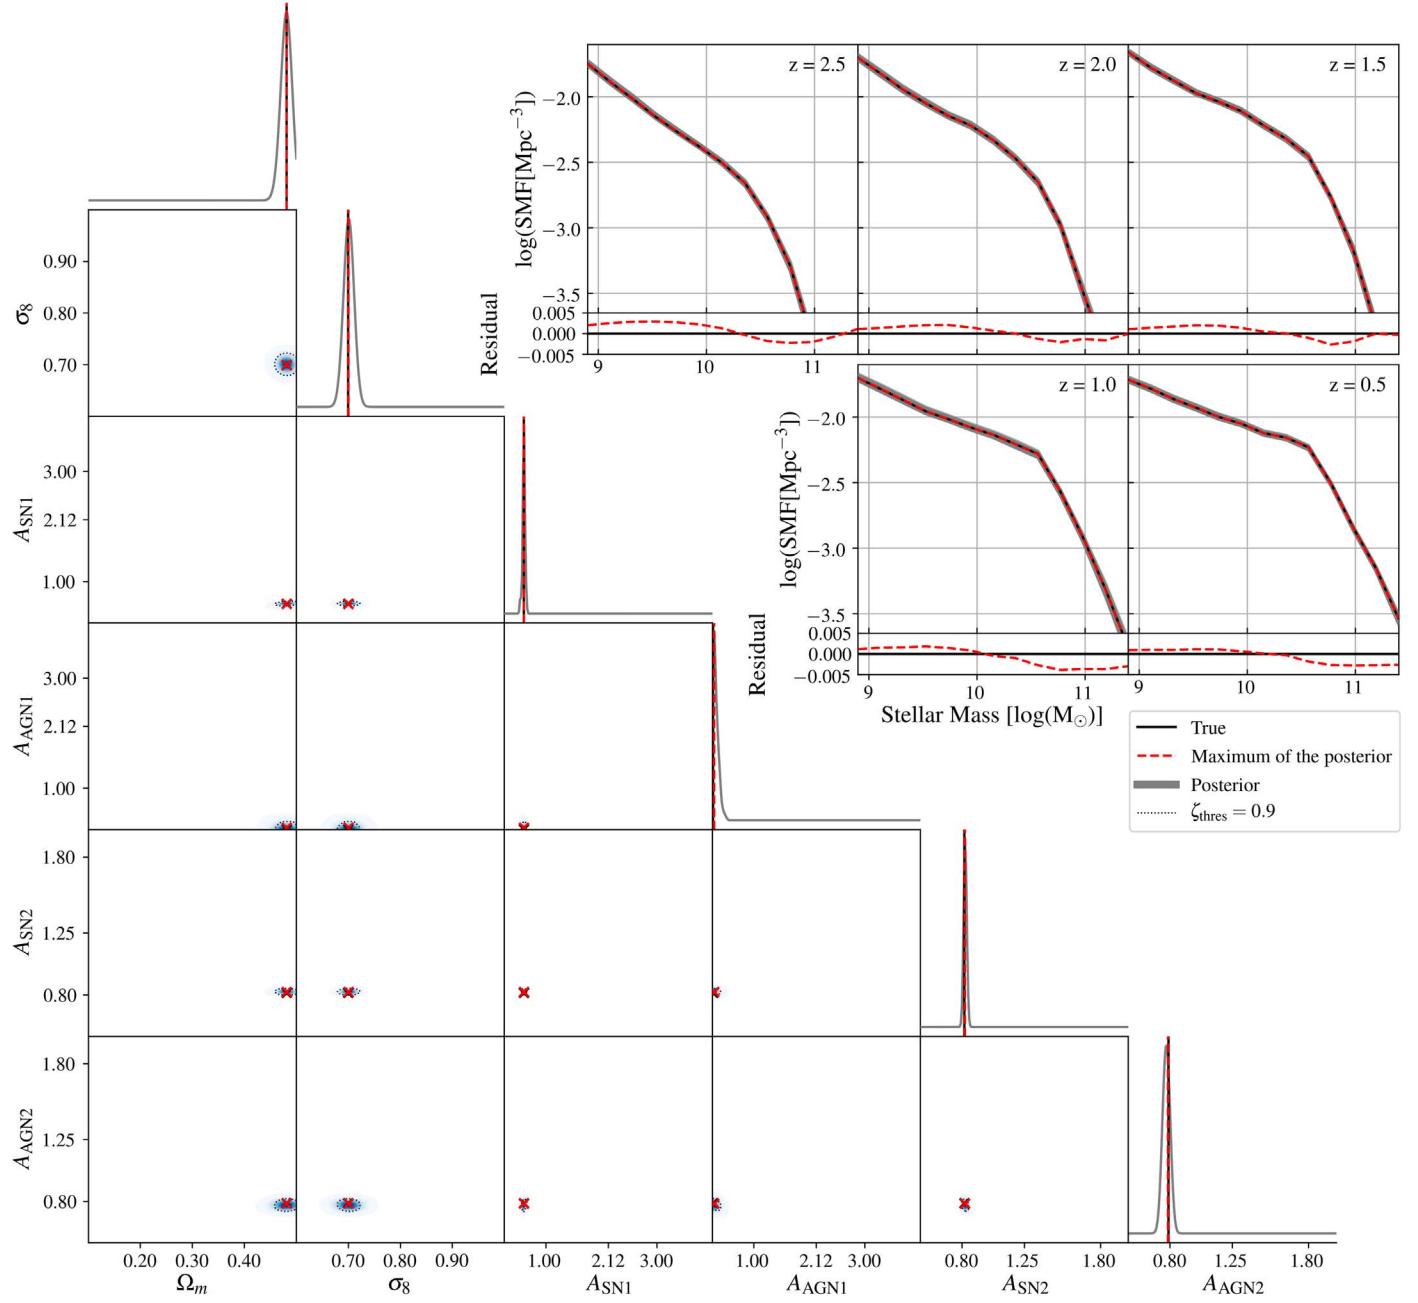


Figure 8. Top right: *Emulated* SMFs from five different redshifts from the inferred posterior (gray), the maximum of the posterior (red dashed), and the emulator-based target SMFs (black solid). Here, the emulator-based target SMFs are generated by the emulator with the same set of parameters that is used in Section 4.1 and Figure 2. Bottom left: Two-dimensional projections of the inferred posterior. The black and red crosshairs represent the values of the target and the maximum of the posterior, respectively. The black solid and red dashed vertical lines indicate the true values and the maximum values of the inferred posterior, respectively. The inferred parameters (red dashed) and true values (black solid) are nearly on top of each other with small errors and variances (see Table 2 for details). The inferred SMF and the target SMF are also in good agreement with a relative error of 0.4% (see Section 5.1).

posterior density, but they also reproduce the target observations within a barely appreciable margin of error.

We use k -means clustering to divide the posterior density into two unimodal distributions. The k -means clustering partitions samples into k clusters in which each sample belongs to the cluster with the nearest mean. Each peak is defined by the center of each cluster. The relative errors of the two SFRDs with respect to the true SFRD are 0.35% for peak 1 and 0.98% for peak 2. Peak 2 has slightly larger errors than peak 1 but they are still less than 1%. Furthermore, peak 1 (red dashed) accurately coincides with the truth (black solid) with small relative errors (see Table 1). We measure the standard

deviations and coefficients of variation with respect to each peak by using the result of the k -means clustering (see Table 1). There is no marked difference between the variances of peak 1 and peak 2. The level of deviation is slightly higher than that of the stably convergent posterior in Figure 2 but similar to those of the unstably convergent ones shown in Figure 3.

Having two strong peaks in the posterior distribution is the result of a physical *degeneracy*, namely a situation where more than a single set of parameters reproduces the same observable. To study this, we propose a definition of degeneracy in a mathematically consistent way using a given posterior

Table 2
Average Relative Errors $\bar{\delta}$, Standard Deviations σ , and Coefficients of Variation c_v for Each Inference in Section 5.1

	$\bar{\delta}$ (%)		σ		c_v (%)	
	Figure 8	Figure 9	Figure 8	Figure 9	Figure 8	Figure 9
	Cnvg.	Stb.	Cnvg.	Stb.	Cnvg.	Stb.
Ω_m	0.15	1.4	0.006	0.007	0.9	0.2
σ_8	0.01	0.3	0.002	0.002	0.2	0.1
$A_{\text{SN}1}$	0.08	1.2	0.024	0.022	2.7	2.6
$A_{\text{AGN}1}$	4.9	34.1	0.076	0.076	17.7	4.9
$A_{\text{SN}2}$	0.72	0.6	0.006	0.006	0.6	0.7
$A_{\text{AGN}2}$	0.12	1.9	0.018	0.028	2.3	20.2

Note. This table includes the values of the well-converged posterior (Cnvg.; see Figure 8) and of tests of the convergence and stability of ILI (Stb.; Figure 9).

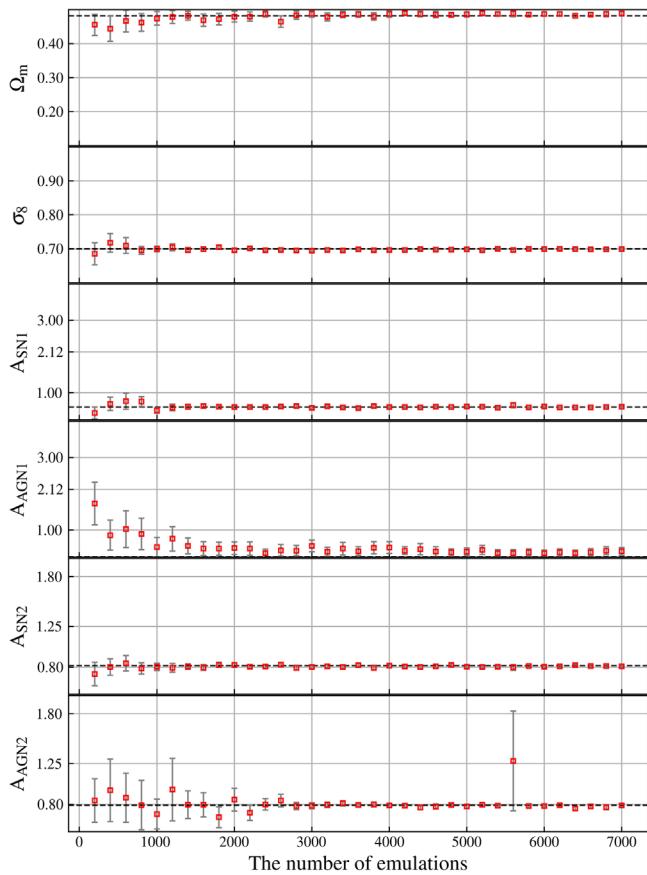


Figure 9. The convergence of each parameter as a function of the number of emulations for ILI from emulator-based target SMFs. The red squares and gray error bars represent the maximum and the standard deviation of the posterior density. The black dotted lines show the true values. The posteriors inferred from the SMFs rapidly and stably converge (see Section 5.1 and Table 2).

distribution. The set of degenerate points in parameter space, Θ_{degen} , is defined such that it satisfies $\int_{\theta \in \Theta_{\text{degen}}} p(\theta) d\theta = \zeta_{\text{thres}}$, where $p(\theta \in \Theta_{\text{degen}}) > p(\theta \notin \Theta_{\text{degen}})$, assuming that the inferred posterior distribution $p(\theta)$ is normalized. Here, ζ_{thres} is a free parameter, and the degenerate parameter set Θ_{degen} collects parameters according to their probability density in descending

order until the integration of the probability over the degenerate set becomes equal to ζ_{degen} (refer to Appendix E for a precise and detailed definition). In this work, we set ζ_{degen} to 0.9. Also, every two-dimensional projection of posterior distributions in this paper includes a contour line for $\zeta_{\text{degen}} = 0.9$.

Shown in Figure 5(c) are the marginals of the *degenerate* posterior distribution that is reconstructed with the probability density only enclosed within $\zeta_{\text{degen}} = 0.9$ through the Gaussian kernel density estimation. On the other hand, the two peaks (red dashed and blue dotted) are identical to the ones in Figure 4, which illustrates details of the posterior distribution with $\zeta_{\text{thres}} = 0.9$ (black dotted contour) in the two-dimensional projection plot (bottom left). The two emulated SFRDs that correspond to the two peaks of the posterior distribution (red dashed and blue dotted) in panel (a) are approximately on top of each other. The distribution of SFRDs (gray) is also sufficiently concentrated with a variance of 0.003 dex that is similar to that of the SFRDs from the stably convergent posterior in Figure 2. Hence, we can conclude that the two peaks are degenerate in terms of the emulated SFRDs.

However, the presence of degeneracy in the emulated SFRD does not necessarily demand that a degeneracy exists in the actual cosmological simulations as well. Therefore, we test this with new simulations. Panel (b) in Figure 5 shows the *simulated* SFRDs obtained with each set of parameters from the two peaks (red dashed and blue dotted). Here, we investigate whether the two sets of parameters are also degenerate in the simulations given the simulation uncertainty (see Section 2.5). We use $\pm 1\sigma$ (gray) and $\pm 3\sigma$ (light gray) regions that correspond to 68.1% and 99.7% confidence levels assuming that the simulation uncertainty follows the Gaussian distribution, respectively. The standard deviations of the Gaussian distribution are directly calculated from the suite of simulations in the CV set along the redshift (refer to Appendix B). In panel (b), the simulated SFRDs from the parameter combinations corresponding to the degenerate peaks are consistent within 1σ . Thus, the two simulated SFRDs are highly likely to be degenerate, sharing the same parameters, while having small discrepancies originating from the simulation uncertainty.

Above all, the emulator result demonstrates clear signs of degeneracy in the SFRD having two approximately identical SFRDs from two different sets of parameters. On the

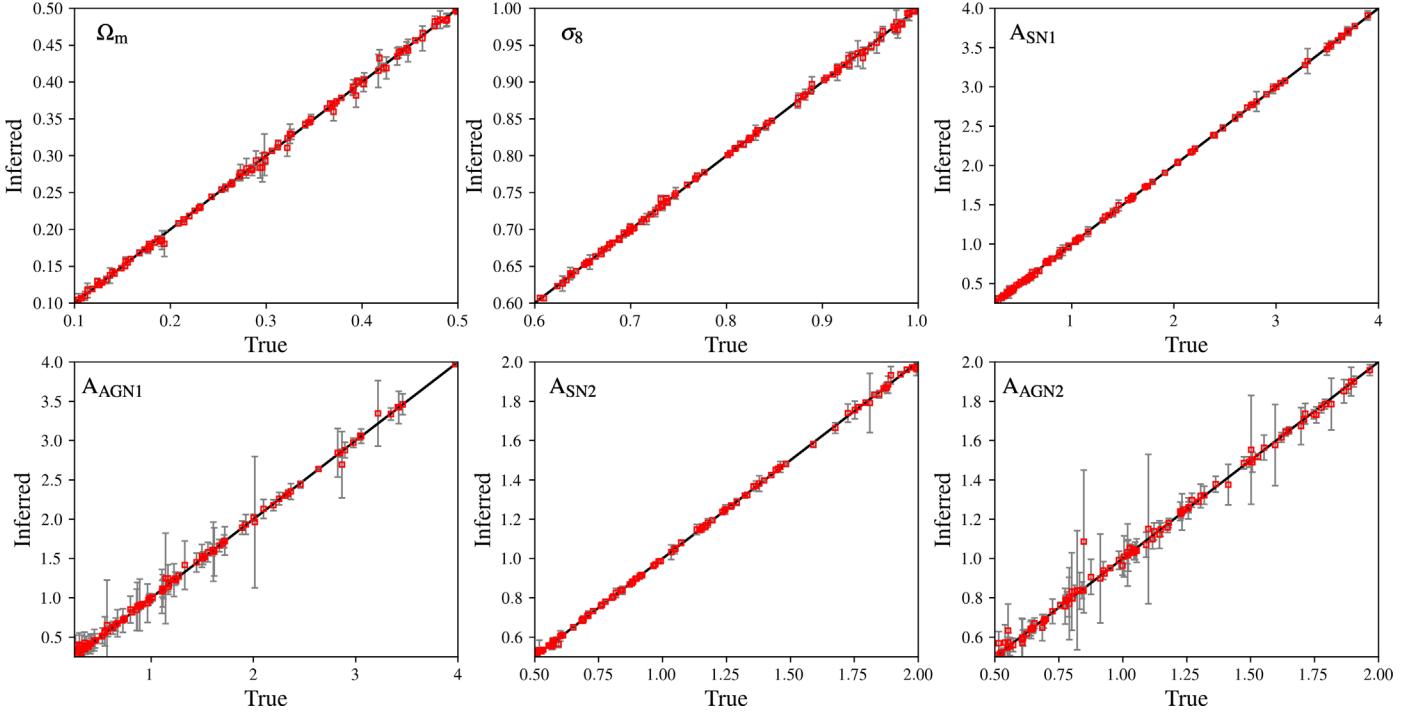


Figure 10. Comparison between the true values and the corresponding inferred values from the posteriors for 100 SMFs generated by the emulator using 100 different combinations of parameters without mock uncertainty. Each red square represents the maximum value of the inferred posterior against its true value. Each gray bar represents the standard deviation of the inferred posterior. The black solid line represents “inferred equals true.” ILI can infer parameters from SMFs accurately with a relative error of $\sim 1\%$ on average across the parameter space (see Section 5.1).

simulation side, cosmological simulations with the two sets of parameters also produce two SFRDs that are close to each other, though the simulation uncertainty leads to small but notable differences in the simulated SFRDs, which are nevertheless not statistically significant. The SFRD lies approximately inside the 64.8% (1σ) confidence regions and is completely enclosed within the 99.7% (3σ) confidence regions of the simulation uncertainty. This is strong evidence that both emulator and cosmological simulations have degeneracies in the SFRD, in spite of the emulator not being a perfect representation of the simulations.

4.1.3. Response to Mock Uncertainty

Thus far, we have performed ILI without including any uncertainty in the forward model, namely the emulator, but such uncertainties do exist in full cosmological simulations, as discussed in Section 2.5. In this section, we include the *simulation uncertainty* that originates from various sources of randomness in the cosmological simulations using the mock uncertainty on top of the emulators (refer to Section 3.3.1 and Appendix B for a technical description of the mock uncertainty).

Shown in Figure 6 are two-dimensional projections of the posterior density inferred from the emulated SFRD (bottom left) with the mock uncertainty. The mock uncertainty that we impose is modeled to have the same standard deviation as the simulation uncertainty (see Appendix B). In comparison to the inferences without the mock uncertainty, the posterior densities inferred with the uncertainty cover much larger regions of parameter space (see Table 1). In comparison to Figure 4, the bimodal peaks are merged to a single oval or banana-shaped distribution (bottom left). Although the deviations of the inferred $A_{\text{AGN}1}$ and $A_{\text{AGN}2}$ become notable, the relative error of

SFRD is 1.0%, and the deviation is only 0.0044 dex. The variance of the inferred SFRDs (0.057 dex) is comparable to the variance of the mock uncertainty (0.061 dex). This demonstrates that the inferred posterior densities and the corresponding observables successfully reproduce the mock uncertainty in terms of the variances. Thus, concerning the uncertainty propagation from observation to parameters, the variance of the inferred parameters can be reliable. On the other hand, the inclusion of the mock uncertainty has led to an increase in the standard deviation of the posterior as well as in the relative errors of the inferred observables with respect to the truth. However, given the size of the standard deviations, the inferred parameters are still accurate. We can relate it to the stochasticity of sampling of the *mock uncertainty*. In every iteration of the ILI, additional training data $(\mathbf{x}, \boldsymbol{\theta})$ is generated from the proposal density. Here, the observable \mathbf{x} is emulated as a function of the sampled parameters together with the mock uncertainty $Z(\eta)$. Due to the finite size of sampling, the mean of observables $\langle \mathbf{x} + Z(\eta) \rangle_{\text{samples}}$ cannot be the same as the ideal (theoretical) mean $\langle \mathbf{x} + Z(\eta) \rangle = \langle \mathbf{x} \rangle$, leading to a bias in the sampled data. Notice that for an infinite number of samples (the ideal case), $\langle Z(\eta) \rangle = 0$ since $Z(\eta)$ is a Gaussian noise that has a mean of zero. The bias in the newly generated training data is highly likely to result in the inaccuracy of the inferred posterior density.

4.2. Inference from the Observed SFRD

We now apply our framework to actual observational data. Here, we perform ILI from the observationally driven SFRD of Leja et al. (2022) with an identical setup to that in the previous section, which includes the mock uncertainty. Figure 7 illustrates two-dimensional projections of the posterior density inferred from the observed SFRD. Note that we do not include

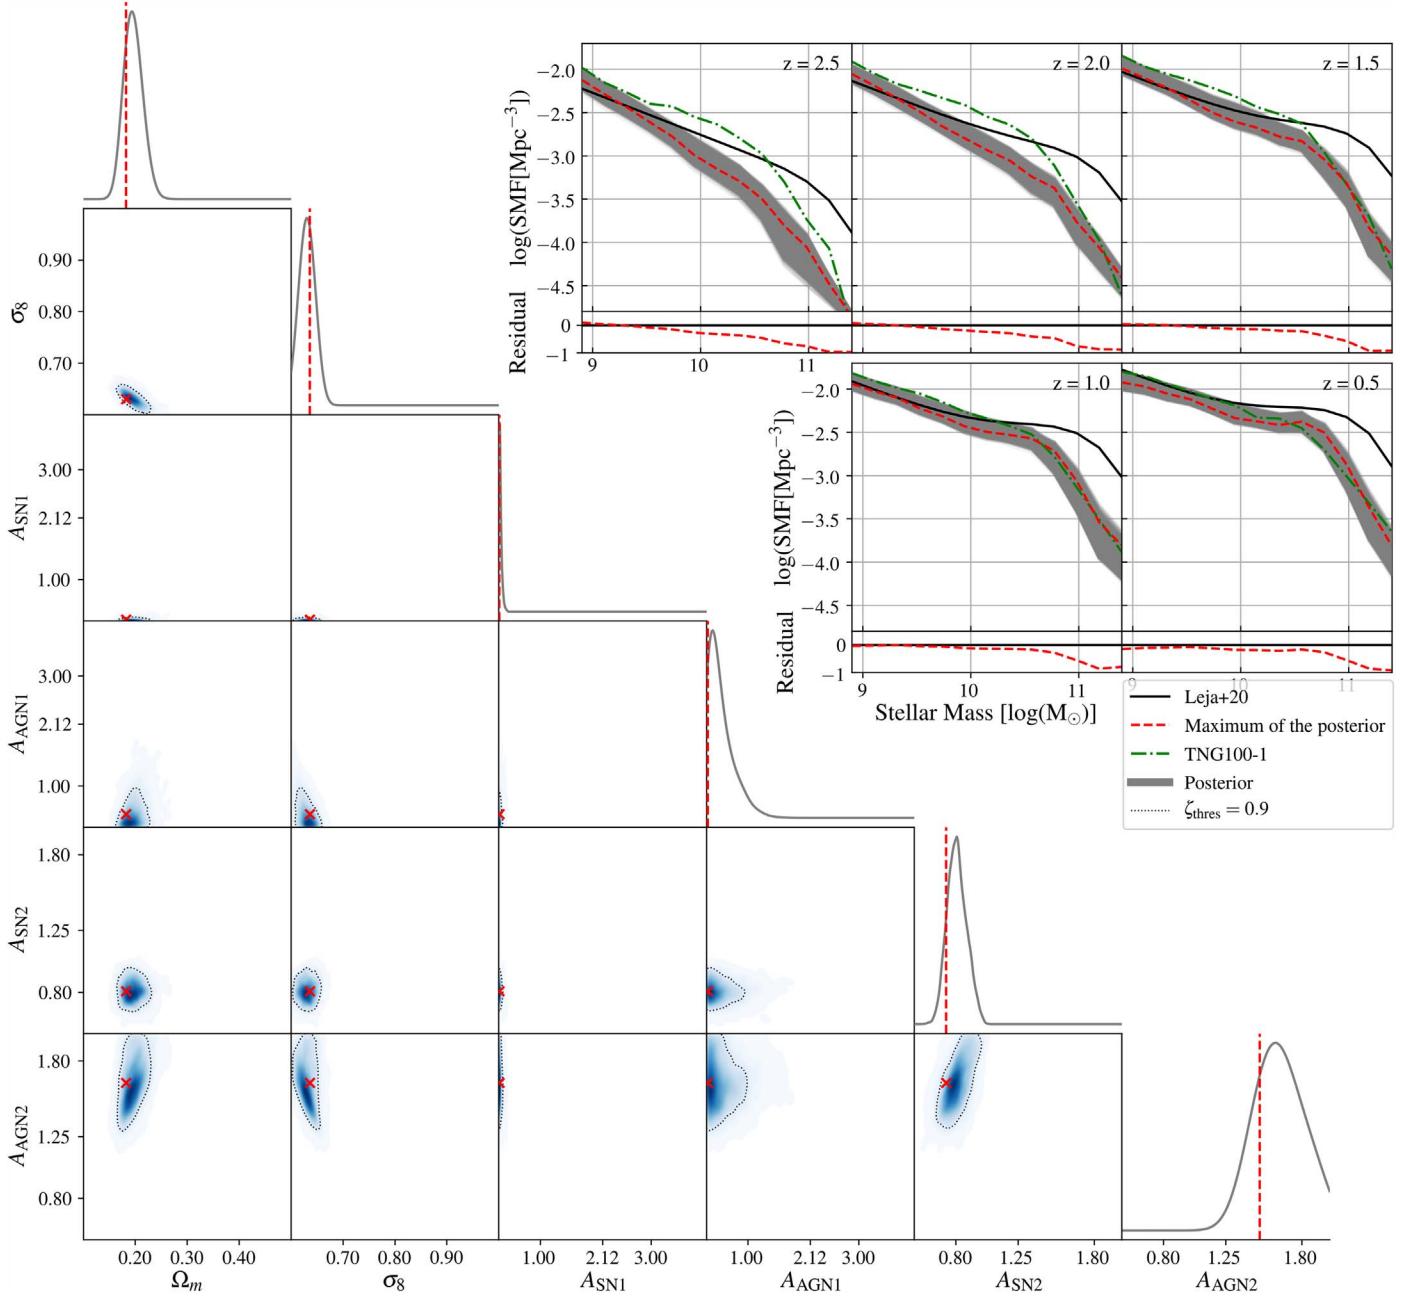


Figure 11. Top right: The SMFs from the inferred posteriors (gray and red) and the observationally driven target SMFs (black). Bottom left: Two-dimensional projections of the inferred posteriors for the observationally driven target SMFs from five different redshifts ($z = 2.5, 2.0, 1.5, 1.0, \text{ and } 0.5$) (Leja et al. 2020). The red vertical lines indicate the maximum values of the inferred posterior. The inferred SMFs and the target SMFs show considerable discrepancy with a relative error of 41.1% (see Section 5.2.1).

the uncertainty of the observation data (see Section 2.3); rather, only the mock uncertainty is adopted. The observed SFRD (black solid) lies completely within the region of the inferred SFRDs (gray), and also the SFRD from the maximum of the posterior (red dashed) matches it with a relative error of 4.1%. The standard deviations and the coefficients of variation are (0.040, 0.039, 0.224, 1.475, 0.078, 0.181) and (8.3%, 5.6%, 37.1%, 566.9%, 9.7%, 23.0%), respectively (for comparison, see Table 1).

This can be thought of as a successful inference given that the emulated fiducial SFRD in Figure 7 (red dotted) shows entirely different trends from the target observation over all redshifts, with an average deviation of ~ 0.1 dex and a relative

error of 22.8%. Here, the emulated fiducial SFRD is an SFRD generated by the emulator from the fiducial parameters. Despite the huge discrepancy between the fiducial SFRD and the target observation, the inferred SFRD follows the observation relatively well with a mean deviation less than 0.02 dex (relative error of 4.1%) and even matching the observed peak precisely. That being said, not only do the inferred parameters (red crosshairs) have significant discrepancy with the standard values (e.g., standard cosmology), but also the posterior distribution itself hardly includes the standard values. This inconsistency mostly stems from several reasons such as resolution effects and degeneracy between the cosmological and astrophysical parameters, but not from the inference

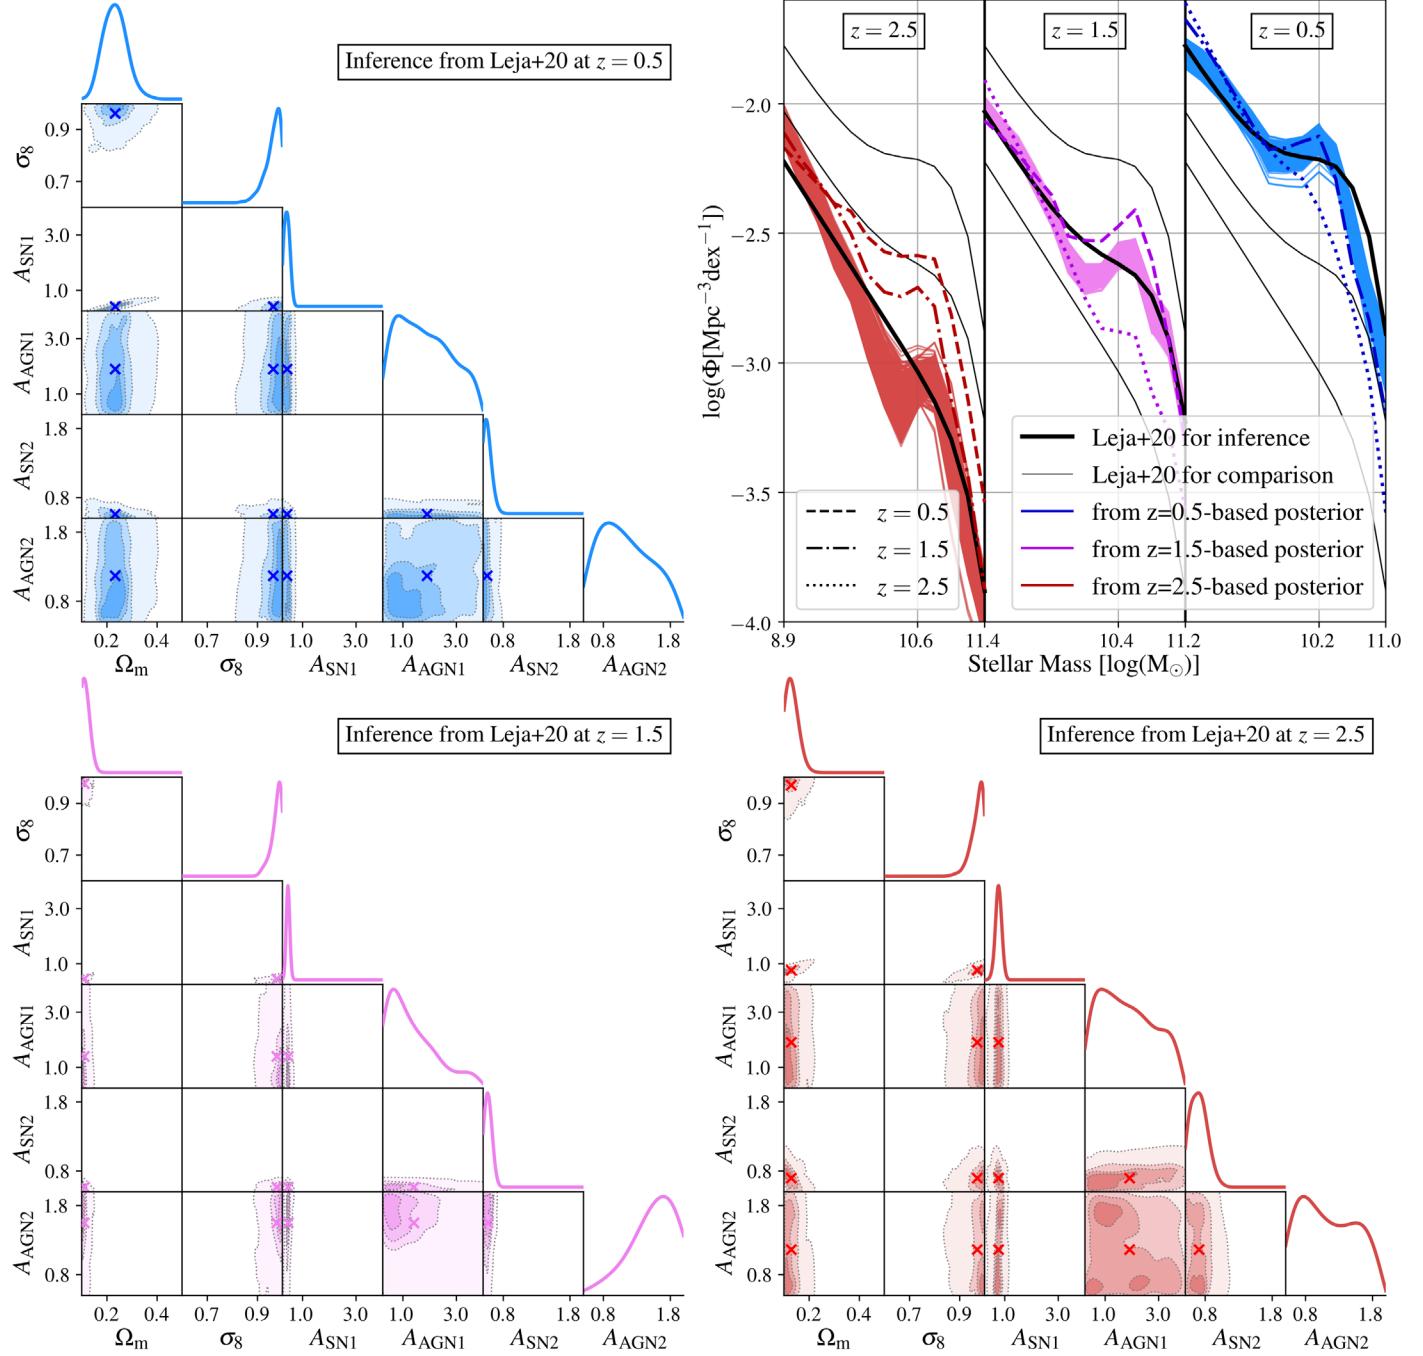


Figure 12. Two-dimensional projections of the posteriors inferred from the individual (observationally driven target) SMFs (Leja et al. 2020) at three different redshifts: $z=0.5$ (top left, sky blue), $z=1.5$ (bottom left, violet), and $z=2.5$ (bottom right, red). The top right panel shows three inference results from the observationally driven target SMFs (thick black solid) at different redshifts $z=2.5$ (left), 1.5 (middle), and 0.5 (right). Here, the thin black solid lines are the observed SMFs that are not involved in the inferences at a particular redshift. The dashed, dotted-dashed, and dotted lines are the emulated SMFs at $z=0.5$, 1.5 , and 2.5 from the maximum of the posteriors inferred from the observed SMFs at $z=0.5$ (dark blue), 1.5 (dark violet), and 2.5 (dark red), respectively. The relative errors of each SMF with respect to each observationally driven target SMF are 17.7% for $z=2.5$, 10.3% for $z=1.5$, and 10.3% for $z=0.5$. Compared to that of the inference from the concatenated SMFs, the accuracy is increased, but still, the inferred SMFs cannot precisely match the observationally driven target SMFs compared to the SFRD (see Section 5.2.2).

procedure (see Section 6.4 for details). In short, the fiducial astrophysical parameters of the CAMELS simulations were obtained by calibrating the TNG100-1 simulation to the observations. However, the fiducial CAMELS simulation and TNG100-1 simulation show nonnegligible discrepancy due to

the resolution effect and simulation box size (Section 6.5.1). This would inevitably lead to a discrepancy between the inferred parameters and the standard values even in the hypothetical case of inference from the same observations against which TNG100-1 was calibrated. In addition, our target

Table 3

The Maximum of the Posteriors, θ_{\max} , in Figures 11 and 12 and the Standard Deviations σ with Respect to the Maximum of the Posteriors

	Ω_m	σ_8	A_{SN1}	A_{AGN1}	A_{SN2}	A_{AGN2}
θ_{\max}	ALL ^a	0.18	0.64	0.26	0.27	0.73
	$z = 0.5$	0.22	0.97	0.41	1.50	0.51
	$z = 1.5$	0.10	1.00	0.47	0.25	0.58
	$z = 2.5$	0.11	1.00	0.72	2.77	0.83
σ	ALL ^a	0.02	0.01	0.03	0.31	0.10
	$z = 0.5$	0.02	0.04	0.11	1.32	0.18
	$z = 1.5$	0.01	0.03	0.04	1.38	0.04
	$z = 2.5$	0.05	0.03	0.11	1.02	0.06

Note.

^a The parameters are inferred from all five redshifts simultaneously (red crosshairs in Figure 11).

SFRD and SMFs (Leja et al. 2020, 2022) are different from those used for calibrating TNG100-1. Lastly, inconsistency of cosmological parameters can be attributed to the compensatory reaction to lack of some astrophysics (e.g., Sections 5.2.2 and 6.4.1).

In the case of inference from observations, the intrinsic limit of the emulators should also be accounted for. The most significant difference between the inferences from the emulated SFRD and observed SFRD is that any emulated SFRD can be predicted precisely by the emulator for sure, whereas we cannot ensure whether there exists a data point in parameter space that can reproduce the observation perfectly. We can define the image of the emulator as the set of all SFRDs and SMFs it can produce (and similarly for simulation predictions). The image of an emulator might not coincide with the set of all physically possible SFRDs, including observed SFRDs: it is a codomain of the SFRDs. In other words, emulators and simulations might not be able to reproduce any possible universe. A necessary condition for successful precision inference is that the observed SFRD is a member of the image of the emulator. Concerning the limit of the emulator and the simulations that the emulator is trained on, there are fundamental issues, e.g., (1) the limited dimensions of the parameter space (domain), (2) the gap between the emulators and the simulations, and (3) the limits of physical models in simulations. The above will be discussed in Section 6.4.

5. Inference from SMFs

We now turn to inference on SMFs as the target observable. We study the dependence of the properties of the ILI and the inferred posteriors on the choice of observable with a comparison to the cosmic SFRD. The latter covers the evolutionary history of the universe ranging from $z = 3$ to $z = 0.5$ whereas a single SMF contains information of only one iteration. To be consistent with the SFRD, we concatenate five SMFs at $z = 0.5, 1.0, 1.5, 2.0$, and 2.5 and each SMF is binned with 13 bins in the mass range $[10^{8.9}, 10^{11.4}] M_{\odot}$ (refer to Section 2.3). Hereafter, “SMFs” denote the five concatenated SMFs from the five different redshifts throughout the paper unless specified otherwise. In Section 5.1, we first investigate the performance and convergence of the inference on the

Table 4

The MI between the Parameters and Observables (First–Fourth Rows)

	Ω_m	σ_8	A_{SN1}	A_{AGN1}	A_{SN2}	A_{AGN2}
<i>SimSFR</i>	1.19	0.62	0.54	0.29	2.11	1.43
<i>SimSMF</i>	12.5	9.27	36.9	1.17	4.69	0.41
<i>EmuSFR</i>	11.4	6.82	30.0	0.7	17.1	0.85
<i>EmuSMF</i>	14.6	12.8	41.4	1.49	5.79	0.30
δ_{sfr} (%)	0.20	0.49	1.36	49.8	0.25	3.65
δ_{smf} (%)	0.15	0.01	0.08	4.9	0.72	0.12

Note. The relative errors δ of the convergent SFRD and SMFs are from Sections 4.1.1 and 5.1, respectively (fifth and sixth rows). *Sim* and *Emu* stand for “simulated” and “emulated.” The values of the MI are multiplied by 100 for convenience and the unit for the relative error is percentage. Note that the higher the MI is, the more strongly the parameter–observable pairs are correlated, whereas the MI is zero when they are independent.

emulator-based target SMFs compared to those of the SFRDs. Then, we perform ILI from observationally driven five concatenated SMFs (Section 5.2.1). Finally, in Section 5.2.2, we study ILI from one *individual* observationally driven SMF (at a single redshift) at a time.

5.1. Inference from Emulated SMFs

We first investigate the performance of ILI and how stably the inferred posterior converges in terms of the SMFs. As in Section 4.1.1, we train emulators that take as input six cosmological and astrophysical parameters θ and predict the SMFs \mathbf{x} (for details, refer to Section 3.3). Using the emulators, we perform ILI to retrieve the posterior density for six cosmological and astrophysical parameters θ given the emulator-based target SMFs \mathbf{x}_0 .

Shown in Figure 8 are two-dimensional projections of the inferred posterior (bottom left) and *emulated* SMFs from the posterior along with the emulator-based target SMFs (top right). Here, the emulator-based target SMFs \mathbf{x}_0 are generated by the emulator with the parameters that we use to generate the emulated SFRD in Section 4.1. A total of 6000 emulations are used to retrieve the posterior density $p(\theta|\mathbf{x}_0)$ given the emulator-based target SMFs \mathbf{x}_0 . The maximum of the posterior density (see footnote 8) (red dashed line or red crosshairs) matches the true value (black solid line or black crosshairs) nearly perfectly for all six parameters (see Table 2 for details). In the top right panel, the SMFs from the maximum of the posterior (red dashed) coincide with the true SMFs (black solid) with a relative error of 0.4%. The SMFs from the full posterior (gray region) have an exceedingly narrow distribution with a mean standard deviation of 0.007 dex. In comparison to those of the SFRD in Section 4.1.1, the relative errors of the SMFs are slightly better except for A_{SN2} , whereas the standard deviations of the SMFs are comparable to those of the SFRD. In general, the error of inference—e.g., relative errors—and the size of the uncertainty region—e.g., standard deviations—can be related to the correlation between the parameters and the observable. We discuss why A_{SN2} can be more precisely predicted from the SFRDs in Section 6.1, where we show how surprisingly the errors presented here are well explained by the correlations.

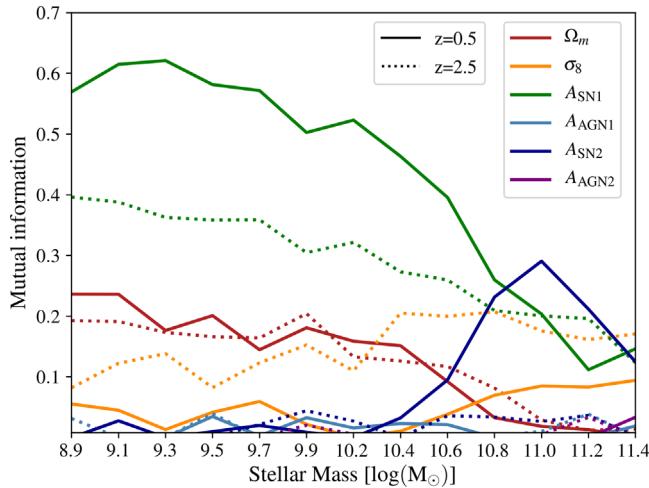


Figure 13. MI of the SMF at $z = 0.5$ (solid) and 2.5 (dotted) and each parameter as a function of stellar mass. Each colored line indicates the MI between one parameter and the SMF.

Figure 9 illustrates the convergence and its stability for the inference from the emulated SMFs. Each panel shows the maxima (red squares) and variances (gray error bars) of the posterior for a given parameter as a function of the (cumulative) number of emulations used for training so far. All six maxima (red squares) stably converge to the true values (black dotted lines) over [4800, 8000] emulations (see Table 2 for details). Hence, the convergence of the inference requires at least ~ 4000 emulations.

In contrast to the case of the SFRD, an unstably convergent stage is absent in the SMFs. To be stably convergent, the SMFs require only 4000 emulations whereas $\sim 40,000$ emulations are needed for the SFRD, an order-of-magnitude difference. Moreover, the average variances at the stably convergent stage, $\bar{\sigma}_{\text{smf}} \sim \mathcal{O}(1)$ for the SMFs and $\bar{\sigma}_{\text{sfrd}} \sim \mathcal{O}(10)$ for the SFRD, demonstrate that the SMFs converge with less fluctuation. Hence, we conclude that the SMFs converge far more rapidly and stably to the truth compared to the SFRD. Furthermore, the NDE of the SFRD has more hidden units ($N_{\text{hid,sfrd}} = 250$) than that of the SMFs ($N_{\text{hid,smf}} = 100$). Given that the more neurons there are, the more easily a parameter converges, we can conclude that the SMF-parameter pairs are more easily and tightly mapped than the SFRD-parameter pairs. This indicates that the degree of correlation between the SMFs and the parameters is stronger than that between the SFRD and the parameters (as discussed further in Section 6.1).

We also investigate the accuracy of the inference for various SMFs over the parameter space. Figure 10 shows the maxima (red squares) and variances (gray error bars) of 100 posterior densities inferred from 100 *emulated* SMFs. Here, the 100 SMFs are generated by the emulators with 100 randomly sampled sets of parameters from the LH set (refer to Section 2). We consistently perform the ILI with a total of 4800 emulations over all 100 SMFs without any further convergence tests. Due to computational cost, the minimal (necessary but perhaps not sufficient) number of emulations for convergence is adopted²⁰, deduced from the previous convergence test in Figure 9. Most

²⁰ For the same reason, we cannot produce this plot for the SFRD since the SFRD requires significantly more simulations than the SMFs and owing to the numerous hidden units of the NDE for the SFRD, the computational time is at least 10 times greater than that of the SMFs.

maxima (red squares) are on top of the ideal prediction line (black solid line) or are indistinguishably close to it. The average relative errors (see footnote 10) of the maximum are $(1.1 \pm 1.2, 0.2 \pm 0.2, 1.3 \pm 2.2, 4.0 \pm 8.5, 0.6 \pm 0.8, 1.8 \pm 3.4)\%$. These values are in line with the relative errors from the convergence test. Thus, the ILI can be performed on the SMFs stably with relatively small, constant errors regardless of the choice of parameters.

5.1.1. No Response to Uncertainty

We perform ILI on the SMFs with the mock uncertainty (for details of the mock uncertainty for the SMFs, refer to Section 3.3 and Appendix B). However, unlike the case of the SFRD in Section 4.1.3, which shows appreciable variances in the inferred posterior, the posterior density inferred from the SMFs with the mock uncertainty is essentially *identical* to the posterior density inferred *without* the mock uncertainty. Since we cannot tell any difference both visually and quantitatively, the corresponding figure for the *with-uncertainty* case is not presented. In Section 6.3, we explain this result by studying in detail how the inferred posteriors respond to various types of uncertainties and which type of uncertainty is suitable for ILI.

5.2. Inference from Observed SMFs

5.2.1. Five Concatenated SMFs

Figure 11 shows two-dimensional projections of the posterior density inferred from the observationally driven target SMFs (Leja et al. 2020) (bottom left) and the corresponding SMFs (top right). The discrepancy between the observation (black solid) and the maximum of the posterior (red dashed) is considerable with a mean relative error of 41.1% over all mass ranges. The discrepancy of the high-mass end of the SMFs is dominant with the relative error for stellar masses $\gtrsim 10^{11} M_{\odot}$ being $\sim 80.9\%$ versus 13.7% at the lower-mass end. The high-mass end ($\gtrsim 8 \times 10^9 M_{\odot}$) of the observed SMFs is located completely outside the region of the SMFs from the full posterior (gray) across all five redshifts. The size of the population of massive galaxies in the observed SMF is far greater than that in the inferred SMFs.

At the low-mass end, even though the observed SMFs lie within the gray region, the emulated and observed SMFs have different, distinct characteristics: (1) different slopes at high redshift and (2) different evolutionary rates. At $z = 2.5$ and $z = 2.0$, the slopes of the inferred SMFs are notably steeper than those of the observed SMFs. The observed SMFs do not show appreciable differences in slope across the redshifts, whereas the slopes of the inferred SMFs become less steep as redshift decreases. Also, there is a significant difference in the evolutionary rates of the SMFs. The evolution of the inferred SMFs is barely notable, but apparent growth can be seen in the observation. For example, the differences between the SMFs at $M_{\star} = 10^{8.9} M_{\odot}$ at $z = 2.5$ and $z = 0.5$ are 0.20 dex for the inference and 0.47 dex for the observation.

5.2.2. Three Individual SMFs

To isolate the problem, we perform inferences from one SMF at a time for each selected redshift ($z = 0.5, 1.5$, and 2.5) separately. Figure 12 illustrates the three posterior densities inferred from the observed SMFs (thick solid lines) at $z = 0.5$ (sky blue), $z = 1.5$ (violet), and $z = 2.5$ (red) and the

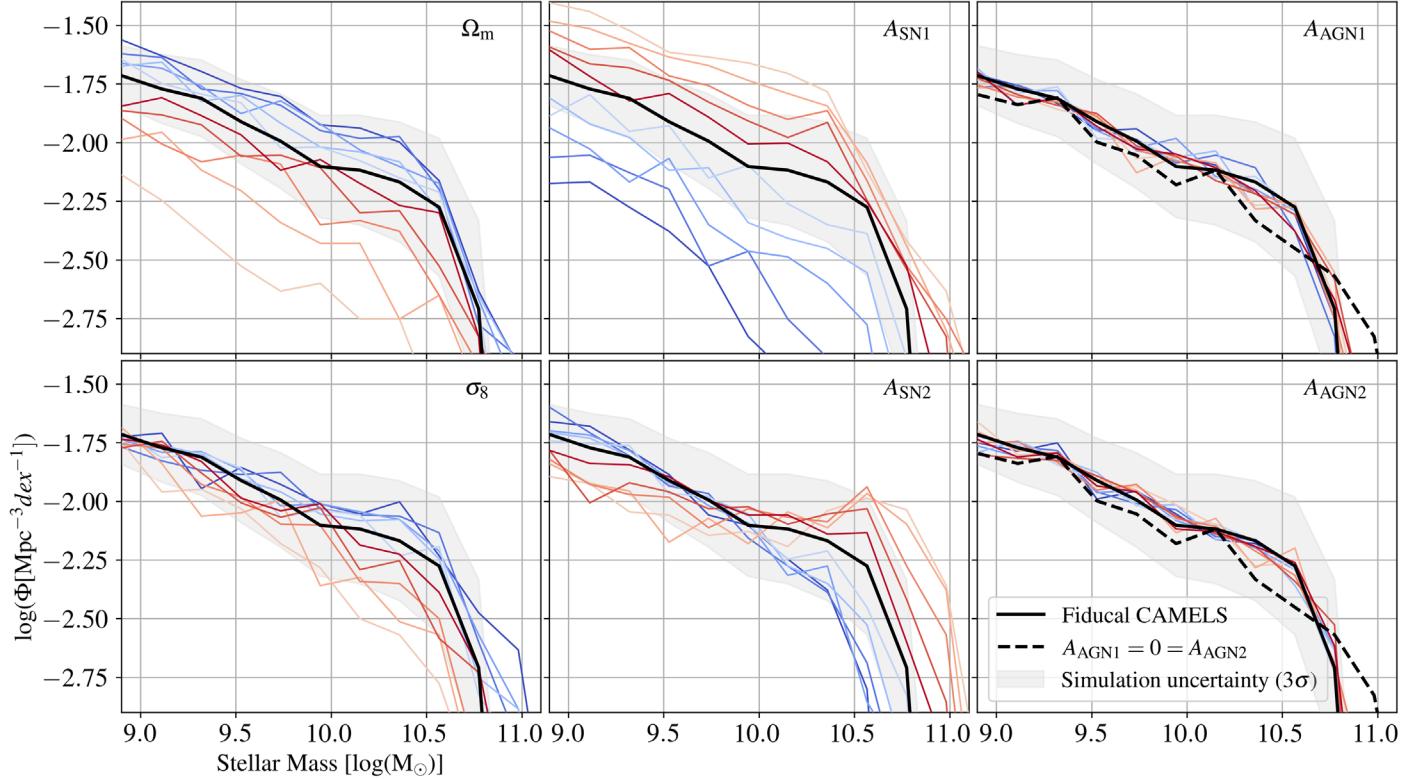


Figure 14. Simulated SMFs at $z = 0.5$ from the 1P set (higher parameter values with bluer curves and lower values with redder curves), from the fiducial parameter combination (black solid), and without AGN feedback (black dashed). The light gray region indicates the 3σ region of the simulation uncertainty. The variations of Ω_m , σ_8 , A_{SN1} , and A_{SN2} (left and middle panels) lead to considerable changes in the SMFs, whereas the variations of A_{AGN1} and A_{AGN2} (right panels) have a negligible impact on the SMFs. However, the inclusion of kinetic AGN feedback obviously has an impact on the SMF by comparison between the weakest AGN feedback (reddest) and no AGN kinetic feedback (black dashed) (see Section 6.1.1 for details).

corresponding SMFs (top right). The separate inferences lead to significantly higher accuracy than the previous inference, with relative errors of the SMFs of 17.7% for $z = 2.5$, 10.3% for $z = 1.5$, and 10.3% for $z = 0.5$ (the relative error of the previous inference is 31.4%). In comparison to that in the previous inference (Figure 11), the loci of the posterior densities are notably different, especially in σ_8 and A_{SN1} (refer to Table 3). The standard deviations of the AGN parameters in Figure 12 are notably larger than those from the inference based on the five concatenated SMFs in Figure 11, whereas the standard deviations of the inferred SMFs are similar. This might indicate that the impact of the AGN parameters on the SMFs is trivial. In general, the impact of the AGN parameters on the SMFs is negligible in this analysis, given the extensive variances that cover almost the entire parameter space. The average variances of A_{AGN1} and A_{AGN2} over the three redshifts are 1.23 and 0.62, whereas the variances of the other parameters are less than 0.01 on average.

The σ_8 values from the separate inferences are considerably higher than those from the combined-SMF inference in Section 5.2.1, whereas Ω_m shows relatively small changes. In both emulators and simulations, the AGN parameters only have a minor effect on the stellar populations of galaxies (Section 6.1). As a compensatory action, the emulated universes exploit the cosmological parameters, especially σ_8 , to control the populations of massive galaxies in the context of structure formation. Meanwhile, to sustain the density of low-mass galaxies, stronger supernova feedback is required, which we can see in Figure 12 as increases in A_{SN1} .

In the top right panel of Figure 12, each of the three subpanels corresponds to the inference based on an observed SMF from a different redshift, $z_{obs} = (0.5, 1.5, 2.5)$ (thick black curves). The SMFs generated by the emulators for $z_{emu} = (0.5, 1.5, 2.5)$ with the maximum of the corresponding posteriors are drawn as dashed, dashed-dotted, and dotted curves, respectively, and in each panel the presented SMFs for the corresponding redshift sample the full posterior in thin colored curves that appear largely as continuous shaded regions. Note that we have three independent emulators that estimate the SMFs at three different redshifts. For example, the emulated SMF at $z = 0.5$ from the maximum of the $z = 1.5$ based posterior (bottom left, violet) is drawn with a dark violet dashed line in the middle subpanel. Each subpanel exhibits the discrepancies of the slope and evolution of the SMFs across different redshifts from a single posterior. Notice that although the separate inferences can predict observations with much improved accuracy as compared to the inference from the combined SMFs, the discrepancies are not resolved. Further discussion of the discrepancies between inferred SMFs and observed SMFs can be found in Section 6.4.1.

On the simulation side, the significant mismatch between the inference and observation can be attributed to the following: (1) the limited volume of simulations, (2) the low resolution of the simulations and the failure of rescaling, (3) limitations of the physical models of the simulations, and (4) inaccuracy of the emulators. The size of the simulation is directly linked to the level of simulation uncertainty that we fail to realize in the SMFs (refer to Sections 5.1.1 and 6.3 for details). Since we cannot account for the simulation uncertainty,

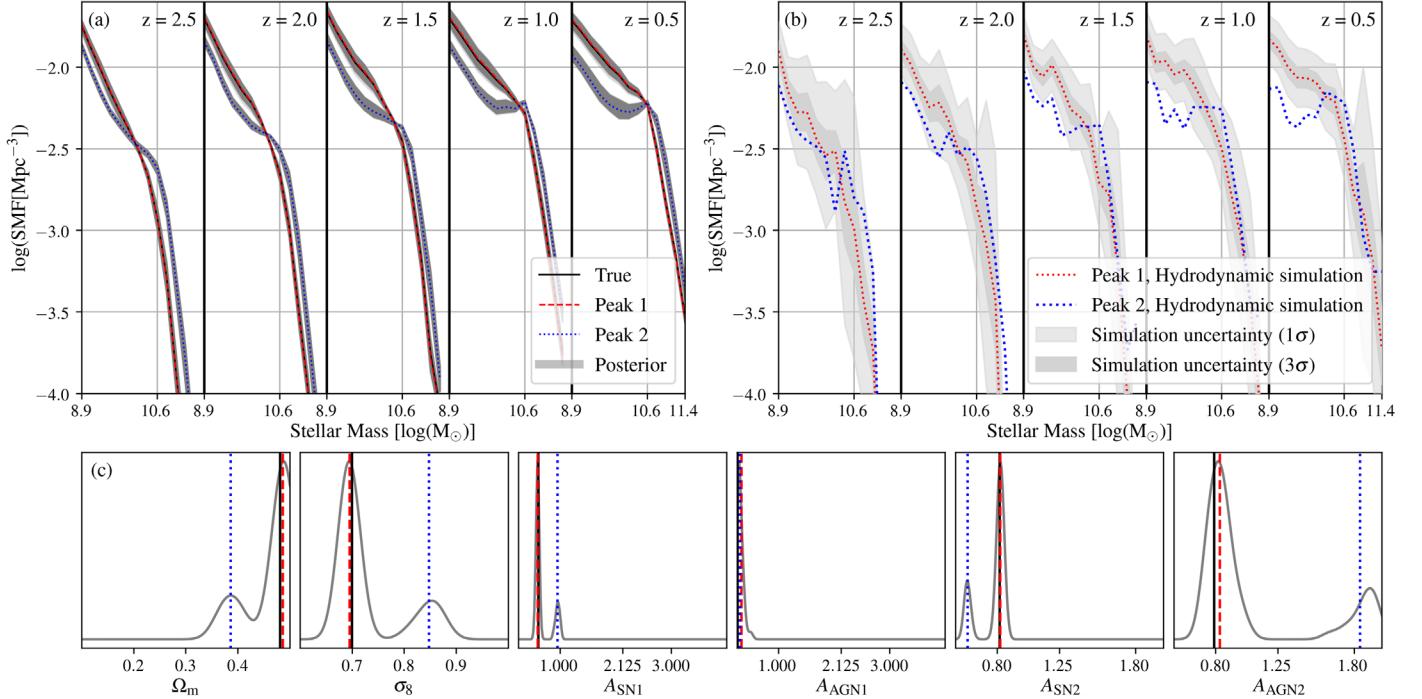


Figure 15. (a) SMFs from the emulator. (b) SMFs from the cosmological simulation. (c) One-dimensional projections of the inferred posterior based on Figure 4. The two peaks that produce nearly identical SFRDs that are degenerate in Figure 4 are drawn with red dashed and blue dotted lines consistently. For details, refer to Section 4.1.2. Note that there exists an offset between the emulated SMFs and simulated SMFs. This is because the emulated SMFs are rescaled, whereas the simulated SMFs come right from the cosmological simulations. Given that both simulated and emulated SMFs from the two peaks lie outside the 3σ region of the simulation uncertainty, the degeneracy can be broken with the SMFs (see Section 6.2).

the inference results suffer from sampling bias. Also, lack of massive galaxies is prevalent in low-resolution simulations (Appendix A in Pillepich et al. 2018a, 2018b; Appendix A in this paper). Even though we apply rescaling to alleviate the resolution effects, the rescaled SMFs are still subject to resolution convergence (refer to Appendix A). Inside the parameter space that we adopt in this work, there might not exist a set of parameters that can reproduce the observations simultaneously across redshifts. Lastly, since we use the emulators as a surrogate for cosmological simulations, the inference results are also limited by the accuracy of the emulators. Note that observational uncertainties, such as systematic biases and physical limits, can also be crucial factors for the mismatch but will not be discussed in this paper.

6. Discussion

6.1. Correlations between Observables and Parameters

Correlations that measure how dependent two variables are on each other play a significant role in mapping between the domain (input) and codomain (output). The accuracy of machine learning can be proportional to the degree of correlation between input and output features. For example, the less correlated they are, the larger the variances are and vice versa. In this work, we have already encountered several trails of the correlation between parameters and observables in the inference results. One of the most evident examples is that the SFRD requires $\mathcal{O}(10^4)$ simulations to reach stable convergence, whereas only $\mathcal{O}(10^3)$ simulations are sufficient for the SMFs (see Sections 4.1.1 and 5.1 for the SFRD and the SMFs, respectively). In other words, with the same size of training data, i.e., the same number of emulations, the SMF

would have attained a more accurate posterior with a smaller variance.

For a quantitative analysis, we measure the correlation between parameters and observables using mutual information (MI), which is a fundamental measure of the interdependence or relationship between two variables. In contrast to linear correlation coefficients, such as the Pearson correlation coefficient, MI captures nonlinear statistical dependencies (Kinney & Atwal 2014). MI is defined by

$$I(X, Y) = D_{\text{KL}}(P_{XY} \parallel P_X \otimes P_Y), \quad (10)$$

where $D_{\text{KL}}(\cdot \parallel \cdot)$ is the Kullback–Leibler divergence²¹, and P_{XY} and P_X are the joint and marginal distributions, respectively. Here, \otimes denotes the outer product that spans the probability space from X and Y to (X, Y) . $I(X, Y)$ quantifies the statistical distance between the joint probability and the product of marginals. The MI is zero if and only if X and Y are independent.

However, the estimation of MI is challenging and only tractable for discrete variables or when probability distributions are known (Paninski 2003). Thus, we adopt the MI regression function in the `sklearn` package (Kraskov et al. 2004; Pedregosa et al. 2011; Ross 2014). Using the package, we estimate the MI for (1) simulated observables and parameters from the LH set, and (2) emulated observables and parameters from the LH set. Here, the observables and the parameters are normalized to reduce the effects of the difference in magnitude of values in the same way that is used to train the emulator (Section 3.3).

²¹ The Kullback–Leibler divergence is defined as $D_{\text{KL}}(P \parallel Q) \equiv \sum_{x \in \mathcal{X}} P(x) \log(P(x)/Q(x))$ or $\int_{-\infty}^{\infty} p(x) \log(p(x)/q(x)) dx$. This is also referred to as the statistical distance between $Q(x)$ and $P(x)$.

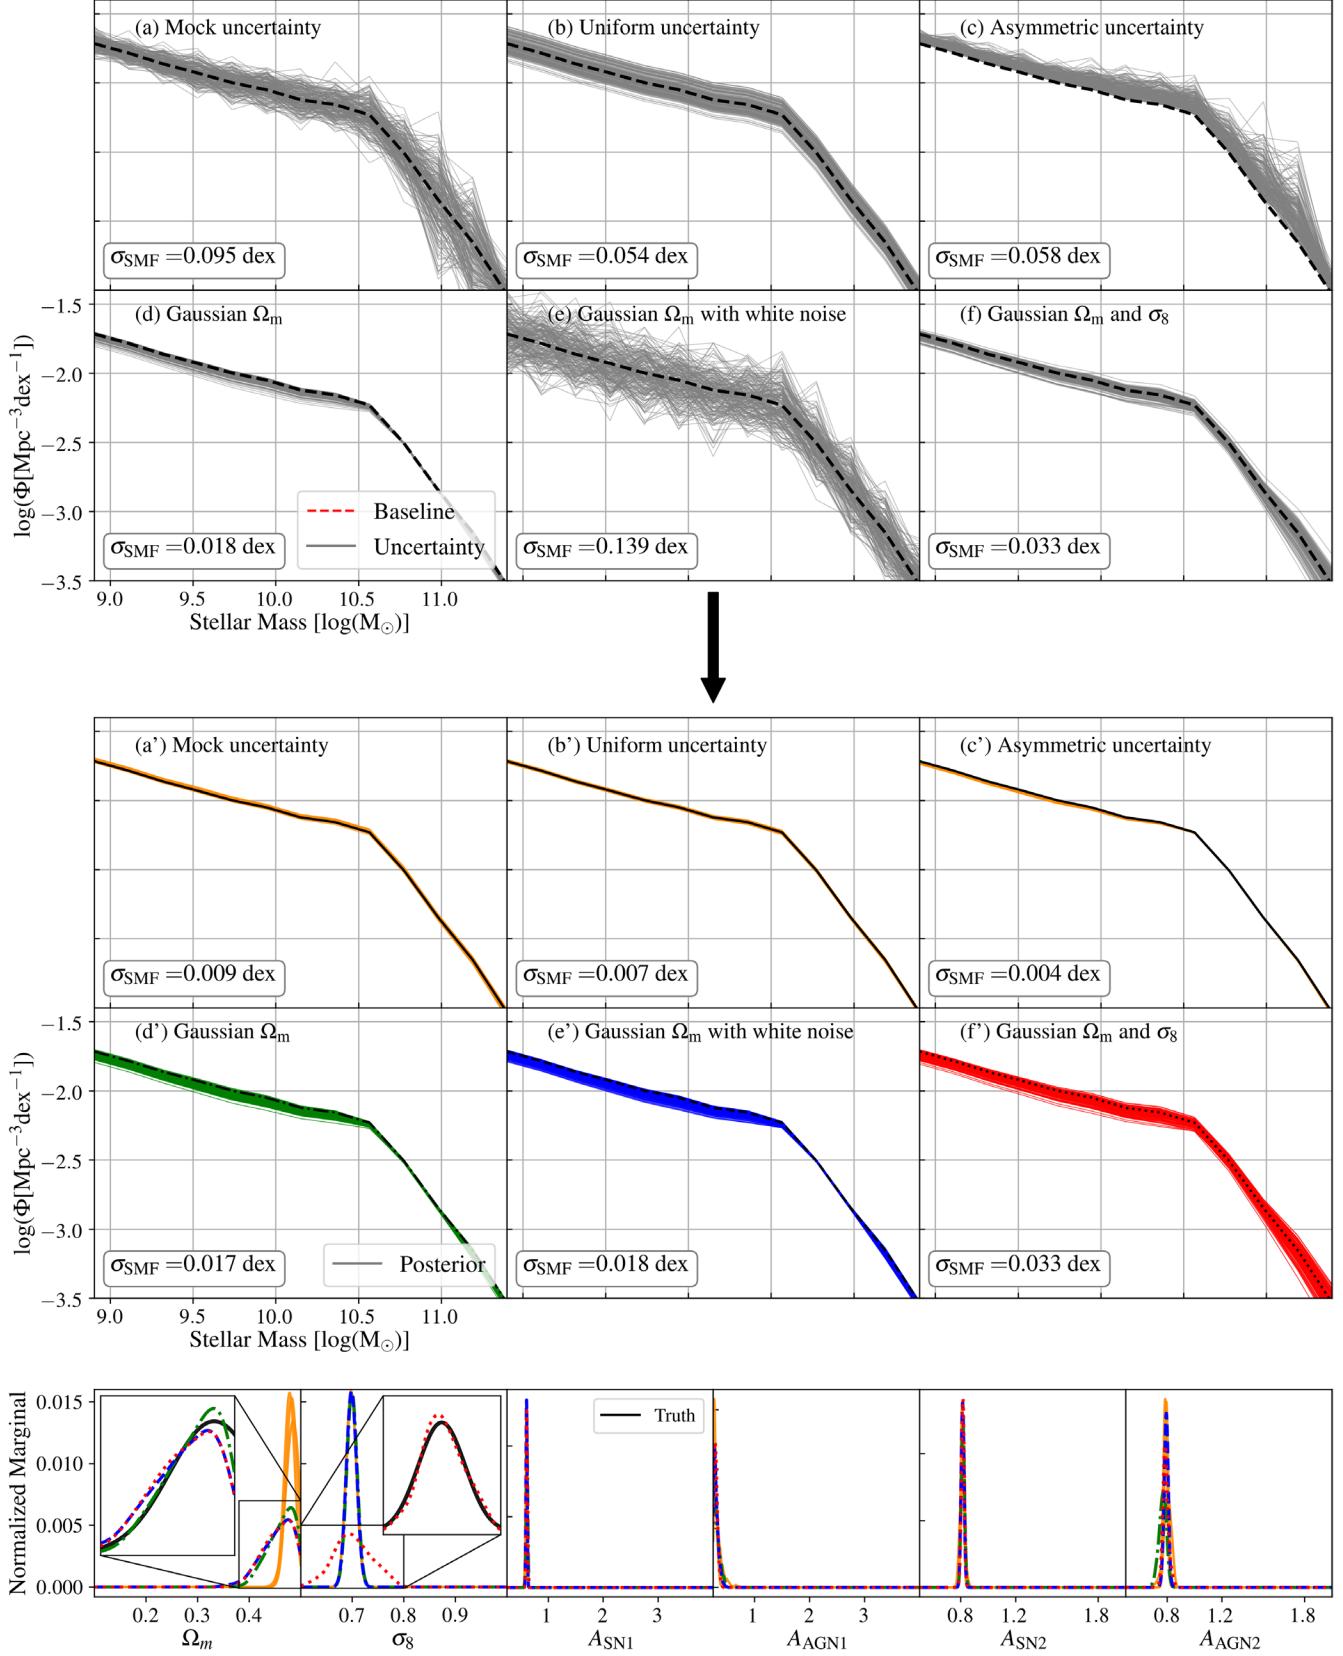


Figure 16. The upper six panels show the SMFs at $z = 0.5$ with six different types of uncertainties. The middle six panels show the inferred SMFs at $z = 0.5$ with six different types of uncertainties. The lower six panels show six marginal distributions of the inferred posteriors from the SMFs with six different types of uncertainties. The marginals are color coded as follows: mock uncertainty, uniform, asymmetric white noise (all three shown in orange), Gaussian Ω_m (green), Gaussian Ω_m with white noise (blue), and Gaussian Ω_m and σ_8 (red). In the zoomed-in panels, the Gaussian distributions implemented in uncertainties (d), (e), and (f) are shown as black solid lines.

Table 4 shows the estimated MI of each observable-parameter pair. In the simulations, the MI of the SMFs (*SimSMF*) is significantly higher than that of the SFRD (*SimSFR*) over all parameters except $A_{\text{AGN}2}$. Interestingly, the gap is drastically reduced in the emulators (*EmuSMF* and *EmuSFR*): the MI goes up for both, but more so for the SFRD, resulting in the narrowing of the gap. The increase of MI in the emulators (especially in the SFRD) can be attributed to the training of the emulators. During training, the emulator (neural network) naturally suppresses noise in the data, making the correlations stronger (Goldfeld et al. 2018; Gabrié et al. 2019). As a result, input and *emulated* output pairs can attain overall more correlation via the emulator than input and *simulated* output pairs. Note that the correlation of the *emulated* pairs should depend on the precision of training. Nonetheless, the relative magnitudes of MI among the six parameter-observable pairs remain the same in terms of order. For instance, the MI of the SMF-parameter pairs is still higher than that of the SFRD-parameter pairs except for $A_{\text{SN}2}$ and $A_{\text{AGN}2}$.

In addition to the gap between simulations and the emulator, the MI among the parameters shows considerable differences in magnitude. $A_{\text{SN}1}$ is found to be the most relevant parameter to both SFRD and SMFs, whereas the AGN feedback parameters ($A_{\text{AGN}1}$ and $A_{\text{AGN}2}$) have MI that is 1 or 2 orders of magnitude lower than that of the other parameters. That is, $A_{\text{SN}1}$ can precisely be inferred or estimated in terms of inference and machine learning. On the other hand, not only are the AGN parameters hardly constrained but also the AGN physics themselves have a negligible impact on the SFRD and the SMFs in the TNG suite of CAMELS simulations that have low resolution and limited volume (refer to Section 6.1.1 for the SMFs and to Figure 9 of Villaescusa-Navarro et al. 2021 for the SFRD). Such impact of relative magnitudes among parameter-SFRD pairs can also be found in Villaescusa-Navarro et al. (2021, Figure 9), which demonstrates that the SFRD is only sensitive to Ω_m , σ_8 , $A_{\text{SN}1}$, and $A_{\text{SN}2}$, and not to $A_{\text{AGN}1}$ or $A_{\text{AGN}2}$.

The values of the MI are also in line with the performance of inference with respect to the parameters. For instance, the MI between $A_{\text{SN}2}$ and the SFRD, which is an exceptional case, is larger than that between $A_{\text{SN}2}$ and the SMFs, which is in line with the relative error of $A_{\text{SN}2}$ in the SFRD (fifth row) being also smaller than that of the SMFs (sixth row). The ratio among the MI of each parameter is inversely proportional to that of the relative error and/or the variance for all inferences presented in this work. In addition, convergence depends on the degree of correlation of the pairs, implying that more training data—i.e., simulations—are required to converge in the SFRD (compare Figures 3 and 9). Lastly, this is also in agreement with previous work. The CAMELS introduction paper (Villaescusa-Navarro et al. 2021) builds a fully connected neural network to predict cosmological and astrophysical parameters taking the SFRD as input. With the neural network, Ω_m , σ_8 , $A_{\text{SN}1}$, and $A_{\text{SN}2}$ are predicted relatively well compared to $A_{\text{AGN}1}$ and $A_{\text{AGN}2}$ (Villaescusa-Navarro et al. 2021).

Figure 13 shows the MI between each parameter and the SMF at $z = 0.5$ (solid) and 2.5 (dotted) as a function of the stellar mass of each bin. At $z = 0.5$, $A_{\text{SN}1}$ is dominant over other parameters especially for $M_\star \lesssim 10^{10.8} M_\odot$, while Ω_m also has appreciable MI. In the higher-mass end ($M_\star \gtrsim 10^{10.8} M_\odot$), $A_{\text{SN}2}$ is the most effective, followed by σ_8 and $A_{\text{SN}1}$, whereas the MI of $A_{\text{AGN}1}$ and that of $A_{\text{AGN}2}$ are negligible across the entire stellar mass range. In the case of the higher redshift ($z = 2.5$),

the order of the relative magnitudes of the MI among the parameters changes. $A_{\text{SN}1}$ is still the most dominant parameter in the low-mass end, but the gap with Ω_m is smaller. The MI of σ_8 is drastically larger overall and dominates the high-mass end, whereas $A_{\text{SN}2}$ becomes negligible. Compared to those at $z = 0.5$ (solid), Ω_m (red dotted) remains approximately the same and σ_8 (yellow dotted) is significantly increased, while $A_{\text{SN}1}$ (green dotted) and $A_{\text{SN}2}$ (blue dotted) are appreciably decreased, which, overall, leads to an increase in the ratio of the MI of the cosmological parameters to that of the astrophysical parameters. This can be indicative that cosmology is more influential at early times than at late times compared to astrophysical phenomena.

The MI can be an indirect but crucial measure of the degree of relative impact of parameters on observables. We can extract physical insights from Figure 13 as follows: (1) The impact of cosmology (Ω_m and σ_8) diminishes at lower redshift in terms of the SMF. (2) The energy budget of supernova-driven winds ($A_{\text{SN}1}$) is of paramount importance in galaxies of stellar mass $\lesssim 10^{10.8} M_\odot$. (3) For massive galaxies ($M_\star \gtrsim 10^{10.8} M_\odot$), the stellar wind's velocity ($A_{\text{SN}2}$) has more effect than its energy budget ($A_{\text{SN}1}$). (4) The portion of dark matter (Ω_m) in the universe has more impact on lower-stellar-mass galaxies than on higher-mass galaxies, (5) whereas density fluctuation (σ_8) is more significant in massive galaxies. (6) The kinetic feedback of black holes ($A_{\text{AGN}1}$ and $A_{\text{AGN}2}$) is weakly related to the SMF regardless of both the redshift and stellar mass of the galaxies. In the following section, the impact of the AGN parameters on the SMFs is discussed further.

6.1.1. Negligible Impact of AGN Kinetic Feedback on Galactic Stellar Mass

We have seen the weak correlation between the observables and AGN feedback parameters from both the variances of the posteriors and the MI. In this section, we directly investigate the impact of AGN parameters on the SMFs, compared with that of other parameters. The AGN parameters control the energy budget of the low-accretion mode (kinetic feedback) of black holes ($A_{\text{AGN}1}$) and its burstiness ($A_{\text{AGN}2}$) (refer to Section 2). The kinetic wind from black holes can play a more crucial role in quenching massive galaxies than thermal feedback can, since the kinetically injected energy is less vulnerable to gas (over)cooling, especially for dense gas, for which the cooling time is short (Bower et al. 2006; Croton et al. 2006; Fabian 2012; Dubois et al. 2013; Rosas-Guevara et al. 2015; Anglés-Alcázar et al. 2017; Zinger et al. 2020; Piotrowska et al. 2022; Wellons et al. 2022). Furthermore, Terrazas et al. (2020) found that in the TNG model, the galaxy can be quenched whenever the cumulative kinetic feedback energy of the central black hole exceeds the gravitational binding energy of the gas within the galaxy. In this regard, the lack of correlation between the AGN parameters and the SMFs in CAMELS-TNG²², even at the higher-mass end, seems puzzling.

Shown in Figure 14 are the SMFs at $z = 0.5$ from the simulations in the 1P set and the SMFs from the simulations with no kinetic feedback ($A_{\text{AGN}1} = 0 = A_{\text{AGN}2}$, black dashed) and the fiducial parameter combination (black solid). Note that the 1P set consists of 61 simulations varying only one

²² Note that CAMELS-SIMBA shows more variations with respect to the AGN parameters.

parameter at a time (Villaescusa-Navarro et al. 2021, Section 3.3.2) and all the simulations presented in this section are performed with the same initial conditions. In each panel, the bluer the SMF, the higher the parameter value, and the redder the SMF, the lower the parameter value. In comparison to the simulation uncertainty (light gray region), four parameters (Ω_m , σ_8 , $A_{\text{SN}1}$, and $A_{\text{SN}2}$) excluding the AGN parameters show significant changes of the SMFs in response to variations of their values, which is in line with the previous section. One interesting feature is that the impact of $A_{\text{SN}2}$ is applied conversely to the low-mass end and the high-mass end of the SMFs. For instance, an increase in the supernova-driven wind speed ($A_{\text{SN}2}$) effectively suppresses the massive population while keeping the lower-mass population similar or increasing it.

In contrast, the changes in the AGN parameters (right panels) have a negligible impact on the SMFs. However, this does not imply that AGN kinetic feedback lacks any effect at all, since even the weakest AGN feedback in the 1P set (the reddest line in the right panels) suppresses the massive population compared to our new simulation that lacks kinetic feedback altogether ($A_{\text{AGN}1} = 0 = A_{\text{AGN}2}$)²³, as expected in many papers (Page et al. 2012; Dubois et al. 2013; Terrazas et al. 2020; Su et al. 2021). In addition, Figure 8 of Pillepich et al. (2018a) shows the clear impact of black hole kinetic feedback on the stellar masses of massive galaxies. The negligible effects of the AGN parameters in the CAMELS simulations can be explained by a combination of the following: (1) Due to the small box size $\sim (25 \text{ Mpc } h^{-1})^3$, the CAMELS simulations usually lack massive galaxies subject to AGN kinetic feedback. (2) The impact of the weakest AGN feedback parameters of the CAMELS simulations is already so effective that further increases of these parameters cannot lead to significant changes in the SMFs.

6.2. Degeneracy Broken with SMFs

We have seen the bimodality of the inferred posterior density and degeneracy in SFRD in Section 4.1.2 and Figures 4 and 5. In this section, we discuss how degeneracy in the SFRD can be broken with both the emulated SMF and the simulated SMF. Figure 15 illustrates the marginal distributions (bottom) that are exactly the same marginals as those in Figure 5, derived by inference on the SFRD, and the corresponding *emulated* SMFs (top left) and the *simulated* SMFs (top right). In contrast to the case of the SFRD, where the two *emulated* SFRDs from the two peaks of the posteriors (red dashed and blue dotted) are nearly on top of each other, the corresponding *emulated* SMFs for the two separate sets of parameters are clearly different, as shown in panel (a). The standard deviation of the SMFs (gray) in panel (a), which is 0.081, is an order of magnitude higher than that of the SMFs from the convergent posterior in Figure 8, which is only 0.007. In addition, the relative error of the SMFs from peak 2 (blue dotted) with respect to the SMFs from peak 1 (red dashed) is 59%, which is far higher than the 0.79% from the SFRDs in Figure 5. Therefore, we can conclude that the SMFs can break the degeneracy of the SFRDs in terms of the emulators.

Also, we perform new simulations with parameter combinations of peak 1 (red dashed) and peak 2 (blue dotted) to study whether the SMFs can break the degeneracy even under the

simulation uncertainty. In Figure 15(b), the red dashed and blue dotted lines indicate the simulated SMFs from the two peaks. Similar to the emulated SMFs, the two simulated SMFs do not coincide with each other. However, since the simulation is affected by simulation uncertainty, such as cosmic variance and butterfly effects (refer to Section 2.5), we include the confidence regions (gray in panel (b)) of the simulation uncertainty. We use $\pm 1\sigma$ (gray) and $\pm 3\sigma$ (light gray) regions that correspond to 68.1% and 99.7% confidence levels for the Gaussian distribution, respectively. The standard deviations of the SMFs are directly calculated from the simulations in the CV set as a function of stellar mass. The confidence regions are drawn with respect to the red dashed curves. Panel (b) (top right) demonstrates that the low-mass ends of the blue dotted SMFs noticeably fall outside the 3σ regions at $z \geq 1.5$. By the definition of 3σ regions as corresponding to a 99.7% confidence level, the blue dotted SMFs have *only a 0.3%* chance that they share the same origin of the red dashed SMFs.

Both the emulated and simulated results support two main conclusions: (1) there exist degeneracies in the SFRD (refer to Section 4.1.2), and (2) the SMFs can break the degeneracies in the SFRD (Figure 15). Taken together with the correlation analysis in Section 6.1, these conclusions further support the notion that the higher the correlation between an observable and other parameters, the more strongly the observable constrains the parameters; here, the SMF is shown to be a stronger constraint than the SFRD.

6.3. Uncertainty in ILI

We have briefly discussed that the mock uncertainty applied to the SMFs has negligible impact on the variance of posteriors in Section 5.1.1 unlike the case of the SFRD in Section 4.1.3. To elucidate the origin and implications of these results, and what type of uncertainty should be adopted for ILI, here we perform tests of ILI using various types of uncertainties. We adopt six different uncertainties as follows: (a) *Mock uncertainty* (that which we have used throughout this paper) is modeled to mimic the simulation uncertainty from the CV set in Appendix B. (b) *Uniform uncertainty* is made of a univariate Gaussian distribution. Random variables are drawn from the Gaussian and added to the SMF uniformly with respect to the stellar mass, leading to overall shifts in normalization. (c) *Asymmetric uncertainty* is simply modeled as the modulus of the mock uncertainty such that the uncertainty only goes in the positive direction. (d) For *Gaussian Ω_m uncertainty*, given a set of parameters (Ω_m^0 , σ_8^0 , $A_{\text{SN}1}^0$, $A_{\text{AGN}1}^0$, $A_{\text{SN}2}^0$, $A_{\text{AGN}2}^0$), we generate the SMFs out of (Ω_m' , σ_8' , $A_{\text{SN}1}'$, $A_{\text{AGN}1}'$, $A_{\text{SN}2}'$, $A_{\text{AGN}2}'$), where $\Omega_m' \sim \mathcal{N}(\Omega_m^0, \sigma_{\Omega_m})$. Here, σ_{Ω_m} is set to 0.04. (e) We model Gaussian Ω_m uncertainty with white noise by adding white Gaussian noise directly to the SMFs with Gaussian Ω_m uncertainty. (f) *Gaussian Ω_m and σ_8 uncertainty* is modeled similarly to *Gaussian Ω_m uncertainty* except that we additionally vary σ_8 as $\sigma_8' \sim \mathcal{N}(\sigma_8^0, \sigma_{\sigma_8})$ with σ_{σ_8} of 0.04. A visual description of the above six uncertainties is shown in the top panels of Figure 16.

We perform ILI from the emulated SMF at $z = 0.5$ that is generated with the same parameters that are used in Section 5.1, together with the six uncertainties. Figure 16 shows the marginals of the inferred posteriors (bottom) and the corresponding SMFs (middle). The corresponding posterior (middle) inferred from each implemented uncertainty (top) ((a)–(f)) is shown in panels (a')–(f'), respectively. Also, the

²³ The new simulation does include thermal AGN feedback.

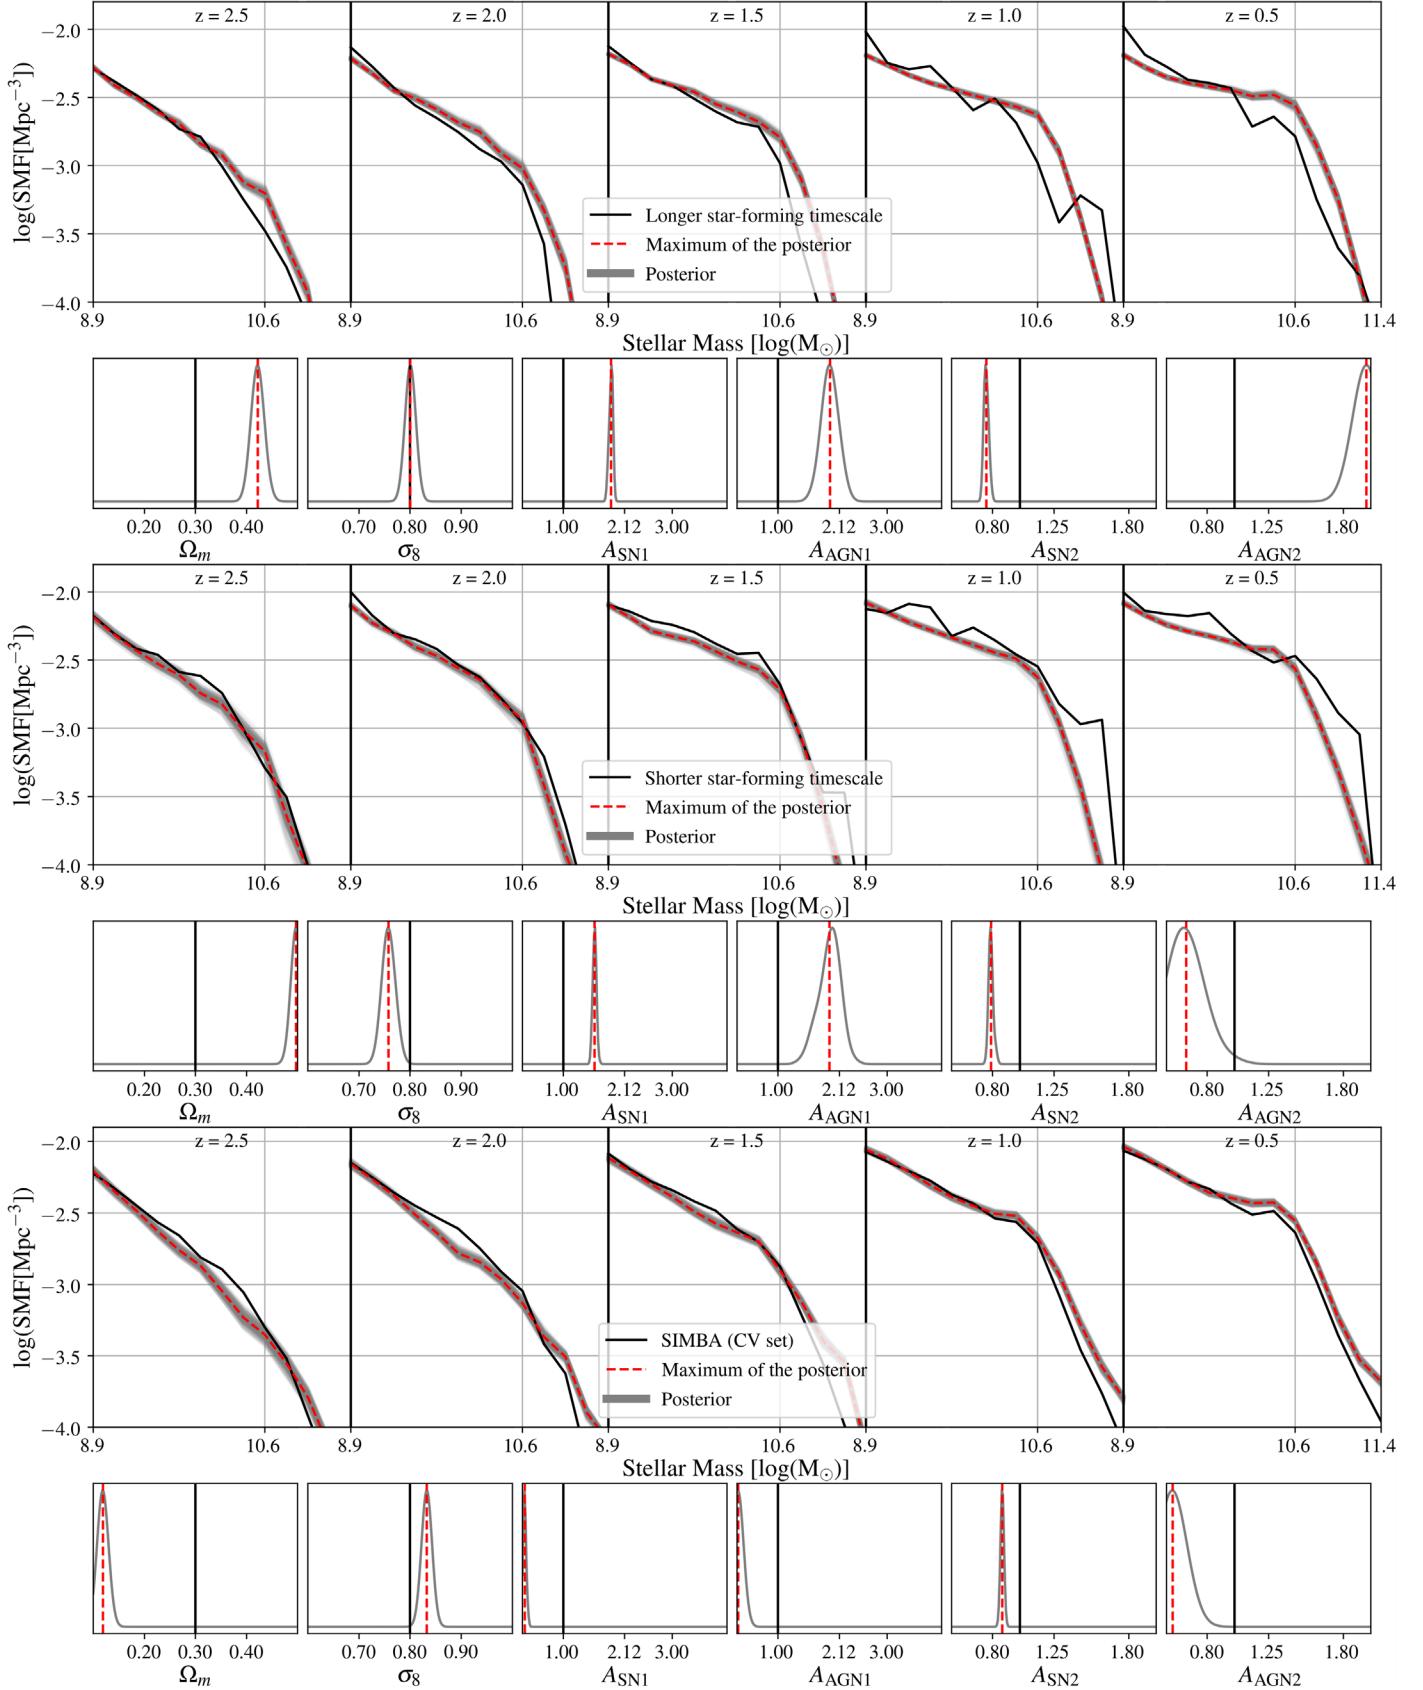


Figure 17. Top: SMFs and the marginal posterior densities that are inferred from the SMFs with the six fiducial parameters and a longer star formation timescale. Middle: Same, but with a shorter star formation timescale. Bottom: SMFs and the marginal posterior densities that are inferred from the average SMFs out of the CV set of the SIMBA, rather than the IllustrisTNG, suite.

Table 5
The Standard Deviations of the Inferred Marginal Distributions in Figure 16

	σ_{Ω_m}	σ_{σ_8}	$\sigma_{A_{\text{SN1}}}$	$\sigma_{A_{\text{AGN1}}}$	$\sigma_{A_{\text{SN2}}}$	$\sigma_{A_{\text{AGN2}}}$
a'	0.005	0.001	0.006	0.084	0.006	0.010
b'	0.007	0.002	0.012	0.056	0.008	0.023
c'	0.002	0.001	0.006	0.027	0.004	0.011
d'	0.025	0.001	0.001	0.055	0.006	0.027
e'	0.030	0.001	0.005	0.042	0.004	0.012
f'	0.028	0.036	0.007	0.030	0.003	0.019

inferred SMFs (middle) and marginal densities (bottom) are consistently color coded as follows: orange solid (a')–(c'), green dotted–dashed (d'), blue dashed (e'), and red dotted (f'). We color (a')–(c') with the same color since the implemented uncertainties (a)–(c) share the same property of the uncertainty injected to the SMFs being *uncorrelated* with the parameters.

In the case of (a)–(c), the standard deviations of the inferred SMFs ((a')–(c')) (first row of the middle panel of Figure 16) are an order of magnitude smaller than those of the implemented uncertainty (first row of the top panel). The marginal distributions ((a')–(c')) (orange) are nearly on top of each other (see Table 5). Most importantly, the standard deviations of both the SMFs and the marginals are close to those of the inference from the SMFs *without uncertainty* in Figure 8 of Section 5.1, which are $\bar{\sigma}_{\text{smf}} = 0.007$ dex and $\sigma_{(\Omega_m, \sigma_8, A_{\text{SN1}}, A_{\text{AGN1}}, A_{\text{SN2}}, A_{\text{AGN2}})} = (0.006, 0.002, 0.024, 0.076, 0.006, 0.018)$. This leads us to the proposition that *uncorrelated uncertainty* has essentially no impact on inference.

In the case of the uncertainties directly injected to Ω_m and σ_8 , the inferred SMFs in (d') and (f') and the implemented uncertainties in (d) and (f) have approximately the same standard deviations (see Figure 16). On the other hand, the standard deviations of the inferred SMFs in (e') are completely different from those of the implemented uncertainty in (e), since the white noise that is uncorrelated with SMFs in (e) is canceled out during the inference. Shown in the zoomed-in panels are the analytic Gaussian distributions (black solid) that are used to generate the uncertainties ((d)–(f)). In Ω_m , the three inferred marginals—green dotted–dashed, blue dashed, and red dotted—are in relatively good agreement with the analytic lines (black solid). Note that the marginal (e') is not affected by the white noise. Due to the range limit of Ω_m , the inferred marginals are skewed such that the probability density near $\Omega_m = 0.5$ drops sharply as compared to the analytic line, whereas in the σ_8 panel, the marginal (f') (red dotted) precisely matches the analytic line (black solid) (see Table 5).

The implemented uncertainties ((d)–(f)) are *correlated* since the uncertainty is directly injected to the parameter such that there exists θ' satisfying $f(\theta') = f(\theta + Z(\eta))$ for $\theta' \in \mathbb{R}$. However, the uncertainty (e) is only a *partially correlated uncertainty* since it includes the *uncorrelated* part as well. In contrast to that of the *uncorrelated* uncertainties ((a)–(c)), the variance of the *correlated* uncertainties ((d) and (f)) can successfully be captured in the posteriors ((d') and (f')). In the case of the *partially correlated* uncertainty (e), the uncorrelated part is canceled out, leaving only the variance of the *correlated uncertainty* in the posterior distribution. This is strong evidence that *uncorrelated uncertainty* has negligible impact on ILI and

highlights the importance of a well-established mock uncertainty that can reproduce the correlation between the simulation uncertainty and the parameters. From another point of view, if the simulation uncertainty is not correlated with the parameters, performing ILI using cosmological simulations can naturally eliminate the effects of the simulation uncertainty. Note that it is not yet proven that the simulation uncertainty is correlated with the parameters.

6.4. Physical Limitations of Simulations

6.4.1. Mismatch between Inferred SMFs and Observed SMFs

We have seen the significant mismatch between the inferred SMFs and the observed SMFs in Section 5.2. There are four primary issues in the SMFs inferred from the observed SMFs shown in Figure 11: (1) The population of massive galaxies ($M_\star \gtrsim 10^{10} M_\odot$) in the inferred SMFs are located far below the observed SMFs across all redshifts. (2) The inferred SMFs have a “shoulder” at $M_\star = [10^{10.5}, 10^{11}] M_\odot$, which does not appear in the observations. (3) The evolution of the emulated SMFs with respect to redshift $d\phi(M, z)/dz$ in the low-mass end is smaller than that of the observed SMFs. (4) The SMF–stellar mass slope changes in time are appreciably different between simulations and observations.

The mismatch in the high-mass end has been a crucial problem in the inference from the five concatenated observed SMFs in Section 5.2.1. The difference between the inferred SMFs and the observed SMFs is approximately 0.3 dex on average (or a relative error of 41%). Concerning physical models in cosmological simulations, there can be two physical mechanisms that can control the populations of massive galaxies: (1) large-scale structure formation and evolution and (2) astrophysical feedback. In the context of structure formation, a sufficient amount of matter (Ω_m) and large density fluctuations (σ_8) can lead to a high abundance of massive halos and galaxies. On the other hand, it is well known that AGN feedback is a dominant factor in quenching star formation in massive galaxies (for a seminal review, see Fabian 2012; Somerville & Dave 2015), whereas stellar feedback is more effective in low-mass galaxies. However, since AGN feedback is ineffective as we have seen in Section 6.1.1, the dependence of the massive-galaxy population on cosmology should considerably increase. The compensation of cosmology can be found in Section 5.2.2. In short, the inferences from the individual SMFs in Figure 12 have recovered the population of massive galaxies by having significantly larger σ_8 compared to the posteriors inferred from the concatenated SMFs in Figure 11. That being said, the inferences cannot find a set of parameters that can reproduce the evolution of SMFs as seen in the observations.

In addition, Figure 12 exhibits several notably different properties for observations and emulators. The emulated SMFs tend to have a “shoulder” at $M_\star = [10^{10.5}, 10^{11}] M_\odot$, which does not appear in the observed SMFs at all. Although we can find this property in the actual hydrodynamic simulation, the rescaling might enhance such properties (refer to Section 2.4 and Appendix A). Figure 18 demonstrates that the rescaled SMFs (red solid) have more notable shoulders than the SMFs without rescaling (red dashed). Thus, such shoulders can be attributed to the rescaling and/or resolution effects.

The rate of evolution of the SMFs in cosmological simulations is relatively low compared to the rate of evolution

in observations. In addition, the slope of the SMFs with respect to stellar mass at the low-mass end hardly changes across redshift in the observations, whereas the slope in the simulations tends to decrease as time goes by. Such discrepancies between simulations and observations can be evidence for the limits of physical models in hydrodynamic simulations and/or the limited dimensions of parameter space. We anticipate that this problem can be at least partially alleviated by introducing additional dimensions in the parameter space.

6.4.2. An Extra Parameter: Star Formation Timescale

In this section, we investigate whether the parameters can be successfully inferred from a simulated observable that comes from a higher-dimensional parameter space that has an *extra* dimension. Thus far, we have performed ILI on six parameters from observables—SFRD or SMFs—that are obtained either by a specific set of six parameters, or from actual observations. Here, we perform ILI on the same six parameters from SMFs that are simulated from the six fiducial parameters while varying one additional subgrid model parameter, the star formation timescale. Two simulations are performed with longer (twice the fiducial) and shorter (half of the fiducial) star formation timescales than the fiducial run (hereafter, LT and ST denote the longer and shorter star formation timescale simulations, respectively). In LT and ST, the same random seed as that in the fiducial run is adopted to minimize the effect of simulation uncertainty through cosmic variance. Lastly, to approximate the emulator counterparts of the SMFs for LT/ST, we calculate the ratios of the SMFs bin by bin from LT/ST to the simulated fiducial SMFs, and then apply them to the emulated SMFs of the fiducial parameters to obtain emulated SMFs for LT/ST. Here, we assume that changes of the star formation timescale do not lead to any significant changes in other physical properties and observables except the SMFs.

Figure 17 shows the marginal posterior densities and SMFs inferred from the emulated SMFs from LT and ST (top and middle panels, respectively). In the case of LT, the peaks of the marginal densities of the inferred posteriors (red dashed) considerably deviate from the fiducial parameters (black solid), except for σ_8 . The low-mass end of the inferred SMFs (red dashed) at $z = 2.5, 2.0$, and 1.5 is in relatively good agreement with the SMFs from LT (black solid), whereas at $z = 1.0$ and 0.5 , the inferred SMFs and the SMFs of LT show a complete mismatch. In addition, LT has a high evolutionary rate of the SMFs and tends to have steeper slopes as redshift decreases, which the inferred SMFs fail to match as in the inference from observations (see Figure 11). Similar to those in LT, the inferred SMFs are in good agreement with the SMFs of ST at $z = 2.5, 2.0$, and 1.5 , whereas the high-mass end of the inferred SMFs at low redshifts is clearly lacking as compared to the SMFs of ST, which we have also seen in the inference from the observation in Figure 11.

We have found that the SMFs from the seven-dimensional parameter space including an extra parameter—star formation timescale—are not necessarily reproduced by points in the conventional six-dimensional parameter space that we have used throughout this work. LT and ST have intrinsic properties that cannot be reproduced from points in the six-dimensional parameter space (e.g., higher rates of evolution in the SMFs, slope evolution in the SMFs, and large populations of massive

galaxies). The results indicate that the introduction of extra dimensions or parameters has the potential to resolve the problem of significant mismatch between the five concatenated observed SMFs and observations.

6.4.3. Inference from SIMBA

In addition to the ILI with an extra parameter, we perform ILI from the SMFs that are obtained by the fiducial simulation from the CV set of the SIMBA suite (refer to Section 2 or Villaescusa-Navarro et al. 2021). Likewise, we rescale the SIMBA SMFs to obtain emulated SIMBA SMFs in the same way as we have in Section 6.4.2, assuming that the same rescaling relation is applicable. Note that although we make inference from a SIMBA-derived observable, we use the same emulators that we have used so far, which are trained on the TNG suite of the CAMELS simulations.

The bottom row of Figure 17 illustrates the inferred SMFs (red dashed and gray) and the emulated SIMBA SMFs (black solid). The inferred SMFs from the maximum of the posteriors (red dashed) seem to match the SIMBA SMFs relatively well even though the relative error is $\sim 21\%$. Unlike those of LT and ST in Section 6.4.2, the inferred SMFs reasonably follow the trends of the SMFs, such as the slopes and rates of evolution. Still, the difference between the inferred parameters and fiducial parameters has not been narrowed. Such deviations between the inferred and fiducial parameters can be attributed to differences of the physical models between TNG and SIMBA.

6.5. Caveats and Physical Interpretation of Inference

This section discusses how the inferred parameters can be interpreted in an astrophysical sense. We defer the physical interpretation of the inference to the last section because the inferred parameters contain very little meaningful physics at this time. Prior to the physical interpretation, it is imperative to understand the key factors that influence inferences on parameters, e.g., the emulator, simulation uncertainty, resolution convergence, and a limited parameter space.

First, we have employed the emulators for computational efficiency, paying the price of discrepancies between the simulations and emulators. The discrepancies inevitably propagate to the inferred posterior and lead to deviations from the posterior that would be inferred if actual cosmological simulations were employed in the inference. Assuming that we can replace the emulators with actual cosmological simulations, the next question shall be, “Will the posterior inferred from the actual cosmological simulations contain robust information on physics?” Or, more specifically, “Will the cosmological parameters inferred from the cosmological simulations, not the emulators, be physical and comparable to preexisting estimations?” To answer these questions, resolution convergence and the limited dimensions of parameter space, along with simulation uncertainty, should be taken into consideration.

6.5.1. Resolution Effect

Resolution effects in hydrodynamic simulations have been discussed in many papers (Lia et al. 2000, 2002; Ceverino & Klypin 2008; Nagamine 2010; Hubber et al. 2013; Anglés-Alcázar et al. 2014; Regan et al. 2014; Pillepich et al. 2018a, 2018b; Snaith et al. 2018; also see Sections 2.4 and

Appendix A in this paper). In general, it is believed that hydrodynamic simulations are sensitive to (spatial and mass) resolution. On the other hand, pure hydrodynamic simulations themselves without subgrid models, such as radiative cooling, are highly likely to be convergent at some level (refer to Hubber et al. 2013 for the resolution convergence of both smoothed particle hydrodynamics and adaptive mesh refinement codes). This implicitly indicates that the resolution effect can likely be attributed to the response of subgrid models to different resolutions. In addition, cosmological hydrodynamic simulations with subgrid models can become convergent once they reach a certain resolution and beyond (Lia et al. 2000, 2002; Ceverino & Klypin 2008; Hopkins et al. 2018). The above arguments suggest that resolution effects can generally be incorporated into subgrid models—as astrophysical parameters. This is often called weak convergence, which requires recalibration for each resolution of simulations, whereas strong convergence requires that the simulation results not change once the subgrid models are fixed (Schaye et al. 2015).

We turn our attention to the resolution effect in this work. As discussed in Section 2.4, we rescale the observables—SFRD and SMF—to minimize the resolution effect according to Pillepich et al. (2018b). However, aside from its inaccuracy, there is another problem with rescaling. We find that the resolution effect has a significant dependence on the position in parameter space, which is in line with the above arguments. In principle, the rescaling must be a function of the parameters θ , whereas we construct the rescaling function only based on the fiducial parameters in this work. Due to computational cost, it is not realistically possible to obtain the full rescaling relation as a function of the parameters by running high-resolution simulations over the entire parameter space (Ho et al. 2022). Hence, the inference results near the fiducial parameters are physically comparable to those of TNG100-1, whose resolution is the target resolution of the rescaling, but other regions of the parameter space may be difficult to grasp.

6.5.2. Limited Parameter Space

In this work, we have seen the limits of the six-dimensional model parameter space, e.g., weak correlation between the AGN parameters and observables in Section 6.1.1, failure of inference from the observed SMF in Section 6.4.1, and the potential of the extended parameter space in Section 6.4.2. The limits of the parameter space can affect not only the performance or accuracy of inference but also the physical interpretation of the inferred posterior. In addition, the limitations of the subgrid model and the limited dimensions of the parameter space can lead to overfitting or over-fine-tuning of the parameters to the target observable, even if the inferred physics is not realistic. For instance, since our emulators (or simulations) are insensitive to AGN feedback, the cosmological and stellar feedback parameters should be adjusted to control the population of the high-mass end in Section 5.2.2. The parameter space of the TNG universe or the SIMBA universe is likely to be insufficient to describe the real universe. Thus, we will ultimately move toward extending the parameter space beyond the current TNG universe so that we can resolve the problems of inference failure and secure the physical meaning of the posteriors.

6.5.3. Physics of Inferred Parameters

The posterior distributions inferred from the observations have shown appreciable discrepancy with what researchers usually expect, especially in cosmological parameters. The discrepancies mostly stem from the physical limitations of the simulations on which the emulators are built, but not from the inference procedure itself. Considering that inference is a procedure that finds the best combinations of parameters that can match a target observable, we can find some problems in (the diversity of) the “combinations,” but not in the “finding.” For instance, in Section 6.4.2, we directly observe how the limits of parameter space or an additional dimension can have an impact on the inference. Also, the lack of physics in one part can lead to an (unwanted) compensatory reaction to another part, in Section 6.4.1. Hence, the inferred parameters in this work are largely affected by intrinsic features of the imperfection of simulations such as a limited parameter space volume, resolution effects, subgrid models, and the simulation box size.

7. Summary

In order to calibrate cosmological simulations against observations, we have employed ILI that enables rigorous Bayesian inference in a computationally efficient way by adopting NDEs that evaluate the likelihood instead of the explicit analytic likelihood used in conventional Bayesian inference (Section 3). In addition, for computational efficiency, we have adopted emulators that are trained on ~ 1000 cosmological simulations from the CAMELS project (specifically, those based on the IllustrisTNG framework) to predict simulated observables, taking as input the cosmological and astrophysical parameters, and used these emulators as surrogates for the cosmological simulations (Section 3.3). Using the emulators, we have conducted ILI on the cosmological and astrophysical parameters (Ω_m , σ_8 , stellar wind feedback, and kinetic black hole feedback) from the cosmic SFRD and SMFs at different redshifts and retrieved six-dimensional posterior distributions of the parameters.

We summarize our results as follows:

1. The posteriors inferred from the emulated SFRD and SMFs converge to their true values with relative errors of less than 1% in either the SFRD or the SMFs (Sections 4.1.1 and 5.1, respectively). However, the SFRD requires an order of magnitude more training data to converge than the SMFs do, having two convergent stages—an unstably convergent stage and a stably convergent stage (Figure 3).
2. In the unstably convergent stage, the posterior distribution is bimodal with two *degenerate* peaks (Section 4.1.2). However, the *degeneracy* in the SFRD, which is also confirmed with new cosmological simulations (i.e., it is not an artifact of the emulator), is broken with the SMFs in both the emulator and the cosmological simulations (Section 6.2). This indicates that the SMFs provide stronger constraints for the parameters.
3. In inferences with the mock uncertainty that we add to the emulators to mimic the simulation uncertainty (Appendix B), the posterior inferred from the emulated SFRD successfully captures the variance of the mock uncertainty in the parameters, whereas the SMFs cannot capture any variances in the posterior. In Section 6.3, we

show that uncorrelated uncertainties have a negligible impact on ILI. More work will be required to build a robust model that can precisely reproduce the simulation uncertainty (note that the actual simulation uncertainty might be also uncorrelated with the parameters). In addition, an emulator or theoretical model that can be used to marginalize the simulation uncertainty will be another direction to future work (refer to Appendix D).

4. Employing ILI on the observationally driven SFRD and SMFs of Leja et al. (2020, 2022), we find the inferred SFRD that matches the target observation with a relative error of only 4.1% (Section 4.2). However, the inferred parameters show notable discrepancy with the fiducial values. Moreover, the posterior distribution for the cosmological parameters barely includes the values of the standard cosmology. Meanwhile, the similarly inferred SMFs show significant discrepancies with the target observed SMFs (Section 5.2).
5. The inconsistency between the inferred parameters and standard values could potentially originate from several causes: the limited simulation box size, the resolution effect, the limited parameter space, and the intrinsic physical limits of cosmological simulations with the TNG framework (Sections 6.4 and 6.5).
6. Using MI, we measure the correlation between parameter–observable pairs and find that the performance of inference for each observable largely depends on the correlation between the parameters and the observable (Section 6.1). Also, we confirm that the amount of information in the SMFs is relatively higher than that in the SFRD, which is in line with our inference results.
7. In both correlations and inferences, we find that the AGN parameters (black hole kinetic feedback) are most weakly correlated with both the SFRD and SMFs. This can be attributed to the relatively insignificant impact of the black hole kinetic feedback parameters, within the range varied in CAMELS, on both the formation of galaxies with high stellar mass and star formation in massive galaxies compared to cosmological parameters and stellar wind feedback (Section 6.1.1).
8. In this work, we refrain from conducting physical interpretations of the inferred parameters because the

inference result is sensitive to emulator accuracy, resolution effects, simulation uncertainty, and inaccuracy from a limited parameter space (Sections 6.5).

9. This work is only a cornerstone of calibrating cosmological simulations against observations and provides considerable insights into future directions. In future work, we will focus on resolution convergence, simulation uncertainty, extension of the parameter space as well as of the number of target observables, and inference (only) with cosmological simulations without emulator bias.

We thank Joel Leja for providing preliminary analysis results. The Flatiron Institute is supported by the Simons Foundation. J.K. acknowledges support from the Samsung Science and Technology Foundation under project No. SSTF-BA1802-04, and from the POSCO Science Fellowship of the POSCO TJ Park Foundation. His work was also supported by the National Institute of Supercomputing and Network/Korea Institute of Science and Technology Information with supercomputing resources including technical support, under grants KSC-2020-CRE-0219 and KSC-2021-CRE-0442. D.A.-A. acknowledges support from NSF grants AST-2009687 and AST-2108944, Chandra X-Ray Observatory grant TM2-23006X, and Simons Foundation award CCA-1018464. D.N. acknowledges funding from Deutsche Forschungsgemeinschaft through an Emmy Noether Research Group (grant No. NE 2441/1-1).

Appendix A Rescaling

In Section 2.4, we have discussed the rescaling of observables to account for the incomplete resolution convergence of our simulations. We adopt the rescaling method introduced in Pillepich et al. (2018b, Appendix A). There, the SMFs of TNG300-1 were rescaled using the ratio between TNG100-2 and TNG100-1. The rescaling was based on the resolutions of TNG100-2 and TNG300-1 being essentially the same, and the SMFs of the two simulations being in good agreement despite the difference in simulation volumes. As a result, the rescaled SMF of TNG300-1 (rTNG300-1) coincides with that of TNG100-1 with high accuracy.

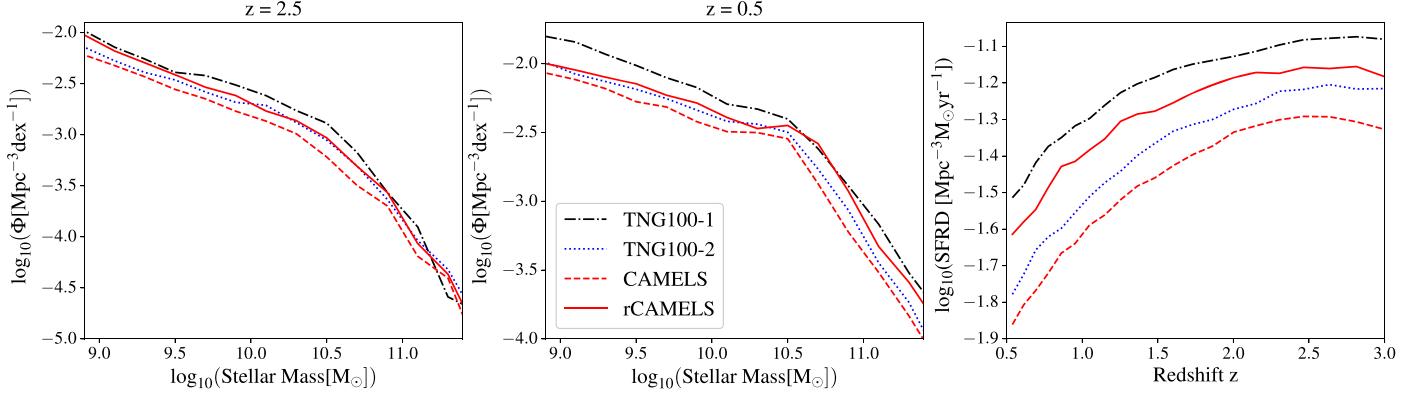


Figure 18. Left and middle: SMFs at $z = 2.5$ and 0.5 . Right: Cosmic SFRDs. The difference between TNG100-1 (black dashed-dotted) and TNG100-2 (blue dotted) illustrates the resolution effect of the TNG simulations. The fiducial CAMELS simulation (red dashed) is rescaled with the ratio between TNG100-1 and TNG100-2. The rescaled CAMELS simulation is shown in red solid lines.

Figure 18 shows the resolution effects in each observable in comparison between TNG100-1 (black dashed-dotted) and TNG100-2 (blue dotted). The fiducial CAMELS (red dashed), computed as the average of the CV set, is overall slightly lower than TNG100-2 in both SMF and SFRD. This is because the resolution of the CAMELS simulations is lower than that of TNG100-2. For dark matter particles, the mass and spatial resolution in CAMELS are $\sim 9.67 \times 10^7 M_\odot h^{-1}$ and ~ 2 kpc comoving, whereas TNG100-2 has a mass resolution of $\sim 5.97 \times 10^7 M_\odot$ and a spatial resolution of ~ 1.48 kpc comoving. In this work, we ignore the discrepancies between the CAMELS simulations and TNG100-2. We rescale the SMF by estimating stellar mass (M_\star) as a function of each bin of halo mass (M_{halo}) as follows:

$$M_{\star, \text{rCAMELS}}(M_{\text{halo}}) = M_{\star, \text{CAMELS}}(M_{\text{halo}}) \times \frac{\langle M_{\star, \text{TNG100-1}}(M_{\text{halo}}) \rangle}{\langle M_{\star, \text{TNG100-2}}(M_{\text{halo}}) \rangle}, \quad (\text{A1})$$

where $\langle \cdot \rangle$ stands for the average over all halos in each bin to which M_{halo} belongs. We multiply the stellar mass of each CAMELS halo by the corresponding rescaling factor (the last fractional term) for that halo mass. The rescaled SMF of the CAMELS simulations (red solid in Figure 18) at $z = 2.5$ (left) is in good agreement with TNG100-1 (black dashed-dotted), whereas the rescaled SMF at $z = 0.5$ (center) deviates appreciably from TNG100-1, especially in the low-mass end. The discrepancy is attributed to the resolution limit that leads to the lower bounds for both the dark matter mass ($\sim 10^8 M_\odot$) and the stellar mass ($\sim 10^7 M_\odot$) of the halos. Since the mass resolution limits of the halos cause a lack of galaxy populations in the vicinity of the limits, construction of the rescaling factor becomes unfeasible in the halo mass range of $[10^8, 10^9] M_\odot$, which largely affects the low-mass end of SMFs. This results in the significant discrepancy between the SMFs of rescaled CAMELS and TNG100-1 in the low-mass end at $z = 0.5$. In this work, we do not employ any post-processing for the zero-stellar-mass galaxies whose stellar mass is not resolved due to the mass resolution limit.

Similarly, we rescale the SFRD by multiplying the SFR of the CAMELS simulations by the rescaling factor as a function

of the bin of halo mass as follows:

$$\text{SFR}_{\text{rCAMELS}}(M_{\text{halo}}) = \text{SFR}_{\text{CAMELS}}(M_{\text{halo}}) \times \frac{\langle \text{SFR}_{\text{TNG100-1}}(M_{\text{halo}}) \rangle}{\langle \text{SFR}_{\text{TNG100-2}}(M_{\text{halo}}) \rangle}, \quad (\text{A2})$$

where dM_\star/dt is the SFR of the galaxy. The rescaled SFRD in Figure 18 (right, red solid) has a consistent offset from TNG100-1 (black dashed-dotted), which is simply attributed to the difference between TNG100-2 (blue dotted) and CAMELS (red dashed).

Appendix B Simulation Uncertainty and Mock Uncertainty

In Sections 2.5 and 4.1.3, we have discussed the simulation uncertainty that originates from various sources of randomness of cosmological simulations. In this section, we study the simulation uncertainty quantitatively using the CV set and model the mock uncertainty. Figure 19 illustrates 27 SFRDs (top left, (a)) and 27 SMFs at $z = 0.5$ (top right, (c)) from 27 cosmological simulations in the CV set, which represents the simulation uncertainty (refer to Section 2 for details of the CV set). We measure the standard deviations of the 27 SFRDs and 27 SMFs as $\sigma_{\text{sim, sfr}} = 0.057$ dex and $\sigma_{\text{sim, smf}} = 0.111$ dex, respectively.

We model the mock uncertainty of the SFRD and SMF using a modified six-dimensional Gaussian noise of the form

$$\mathbf{Z}(\mathbf{x}) = \mathbf{C}^T \frac{\exp(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu}))}{\sqrt{(2\pi)^6 |\boldsymbol{\Sigma}|}}$$

with $\mathbf{C} = \mathbf{I}$ for the SFRD and $[0.3, 0.34, 0.38, 0.43, 0.47, 0.51, 0.56, 0.6, 0.7, 0.9, 1.5, 1.7, 0.5]$ for the SMF, $\boldsymbol{\mu} = \mathbf{0}$, and $\boldsymbol{\Sigma} = \sigma e^{\boldsymbol{\gamma} \boldsymbol{\Gamma}}$, where

$$\boldsymbol{\Gamma} = \begin{bmatrix} 0 & 1 & 2 & 3 & 4 & 5 \\ 1 & 0 & 1 & 2 & 3 & 4 \\ 2 & 1 & 0 & 1 & 2 & 3 \\ 3 & 2 & 1 & 0 & 1 & 2 \\ 4 & 3 & 2 & 1 & 0 & 1 \\ 5 & 4 & 3 & 2 & 1 & 0 \end{bmatrix}.$$

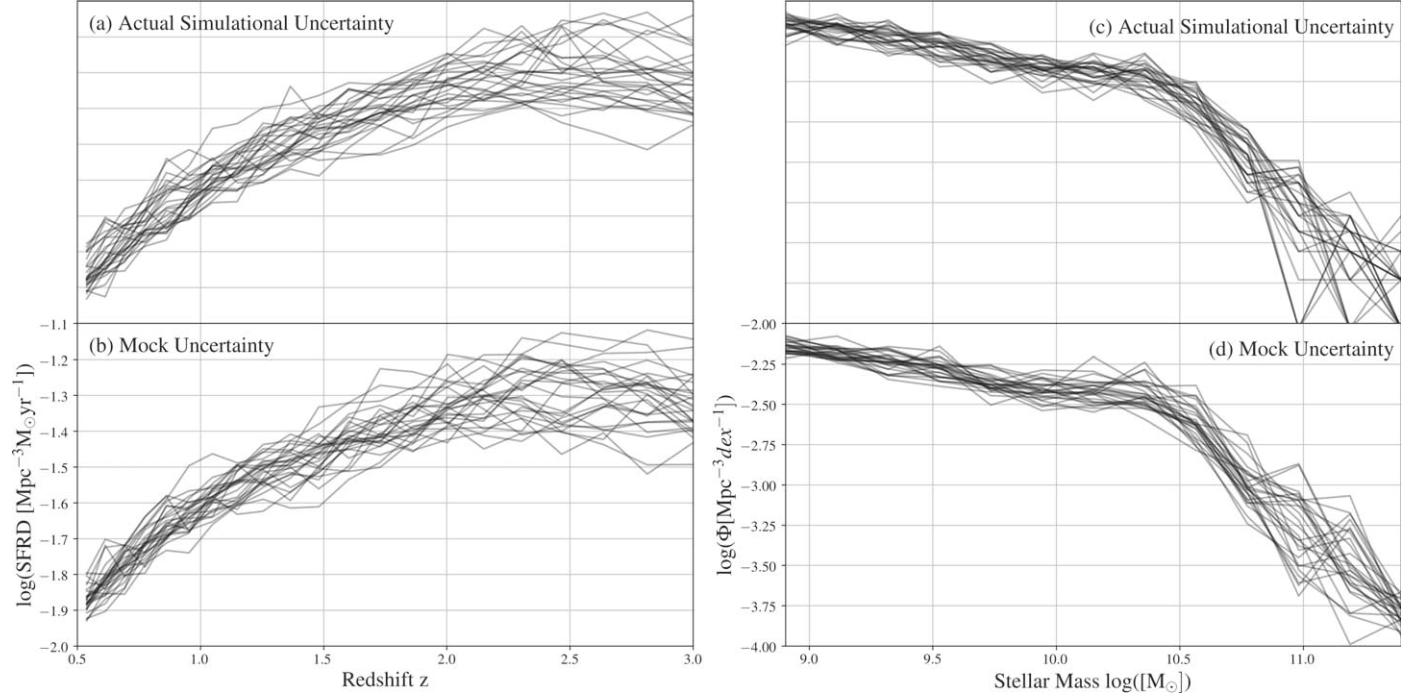


Figure 19. (a) Cosmic SFRDs from 27 simulations of the CV set. (b) Mean of the 27 cosmic SFRDs from the CV set + mock uncertainty. (c) SMFs at $z = 0.5$ from 27 simulations of the CV set. (d) Mean of the 27 SMFs at $z = 0.5$ from the CV set + mock uncertainty. Each panel consists of 27 different curves.

Here, (σ, γ) are $(0.02, 0.2)$ for the SFRD and $(0.02, 2)$ for the SMF. Note that the parameters are empirical. The parameters of the mock uncertainty are tuned such that it can have a visually similar form and similar standard deviations to the simulation uncertainty. In Figure 19, panels (b) and (d) exhibit the 27 SFRDs and 27 SMFs with mock uncertainty. Here, the 27 SFRDs and 27 SMFs are generated by adding the mock uncertainty to the mean of the 27 SFRDs and 27 SMFs from the CV set, respectively. The mock-uncertainty-implemented SFRDs and SMFs are visually in good agreement with the SFRDs and SMFs from the simulation uncertainty. The standard deviations of the mock uncertainty of the SFRD and SMFs are $\sigma_{\text{mock,sfr}} = 0.061$ dex and $\sigma_{\text{mock,smf}} = 0.096$ dex, which approximate those of the simulation uncertainty. Note that in principle, the mock uncertainty depends on the cosmological and astrophysical parameters. Nevertheless, we model the mock uncertainty as if the simulation uncertainty were consistent over the entire parameter space since it is computationally impossible to perform simulations across parameter space to

obtain the simulation uncertainty as a function of the parameters.

Appendix C Architecture of Emulators

In this section, we describe the details of the neural network architectures used for the emulators (Tables 6–11). The structures and hyperparameters of the emulators are automatically optimized using Optuna.

Table 6
SFRD

Layer	Number of Neurons	Dropout	Activation Functions
Input	6		
Fully connected	871	0.25	Leaky ReLU
Fully connected	100	0.26	Leaky ReLU
Output	21		

Note. The learning rate and weight decay are 0.0028 and 1.25e-06.

Table 7
SMFs at $z = 2.5$

Layer	Number of Neurons	Dropout	Activation Functions
Input	6		
Fully connected	451	0.31	Leaky ReLU
Fully connected	727	0.76	Leaky ReLU
Fully connected	851	0.69	Leaky ReLU
Fully connected	890	0.31	Leaky ReLU
Fully connected	825	0.61	Leaky ReLU
Output	21		

Note. The learning rate and weight decay are 0.00015 and 0.00013.

Table 8
SMFs at $z = 2.0$

Layer	Number of Neurons	Dropout	Activation Functions
Input	6		
Fully connected	283	0.28	Leaky ReLU
Fully connected	738	0.32	Leaky ReLU
Output	21		

Note. The learning rate and weight decay are 0.00051 and 0.00031.

Table 9
SMFs at $z = 1.5$

Layer	Number of Neurons	Dropout	Activation Functions
Input	6		
Fully connected	828	0.26	Leaky ReLU
Fully connected	845	0.25	Leaky ReLU
Fully connected	567	0.27	Leaky ReLU
Output	21		

Note. The learning rate and weight decay are 2.4e-05 and 5.7e-05.

Table 10
SMFs at $z = 1.0$

Layer	Number of Neurons	Dropout	Activation Functions
Input	6		
Fully connected	244	0.78	Leaky ReLU
Fully connected	626	0.22	Leaky ReLU
Output	21		

Note. The learning rate and weight decay are 0.00040 and 0.00013.

Table 11
SMFs at $z = 0.5$

Layer	Number of Neurons	Dropout	Activation Functions
Input	6		
Fully connected	890	0.20	Leaky ReLU
Fully connected	298	0.23	Leaky ReLU
Fully connected	836	0.31	Leaky ReLU
Fully connected	600	0.62	Leaky ReLU
Output	21		

Note. The learning rate and weight decay are 4.3e-05 and 3.9e-07.

Appendix D

Marginalization of Simulation Uncertainty in Emulator

This section discusses how emulators marginalize the simulation uncertainty in the LH set that the emulators are trained on. The LH set, consisting of 1000 cosmological simulations, suffers from simulation uncertainty since the initial conditions of the 1000 simulations are all different—that is, it is dominated by cosmic variance, since the scale of our box size is $25 \text{ Mpc } h^{-1}$. The simulation uncertainty plays the role of intrinsic noise in the simulated data in the training of the emulators, which potentially degrades the accuracy of the emulators. Thus, we study whether the emulators can marginalize the noise during training or how much the emulator suffers from the simulation uncertainty. If the emulators can marginalize the simulation uncertainty completely, then the emulator prediction should be equivalent to the mean of the results of a hypothetical suite with a large number of emulations for each point in parameter space—hereafter the uncertainty-marginalized ideal simulation.

We first describe the simulated observable g (SFRD or SMF) as follows:

$$g = g(\boldsymbol{\theta}, \delta(\lambda, \boldsymbol{\theta})), \quad (\text{D1})$$

where $\boldsymbol{\theta}$ is the set of parameters and $\delta(\lambda, \boldsymbol{\theta})$ describes the initial conditions with a random seed λ for cosmological simulations. Here, the marginalization of simulation uncertainty—i.e., the case where the emulator prediction is equivalent to the uncertainty-marginalized ideal simulation—can be written as $\bar{g}(\boldsymbol{\theta}) = \langle g(\boldsymbol{\theta}, \delta(\lambda, \boldsymbol{\theta})) \rangle_{\lambda}$ by averaging g over the sources of randomness λ . The physical analog of the mean of the simulation uncertainty is the mean of the cosmic variance obtained by performing an infinite-volume simulation (only if we ignore the effects of the long-wave limit in the power spectrum in the cosmic variance).

The emulator prediction f can be written with respect to the ideally marginalized prediction—the uncertainty-marginalized ideal simulation—in the form

$$f = f(\boldsymbol{\theta}) = \bar{g}(\boldsymbol{\theta}) + \epsilon(\boldsymbol{\theta}) + m(\boldsymbol{\theta}). \quad (\text{D2})$$

Here, we separate the inaccuracy of the emulator into the training error $\epsilon(\boldsymbol{\theta})$ and the marginalization error $m(\boldsymbol{\theta})$. Then, the

bias of the emulator with respect to the simulations in the LH set, b_{LH} , can be written as

$$b_{\text{LH}} = \langle g(\theta, \delta(\lambda_\theta, \theta)) - f(\theta) \rangle_{\theta_{\text{LH}}} = \langle g(\theta, \delta(\lambda_\theta, \theta)) \rangle_{\theta_{\text{LH}}} - \langle \bar{g}(\theta) \rangle_{\theta_{\text{LH}}} - \langle \epsilon(\theta) \rangle_{\theta_{\text{LH}}} - \langle m(\theta) \rangle_{\theta_{\text{LH}}}, \quad (\text{D3})$$

where θ_{LH} denotes $\theta \in \Theta_{\text{LH}}$ and Θ_{LH} is a set of 1000 parameters in the LH set. Here, we write λ as λ_θ because in the LH set, λ is already determined depending on θ . Assuming that the average of the simulations over the LH set can not only average the simulations from different parameters but also marginalize their simulation uncertainties such that it can be approximately equal to the average of the uncertainty-marginalized ideal simulations over the LH set, i.e., $\langle g(\theta, \delta(\lambda_\theta, \theta)) \rangle_{\theta_{\text{LH}}} \sim \langle \langle g(\theta, \delta(\lambda, \theta)) \rangle_\lambda \rangle_{\theta_{\text{LH}}} \equiv \langle \bar{g}(\theta) \rangle_{\theta_{\text{LH}}}$, we obtain

$$b_{\text{LH}} \simeq -\langle \epsilon(\theta) \rangle_{\theta_{\text{LH}}} - \langle m(\theta) \rangle_{\theta_{\text{LH}}}. \quad (\text{D4})$$

The empirical biases b_{LH} for the SFRD and SMFs are 0.0026 dex and -0.0014 dex, respectively. Compared to the standard deviations of the simulation uncertainty for the SFRD and SMFs (~ 0.06 dex and ~ 0.1 dex), the emulators have relatively small bias, which indicates that the mean of the emulators and the mean of the simulations are in good agreement.

Second, we estimate the variance of emulator prediction with respect to the uncertainty-marginalized ideal simulation in terms of the LH set as follows:

$$\begin{aligned} \hat{\sigma}_{\text{LH}}^2 &= \langle (\bar{g}(\theta) - f(\theta))^2 \rangle_{\theta_{\text{LH}}} \\ &= \langle (\bar{g}(\theta) - g(\theta, \delta(\lambda, \theta)) + g(\theta, \delta(\lambda, \theta)) - f(\theta))^2 \rangle_{\theta_{\text{LH}}}. \end{aligned} \quad (\text{D5})$$

With the definitions $A \equiv \bar{g}(\theta) - g(\theta, \delta(\lambda_\theta, \theta))$, $B \equiv g(\theta, \delta(\lambda_\theta, \theta)) - f(\theta)$, and $C \equiv \bar{g}(\theta) - f(\theta)$, we obtain

$$\begin{aligned} \hat{\sigma}_{\text{LH}}^2 &= \langle C^2 \rangle_{\theta_{\text{LH}}} = \langle A^2 \rangle_{\theta_{\text{LH}}} + \langle B^2 \rangle_{\theta_{\text{LH}}} + 2\langle AB \rangle_{\theta_{\text{LH}}} \\ &= \langle A^2 \rangle_{\theta_{\text{LH}}} + \langle B^2 \rangle_{\theta_{\text{LH}}} + 2\langle A(C - A) \rangle_{\theta_{\text{LH}}} \\ &= -\langle A^2 \rangle_{\theta_{\text{LH}}} + \langle B^2 \rangle_{\theta_{\text{LH}}} + 2\langle AC \rangle_{\theta_{\text{LH}}} \\ &\simeq -\langle A^2 \rangle_{\theta_{\text{LH}}} + \langle B^2 \rangle_{\theta_{\text{LH}}}. \end{aligned} \quad (\text{D6})$$

In the third line, $\langle AC \rangle_{\theta_{\text{LH}}}$ is the covariance of A and C , where A and C are the simulation uncertainty and the deviation of the emulator from the ideally marginalized prediction, respectively. Since the error of the emulator and the simulation uncertainty are, in principle, fully independent, we assume that $\langle AC \rangle_{\theta_{\text{LH}}}$ approximately vanishes. $\langle A^2 \rangle_{\theta_{\text{LH}}}$ is the variance of the simulation uncertainty and $\langle B^2 \rangle_{\theta_{\text{LH}}}$ is the variance of the emulator with respect to the simulation in the LH set, both of which are measurable quantities. $\langle A^2 \rangle_{\theta_{\text{LH}}}$ and $\langle B^2 \rangle_{\theta_{\text{LH}}}$ for the SFRD are 0.036 dex 2 and 0.004 dex 2 , and $\langle A^2 \rangle_{\theta_{\text{LH}}}$ and $\langle B^2 \rangle_{\theta_{\text{LH}}}$ for the SMFs are 0.369 dex 2 and 0.035 dex 2 . Thus, the $\hat{\sigma}_{\text{LH}}^2$ for the SFRD and SMFs are 0.032 dex 2 and 0.334 dex 2 , respectively. As a result, the emulators have negligibly small biases but deviate from the mean of the simulation uncertainty at a similar level to the variance of the simulation uncertainty.

Therefore, we can conclude that the emulators cannot properly marginalize the simulation uncertainty.

Appendix E Definition of Degeneracy

We have encountered degeneracy in Section 4.1.2, which discusses the bimodality of the SFRD. In this section, we probabilistically discuss the relation between uncertainty and degeneracy and mathematically define degeneracy for our purposes. In terms of parameter–observable pairs, degeneracy originates from *indistinguishability* among these pairs. Given an arbitrary observation \mathbf{x}_0 , if there exists a set of parameters Θ_{degen} such that $f(\theta)|_{\theta \in \Theta_{\text{degen}}} = \mathbf{x}_0$, where $f(\theta)$ is, e.g., a theoretical model, a simulation, or a fast approximation method that predicts observable \mathbf{x} as a function of the parameters θ , then one *cannot pinpoint* the parameters θ from which the given observation \mathbf{x}_0 comes. In this case, the pairs $\{(\theta, \mathbf{x}_0) | \theta \in \Theta_{\text{degen}}\}$ are said to be degenerate with respect to observation \mathbf{x}_0 . This can be usually seen in quantum systems such as the spin triplet state under no magnetic field. In consideration of an arbitrary uncertainty $\mathbf{Z}(\eta)$ in the observable²⁴, the pairs can be written as either $(\theta, \mathbf{x} + \mathbf{Z}(\eta))$ or $(\theta + \epsilon(\eta), \mathbf{x})$,²⁵ where η is a random seed (Section 4.1.3 and Appendix B). One can notice that \mathbf{Z} and ϵ are basically random variables that require a probabilistic treatment. In the following section, degeneracy will be discussed in a probabilistic manner.

The ideal probabilistic inference naturally traces the propagation of uncertainty in the observation onto each parameter (see footnote 16). However, consistent, robust confinement for the posterior density is necessary to define a finite region of the degenerate parameter space since an arbitrary inferred posterior density is generally well defined over the entire parameter space.

We define degeneracy as follows.

Given an arbitrary probability density $p(\theta)$ and ζ_{thres} , there exist p_{thres} and Θ_{degen} satisfying $\Theta_{\text{degen}} = \{\theta | p(\theta) \geq p_{\text{thres}}\}$ such that

$$\text{(continuous)} \quad \frac{\int_{\theta \in \Theta_{\text{degen}}} p(\theta) d\theta}{\int_{\theta \in \mathcal{V}} p(\theta) d\theta} = \zeta_{\text{thres}} \quad (\text{E1})$$

or

$$\text{(discrete)} \quad \frac{\sum_{\theta \in \Theta_{\text{degen}}} p(\theta) \Delta\theta}{\sum_{\theta \in \mathcal{V}} p(\theta) \Delta\theta} = \zeta_{\text{thres}}, \quad (\text{E2})$$

²⁴ For example, the cosmological simulations f reproduce different output \mathbf{x} depending on the initial conditions η even with the same set of physical and free parameters θ , i.e., $f(\theta) = \mathbf{x} + \mathbf{Z}(\eta)$, where $\mathbf{Z}(\eta)$ stands for cosmic variance.

²⁵ Given a model $f(\theta)$ that predicts observable $\mathbf{Z}(\eta)$ can be propagated onto parameters θ as follows: $\mathbf{x} + \mathbf{Z}(\eta) \xrightarrow[\mathbf{x} = f(\theta)]{\theta = g(\mathbf{x})} \theta + \epsilon(\eta)$, where η is a random seed. Here, $g(\mathbf{x})$ can be a set of locally defined functions that satisfy $\mathbf{x} = f(\theta)$.

where $0 < \zeta_{\text{thres}} < 1$ and \mathcal{V} stands for the entire parameter space. Here, the parameters θ in Θ_{degen} are said to be degenerate. Note that the threshold value ζ_{thres} is a free parameter.

In the case of a high-dimensional problem with an intractable probability distribution, it is practically impossible to numerically integrate the probability distribution $p(\theta)$ over \mathcal{V} or Θ_{degen} even if Θ_{degen} is known. In low-dimensional problems, simple quadrature methods can work. For example, we can evaluate a function on a fixed grid of points, then apply the trapezoid rule. However, in high dimensions, the number of grid points grows exponentially.²⁶ Hence, MCMC methods are widely adopted to integrate high-dimensional functions.

We utilize the MCMC sampling to estimate the degenerate parameter space. MCMC provides millions of parameter sets such that the ratio of the number of parameter sets in each bin, $N(\Delta\theta_{\text{bin}})$, to the total number of samplings, N_{total} , represents the approximate probability of that bin. In other words, $p(\theta)\Delta\theta_{\text{bin}} \simeq N(\Delta\theta_{\text{bin}})/N_{\text{total}}$, where $\Delta\theta_{\text{bin}}$ is the size of each bin. We can rewrite the left-hand side of Equation (E2) into the form

$$\frac{\sum_{\theta \in \Theta_{\text{degen}}} p(\theta)\Delta\theta}{\sum_{\theta \in \mathcal{V}} p(\theta)\Delta\theta} = \frac{\sum_{\theta \in \Theta_{\text{degen}}} N(\Delta\theta_{\text{bin}})}{N_{\text{total}}}. \quad (\text{E3})$$

Then,

$$\sum_{\theta \in \Theta_{\text{degen}}} N(\Delta\theta_{\text{bin}}) = \zeta_{\text{thres}} N_{\text{total}}. \quad (\text{E4})$$

Together with $\Theta_{\text{degen}} = \{\theta \mid p(\theta) \geq p_{\text{thres}}\}$, we can write $\Theta_{\text{degen}} = \arg \max_{D \in \mathcal{D}} p_{\text{thres}}; \mathcal{D} = \{D \mid \tilde{p}(\theta) > p_{\text{thres}} \forall \theta; \theta \in D, n(D) = \zeta_{\text{thres}} N_{\text{total}}, D \subset \Theta_{\text{MCMC}}\}$, where Θ_{MCMC} , \tilde{p} , and $n(\cdot)$ are the set of sampled parameters from MCMC; an arbitrary surrogate posterior function, such as the NDE or kernel density estimation; and the cardinality of the set. The set of degenerate points in the parameter space, Θ_{degen} , is a subset of Θ_{MCMC} such that the sum of the probability of the parameters in the subset is equal to ζ_{thres} (refer to Equation (E2)). In practice, we identify Θ_{degen} among Θ_{MCMC} as follows: (1) estimate $\tilde{p}(\theta)$ for all $\theta \in \Theta_{\text{MCMC}}$; (2) pair θ and $\tilde{p}(\theta)$ into $(\theta, \tilde{p}(\theta))$; (3) sort $(\theta, \tilde{p}(\theta))$ in descending order along $\tilde{p}(\theta)$; and (4) identify the first $\zeta_{\text{thres}} N_{\text{total}}$ of θ of the sorted pairs as Θ_{degen} .

ORCID iDs

Yongseok Jo [ID](https://orcid.org/0000-0003-3977-1761) <https://orcid.org/0000-0003-3977-1761>
 Shy Genel [ID](https://orcid.org/0000-0002-3185-1540) <https://orcid.org/0000-0002-3185-1540>
 Benjamin Wandelt [ID](https://orcid.org/0000-0002-5854-8269) <https://orcid.org/0000-0002-5854-8269>
 Francisco Villaescusa-Navarro [ID](https://orcid.org/0000-0002-4816-0455) <https://orcid.org/0000-0002-4816-0455>
 Greg L. Bryan [ID](https://orcid.org/0000-0003-2630-9228) <https://orcid.org/0000-0003-2630-9228>
 Daniel Anglés-Alcázar [ID](https://orcid.org/0000-0001-5769-4945) <https://orcid.org/0000-0001-5769-4945>
 Daniel Foreman-Mackey [ID](https://orcid.org/0000-0002-9328-5652) <https://orcid.org/0000-0002-9328-5652>
 Dylan Nelson [ID](https://orcid.org/0000-0001-8421-5890) <https://orcid.org/0000-0001-8421-5890>
 Ji-hoon Kim [ID](https://orcid.org/0000-0003-4464-1160) <https://orcid.org/0000-0003-4464-1160>

²⁶ In our case, the six-dimensional parameter space requires $(10^2)^6 = 10^{12}$ grid points, where we space each axis with 100 grid points, which means that 10^{12} arithmetic calculations at least are needed.

References

Aihara, H., Arimoto, N., Armstrong, R., et al. 2018, *PASJ*, **70**, S4
 Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. 2019, arXiv:1907.10902
 Alsing, J., Charnock, T., Feeney, S., & Wandelt, B. 2019, *MNRAS*, **488**, 4440
 Alsing, J., Wandelt, B. D., & Feeney, S. M. 2018, *MNRAS*, **477**, 2874
 Anglés-Alcázar, D., Davé, R., Faucher-Giguère, C.-A., Özel, F., & Hopkins, P. F. 2017, *MNRAS*, **464**, 2840
 Anglés-Alcázar, D., Davé, R., Özel, F., & Oppenheimer, B. D. 2014, *ApJ*, **782**, 84
 Baldry, I. K., Driver, S. P., Loveday, J., et al. 2012, *MNRAS*, **421**, 621
 Benson, A. J. 2014, *MNRAS*, **444**, 2599
 Bird, S., Ni, Y., DiMatteo, T., et al. 2022, *MNRAS*, **512**, 3703
 Bondi, H. 1952, *MNRAS*, **112**, 195
 Bondi, H., & Hoyle, F. 1944, *MNRAS*, **104**, 273
 Bower, R. G., Benson, A. J., Malbon, R., et al. 2006, *MNRAS*, **370**, 645
 Boylan-Kolchin, M., Springel, V., White, S. D. M., Jenkins, A., & Lemson, G. 2009, *MNRAS*, **398**, 1150
 Brunner, H., Liu, T., Lamer, G., et al. 2022, *A&A*, **661**, A1
 Ceverino, D., & Klypin, A. 2008, in IAU Symp. 245, Formation and Evolution of Galaxy Bulges, ed. M. Bureau, E. Athanassoula, & B. Barbuy (Cambridge: Cambridge Univ. Press), 33
 Crammer, K., Brehmer, J., & Loupou, G. 2020, *PNAS*, **117**, 30055
 Croton, D. J., Springel, V., White, S. D. M., et al. 2006, *MNRAS*, **365**, 11
 Dalmasso, N., Zhao, D., Izbicki, R., & Lee, A. B. 2021, arXiv:2107.03920
 Dave, R., Anglés-Alcázar, D., Narayanan, D., et al. 2019, *MNRAS*, **486**, 2827
 Dave, R., Thompson, R., & Hopkins, P. F. 2016, *MNRAS*, **462**, 3265
 Dey, A., Schlegel, D. J., Lang, D., et al. 2019, *AJ*, **157**, 168
 Dubois, Y., Gavazzi, R., Peirani, S., & Silk, J. 2013, *MNRAS*, **433**, 3297
 Dubois, Y., Pichon, C., Welker, C., et al. 2014, *MNRAS*, **444**, 1453
 Durkan, C., Murray, I., & Papamakarios, G. 2020, arXiv:2002.03712
 Elliott, E. J., Baugh, C. M., & Lacey, C. G. 2021, *MNRAS*, **506**, 4011
 Fabian, A. C. 2012, *ARA&A*, **50**, 455
 Feng, Y., DiMatteo, T., Croft, R. A., et al. 2016, *MNRAS*, **455**, 2778
 Feroz, F., & Hobson, M. P. 2008, *MNRAS*, **384**, 449
 Fontanot, F., Calabro, A., Talia, M., et al. 2021, *MNRAS*, **504**, 4481
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Gabré, M., Manoel, A., Luneau, C., et al. 2019, *JSMTE*, **12**, 124014
 Gallazzi, A., Charlot, S., Brinchmann, J., White, S. D. M., & Tremonti, C. A. 2005, *MNRAS*, **362**, 41
 Genel, S., Vogelsberger, M., Springel, V., et al. 2014, *MNRAS*, **445**, 175
 Genel, S., Bryan, G. L., Springel, V., et al. 2019, *ApJ*, **871**, 21
 Gerardi, F., Feeney, S. M., & Alsing, J. 2021, *PhRvD*, **104**, 083531
 Goldfeld, Z., van den Berg, E., Greenewald, K., et al. 2018, arXiv:1810.05728
 Hermans, J., Begy, V., & Loupou, G. 2019, arXiv:1903.04057
 Hirschmann, M., Dolag, K., Saro, A., et al. 2014, *MNRAS*, **442**, 2304
 Ho, M.-F., Bird, S., & Shelton, C. R. 2022, *MNRAS*, **509**, 2551
 Hopkins, P. F., Wetzel, A., Keres, D., et al. 2018, *MNRAS*, **480**, 800
 Hoyle, F., & Lyttleton, R. A. 1939, *PCPS*, **35**, 405
 Hubber, D. A., Falle, S. A. E. G., & Goodwin, S. P. 2013, *MNRAS*, **432**, 711
 Huppenkothen, D., & Bachetti, M. 2022, *MNRAS*, **511**, 5689
 Karamanis, M., & Beutler, F. 2020, arXiv:2002.06212
 Karamanis, M., Beutler, F., & Peacock, J. A. 2021, *MNRAS*, **508**, 3589
 Keller, B. W., Wadsley, J. W., Wang, L., & Krijssen, J. M. D. 2019, *MNRAS*, **482**, 2244
 Khandai, N., DiMatteo, T., Croft, R., et al. 2015, *MNRAS*, **450**, 1349
 Kinney, J. B., & Atwal, G. S. 2014, *PNAS*, **111**, 3354
 Klypin, A., Yepes, G., Gottlöber, S., Prada, F., & Heß, S. 2016, *MNRAS*, **457**, 4340
 Klypin, A. A., Trujillo-Gomez, S., & Primack, J. 2011, *ApJ*, **740**, 102
 Kodi Ramanah, D., Wojtak, R., & Arendse, N. 2021, *MNRAS*, **501**, 4080
 Kormendy, J., & Gebhardt, K. 2001, in AIP Conf. Proc. 586, 20th Texas Symp. on Relativistic Astrophysics, ed. J. C. Wheeler & H. Martel (Melville, NY: AIP), 363
 Kormendy, J., & McClure, R. D. 1993, *AJ*, **105**, 1793
 Kraskov, A., Stogbauer, H., & Grassberger, P. 2004, *PhRvE*, **69**, 066138
 Kuijken, K., Heymans, C., Dvornik, A., et al. 2019, *A&A*, **625**, A2
 Laigle, C., McCracken, H. J., Ilbert, O., et al. 2016, *ApJS*, **224**, 24
 Le Brun, A. M. C., McCarthy, I. G., Schaye, J., & Ponman, T. J. 2014, *MNRAS*, **441**, 1270
 Lee, J., Shin, J., Snaith, O. N., et al. 2021, *ApJ*, **908**, 11
 Legin, R., Hezaveh, Y., Perreault Levasseur, L., & Wandelt, B. 2021, arXiv:2112.05278
 Leja, J., Speagle, J. S., Johnson, B. D., et al. 2020, *ApJ*, **893**, 111
 Leja, J., Johnson, B. D., Conroy, C., et al. 2019, *ApJ*, **877**, 140
 Leja, J., Speagle, J. S., Ting, Y.-S., et al. 2022, *ApJ*, **936**, 165

Lemos, P., Jeffrey, N., Whiteway, L., et al. 2021, *PhRvD*, **103**, 023009

Lequeux, J., Peimbert, M., Rayo, J. F., Serrano, A., & Torres-Peimbert, S. 1979, *A&A*, **500**, 145

Lia, C., Carraro, G., & Salucci, P. 2000, *A&A*, **360**, 76

Lia, C., Portinari, L., & Carraro, G. 2002, *MNRAS*, **330**, 821

Lu, Y., Mo, H. J., Katz, N., & Weinberg, M. D. 2012, *MNRAS*, **421**, 1779

Lu, Y., Mo, H. J., Lu, Z., Katz, N., & Weinberg, M. D. 2014, *MNRAS*, **443**, 1252

Madau, P., & Dickinson, M. 2014, *ARA&A*, **52**, 415

Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, **480**, 5113

McCarthy, I. G., Schaye, J., Bird, S., & Le Brun, A. M. C. 2017, *MNRAS*, **465**, 2936

McGaugh, S. S., Lelli, F., Schombert, J. M., et al. 2021, *AJ*, **162**, 202

McGaugh, S. S., Schombert, J. M., Bothun, G. D., & de Blok, W. J. G. 2000, *ApJL*, **533**, L99

McLeod, D. J., McLure, R. J., Dunlop, J. S., et al. 2021, *MNRAS*, **503**, 4413

Merritt, D., & Ferrarese, L. 2001, *ApJ*, **547**, 140

Mohamed, S., & Lakshminarayanan, B. 2016, arXiv:1610.03483

Mowla, L. A., Dokkum, P. v., Brammer, G. B., et al. 2019, *ApJ*, **880**, 57

Nagamine, K. 2010, *AdAst*, **2010**, 651621

Naiman, J. P., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, **477**, 1206

Nelson, D., Pillepich, A., Genel, S., et al. 2015, *A&C*, **13**, 12

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

Noeske, K. G., Weiner, B. J., Faber, S. M., et al. 2007, *ApJL*, **660**, L43

Oh, B. K., Smith, B. D., Peacock, J. A., & Khochfar, S. 2020, *MNRAS*, **497**, 5203

Page, M. J., Symeonidis, M., Vieira, J. D., et al. 2012, *Natur*, **485**, 213

Paninski, L. 2003, *Neural Comput.*, **15**, 1191

Papamakarios, G., Sterratt, D. C., & Murray, I. 2018, arXiv:1805.07226

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825

Pillepich, A., Springel, V., Nelson, D., et al. 2018a, *MNRAS*, **473**, 4077

Pillepich, A., Nelson, D., Hernquist, L., et al. 2018b, *MNRAS*, **475**, 648

Piotrowska, J. M., Bluck, A. F. L., Maiolino, R., & Peng, Y. 2022, *MNRAS*, **512**, 1052

Ragagnin, A., Dolag, K., Biffi, V., et al. 2017, *A&C*, **20**, 52

Regan, J. A., Johansson, P. H., & Haehnelt, M. G. 2014, *MNRAS*, **439**, 1160

Rosas-Guevara, Y. M., Bower, R. G., Schaye, J., et al. 2015, *MNRAS*, **454**, 1038

Ross, B. C. 2014, *PLoS*, **9**, e87357

Schaye, J., Crain, R. A., Bower, R. G., et al. 2015, *MNRAS*, **446**, 521

Shen, S., Mo, H. J., White, S. D. M., et al. 2003, *MNRAS*, **343**, 978

Skelton, R. E., Whitaker, K. E., Momcheva, I. G., et al. 2014, *ApJS*, **214**, 24

Snaith, O. N., Park, C., Kim, J., & Rosdahl, J. 2018, *MNRAS*, **477**, 983

Somerville, R. S., & Dave, R. 2015, *ARA&A*, **53**, 51

Speagle, J. S., Steinhardt, C. L., Capak, P. L., & Silverman, J. D. 2014, *ApJS*, **214**, 15

Springel, V. 2010, *MNRAS*, **401**, 791

Springel, V., White, S. D. M., Jenkins, A., et al. 2005, *Natur*, **435**, 629

Springel, V., Pakmor, R., Pillepich, A., et al. 2018, *MNRAS*, **475**, 676

Su, K.-Y., Hopkins, P. F., Bryan, G. L., et al. 2021, *MNRAS*, **507**, 175

Sugiyama, M., Suzuki, T., & Kanamori, T. 2012, *Density Ratio Estimation in Machine Learning* (Cambridge: Cambridge Univ. Press)

Tal, T., Dekel, A., Oesch, P., et al. 2014, *ApJ*, **789**, 164

Tam, S.-I., Umetsu, K., & Amara, A. 2022, *ApJ*, **925**, 145

Tejero-Cantero, A., Boelts, J., Deistler, M., et al. 2020, *JOSS*, **5**, 2505

Terrazas, B. A., Bell, E. F., Pillepich, A., et al. 2020, *MNRAS*, **493**, 1888

Tremmel, M., Karcher, M., Governato, F., et al. 2017, *MNRAS*, **470**, 1121

Trotta, R., Johannesson, G., Moskalenko, I. V., et al. 2011, *ApJ*, **729**, 106

Tully, R. B., & Fisher, J. R. 1977, *A&A*, **500**, 105

van der Wel, A., Franx, M., van Dokkum, P. G., et al. 2014, *ApJ*, **788**, 28

Venholo, A., Peletier, R., Laurikainen, E., et al. 2018, *A&A*, **620**, A165

Villaescusa-Navarro, F., Angles-Alcazar, D., Genel, S., et al. 2021, *ApJ*, **915**, 71

Villanueva-Domingo, P., Villaescusa-Navarro, F., Genel, S., et al. 2021, arXiv:2111.14874

Vogelsberger, M., Marinacci, F., Torrey, P., & Puchwein, E. 2020, *NatRP*, **2**, 42

Vogelsberger, M., Genel, S., Springel, V., et al. 2014a, *MNRAS*, **444**, 1518

Vogelsberger, M., Genel, S., Springel, V., et al. 2014b, *Natur*, **509**, 177

Wang, Z., Xu, H., Yang, X., et al. 2021, *SCPMA*, **64**, 289811

Weinberger, R., Springel, V., & Pakmor, R. 2020, *ApJS*, **248**, 32

Weinberger, R., Springel, V., Hernquist, L., et al. 2017, *MNRAS*, **465**, 3291

Wellons, S., Faucher-Giguère, C.-A., Hopkins, P. F., et al. 2022, arXiv:2203.06201

Zhao, X., Mao, Y., Cheng, C., & Wandelt, B. D. 2022a, *ApJ*, **926**, 151

Zhao, X., Mao, Y., & Wandelt, B. D. 2022b, *ApJ*, **933**, 236

Zinger, E., Pillepich, A., Nelson, D., et al. 2020, *MNRAS*, **499**, 768