Developing a new physics-first galaxy formation model for cosmological simulations

Jan David Burger
Cosmic Horizons, Charlottesville VA,
07/15/2026, (with Matt Ho and Volker Springel)

MAX-PLANCK-INSTITUT
FÜR ASTROPHYSIK

The Illustris-TNG model as a baseline

The TNG-model (Pillepich 2018) remains one of the most successful galaxy formation models until today, including subgrid models for:

- BH accretion and feedback
- Star formation and cooling
- Stellar winds
- Magnetohydrodynamics

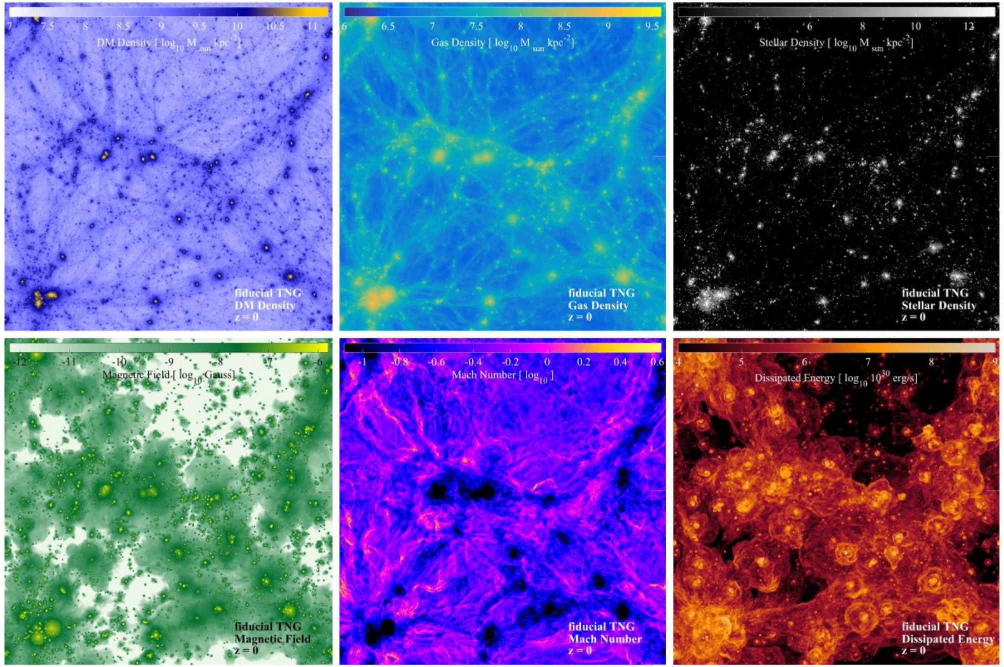

Pillepich et al. 2018

Fiducial TNG scaling relations

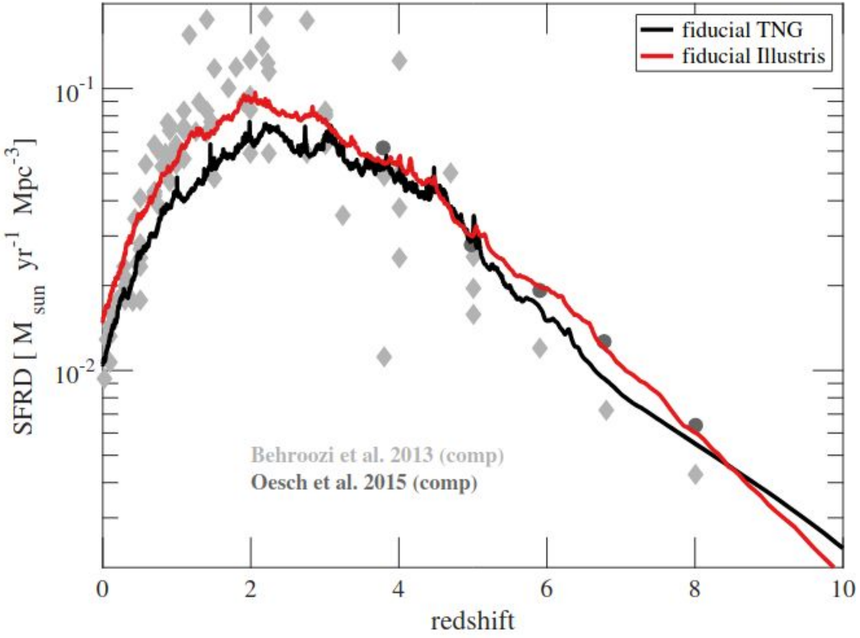

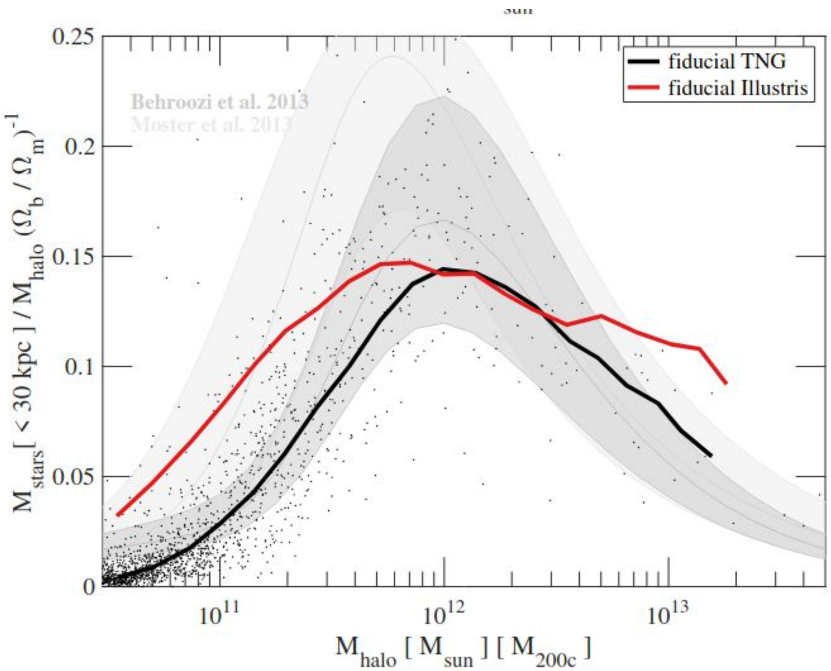

Pillepich et al. 2018

Fiducial TNG scaling relations

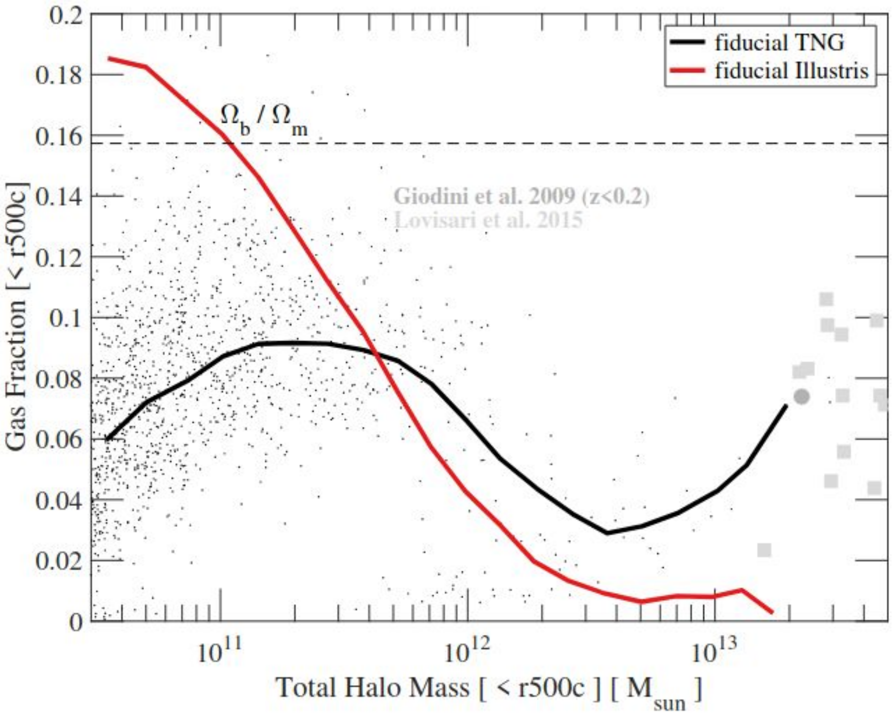

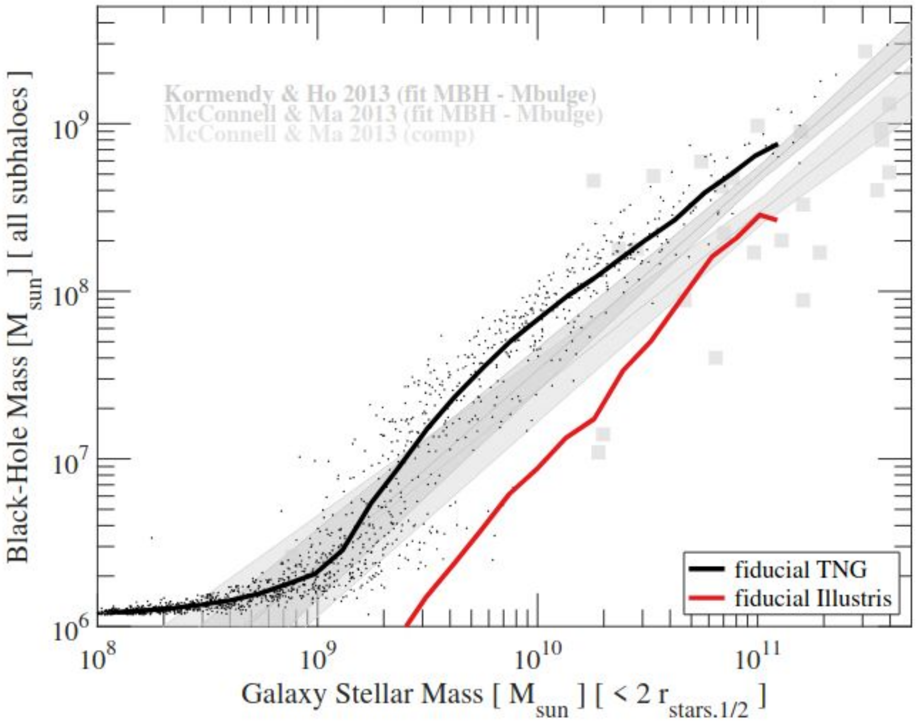

Pillepich et al. 2018

The TNG model today

TNG remains a reference standard for galaxy formation models, but:

- Star formation model includes a relatively ad-hoc ISM description
- Stellar feedback is included through hydrodynamically decoupled winds
- Black holes are fixed to the center of a halo’s potential and seeded in fairly heavy haloes
- Black hole accretion model assumes spherical symmetry (Bondi-Hoyle accretion) and black hole feedback cannot explain the low observed gas fractions in massive haloes
- There is a growing desire for new, physics-motivated subgrid models that improve our understanding (and maybe the match with observations)

The Learning the Universe collaboration

Within the LtU collaboration, part of our effort focuses on the development of new, physics-informed subgrid models for galaxy formation.

- The Arkenstone galactic wind model (Bennet et al. 2025)
- The PRFM-based star formation model (Burger et al. 2026)
- The BRAHMA model for black hole seeding (Bhowmick et al. 2024)
- A new dynamical friction model for black holes (Genina et al. 2024)
- Blended (Illustris + TNG) black hole feedback (Chen et al. in prep.)

A complete new galaxy formation model requires the joint calibration of the combination of all of these subgrid models!

A case study on torque-limited black hole accretion

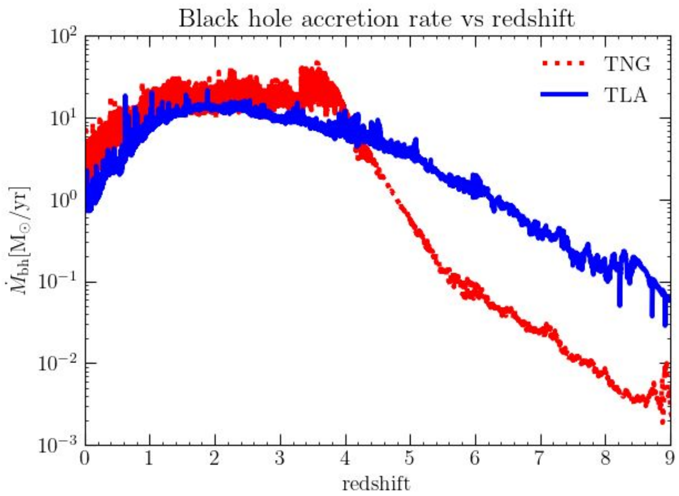

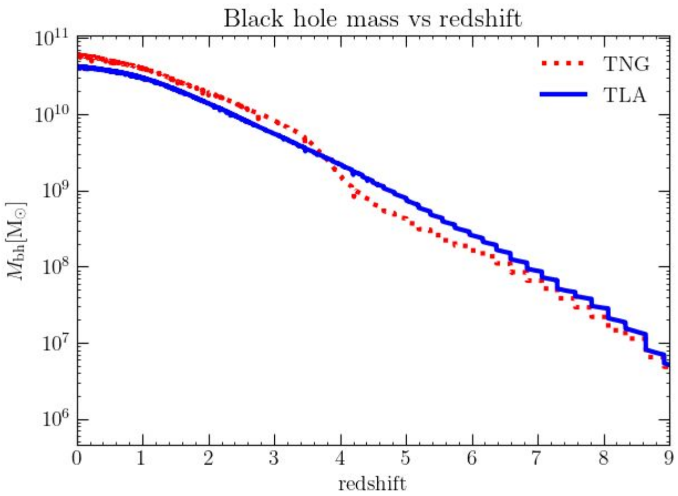

- Torque-limited accretion allows for higher accretion rates than Bondi accretion at high redshift
- This can potentially help to explain luminous high-z quasars
- The model has to be calibrated together with the TNG feedback model to match observations

The TLA model and its equations

Torque-limited accretion, as introduced in Angles-Alcazar et al. 2013, calculates two different BH accretion rates for hot gas and cold gas:

$$\dot{M}_{\rm Bondi} = \epsilon_m \frac{4\pi G^2 M_{\rm BH}^2 \rho}{(v^2 + c_s^2)^{(3/2)}}$$

$$\dot{M}_{\rm Torque} = \epsilon_{\rm T} f_{\rm d}^{5/2} \times \left(\frac{M_{\rm BH}}{10^8 {\rm M}_\odot}\right)^{1/6} \times \left(\frac{M_{\rm enc}(R_0)}{10^9 {\rm M}_\odot}\right)\left(\frac{R_0}{100\,{\rm pc}}\right)^{-3/2} \times \left(1 + \frac{f_0}{f_{\rm gas}}\right)^{-1} {\rm M}_\odot\,{\rm yr}^{-1}$$

Bondi rate normalization for hot gas

TLA normalization for cold gas

- Cold and hot gas separated by temperature threshold
- Disk fraction calculated within R_0
- Gas fraction in the disk as well
- f_0 is a normalization parameter

The TNG feedback model

The TNG feedback model decides which feedback channel to use based on whether the BH is in high-accretion state or low accretion state:

$$\frac{\dot{M}_{\mathrm{Bondi}}}{\dot{M}_{\mathrm{Edd}}} \geq \chi \qquad \chi = \min\left[\chi_0 \left(\frac{M_{\mathrm{BH}}}{10^8\,\mathrm{M}_\odot}\right)^\beta, 0.1\right]$$

Quasarthreshold parameter, fixing the transition between quasar and radio mode

•

Note that varying the threshold mass would lead to a degenerate effect, and that β is fixed to 2.

The TNG feedback model

In the quasar mode, the injected energy is regulated by the feedback efficiency parameter

$$\Delta \dot{E}_{\rm high} = \epsilon_{\rm f,high} \epsilon_{\rm r} \dot{M}_{\rm BH} c^2$$

Blackhole feedback efficiency parameter that regulates the quasar mode

For the radio mode, the efficiency function is more complicated:

Minimum density factor for radio feedback

$$\Delta \dot{E}_{\rm low} = \epsilon_{\rm f,kin} \dot{M}_{\rm BH} c^2 \quad \epsilon_{\rm f,kin} = \min\left(\frac{\rho}{f_{\rm thresh}\,\rho_{\rm SFthresh}}, 0.2\right)$$

The TLA calibration suite

- We generate 150 latin hypercube samples (in log-space) of our 5 selected parameters. We vary:
    - Bondi rate normalization and TLA normalization from 0.1 to 10
    - The quasar threshold from 0.0006 to 0.006
    - Quasar feedback efficiency from 0.03 to 0.3
    - Minimum density factor for radio feedback from 0.002 to 0.5
- We then run 150 hydro simulations in (50 Mpc)^3 boxes at TNG-300 resolution to calibrate:
    - We perform PCA on 6 measured scaling relations and then train Gaussian process emulators on the principal components
    - Using the emulator, we carry out sensitivity studies
    - Finally, we attempt to recover withheld parameter vectors using MCMC

TLA: Stellar-to-halo mass relation

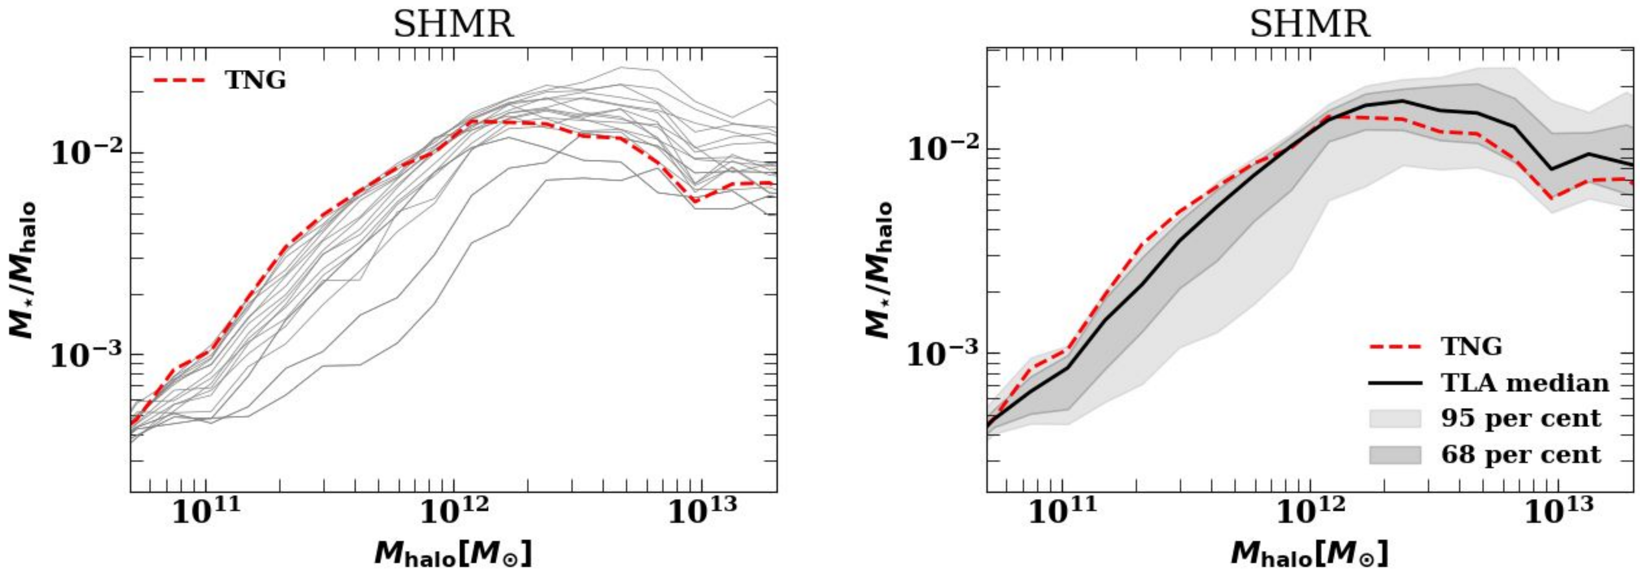

- Stellar-to-halo mass ratios that are covered by the simulation suite encompass the TNG line
- No individual line, however, folllows the shape of the TNG stellar-to-halo-mass relation

TLA: black hole vs stellar mass

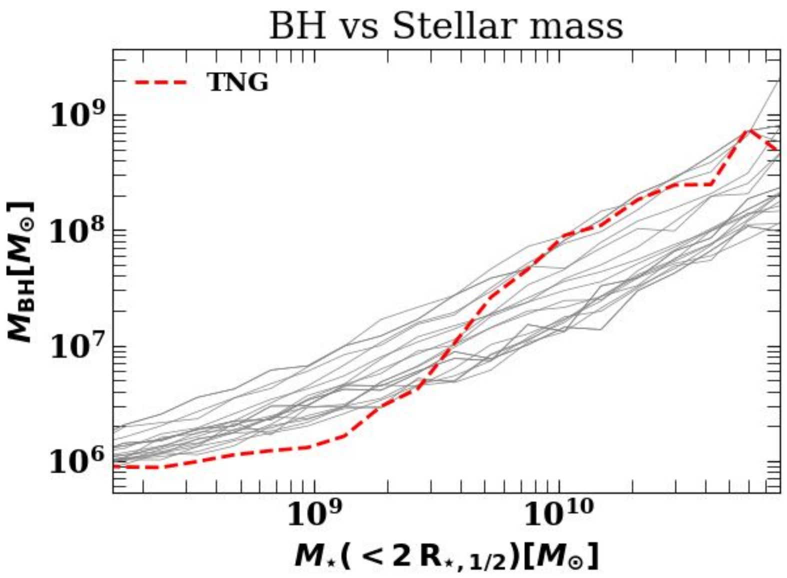

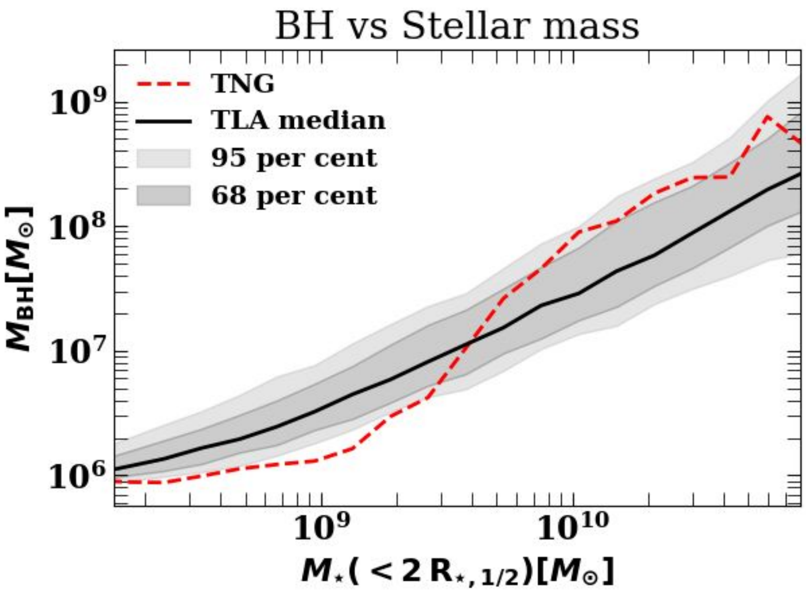

- In all TLA runs, black hole mass scales with stellar mass almost as a pure power law
- This is in contrast to the S shape obtained in the TNG run

TLA: Gas fractions

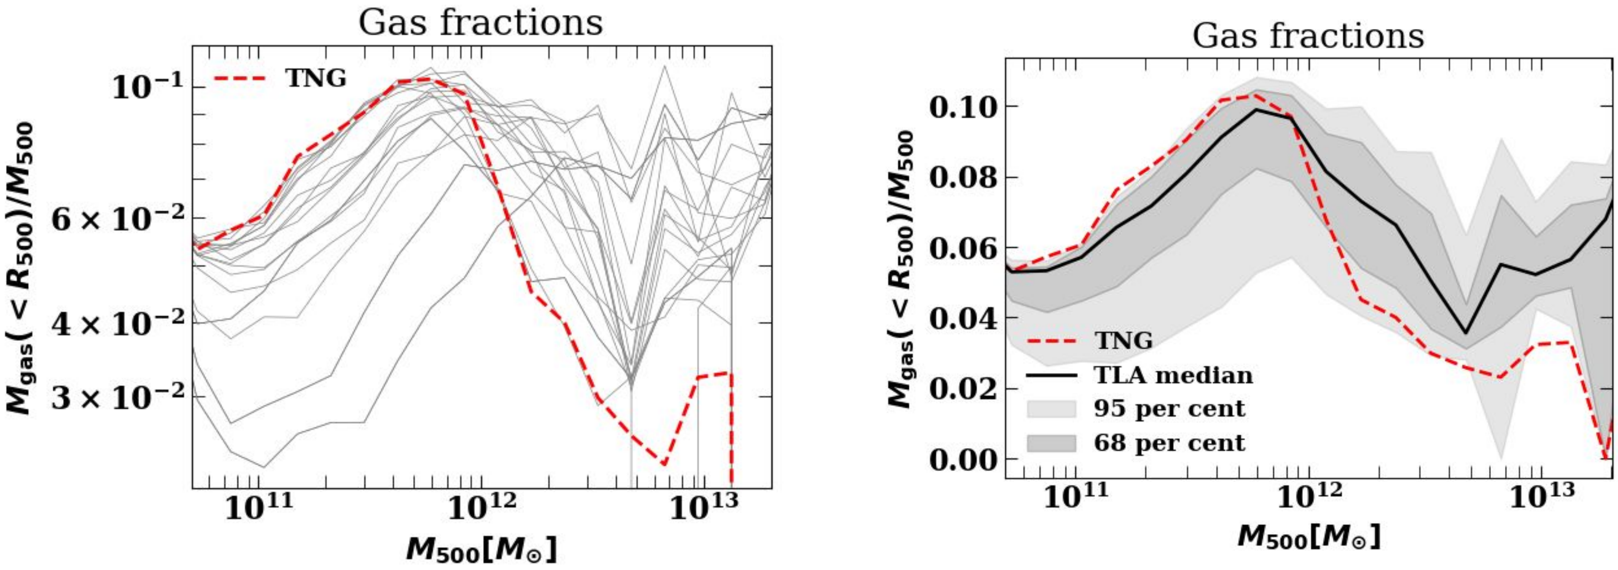

- Simulated gas fractions cover a large range, but at the high mass end most of them fall above the TNG line
- This indicates that possibly, a more substantial change in the feedback model is needed (see upcoming work of Chen et al.)

TLA: Quenched fractions

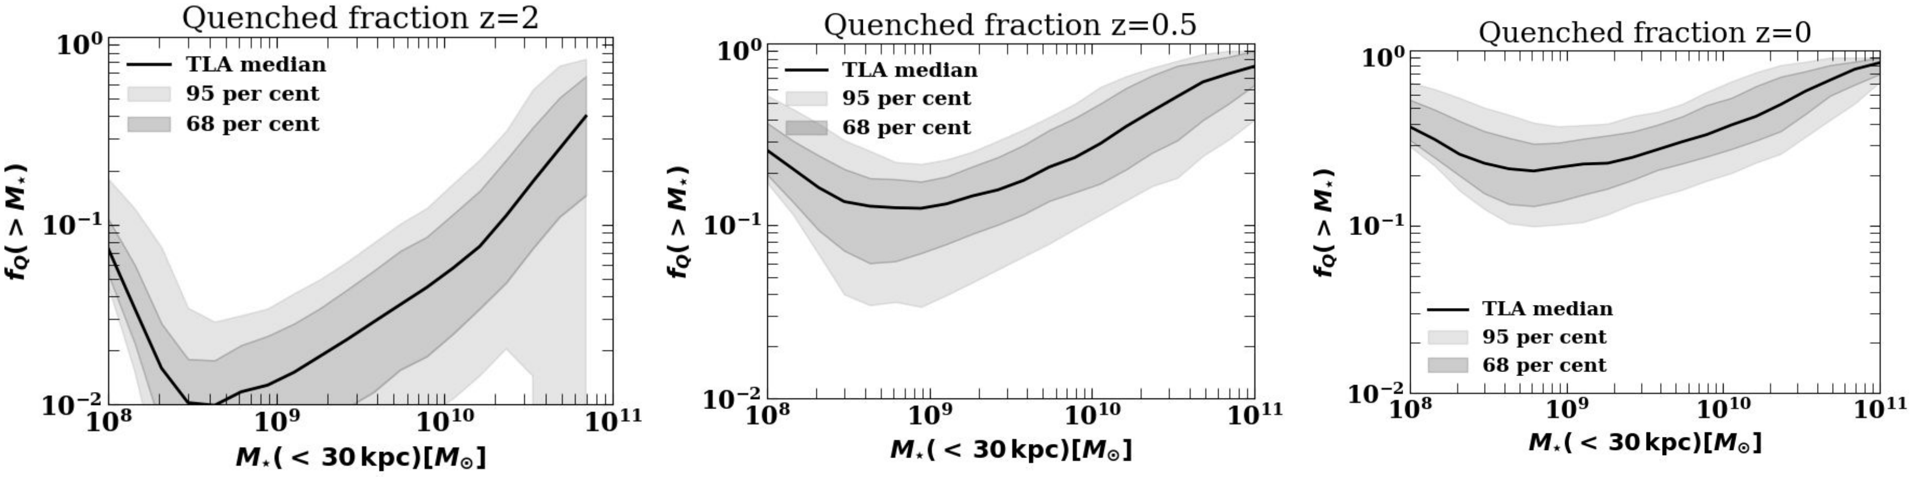

- Galaxies become increasingly quenched towards lower redshift
- There is a large spread in quenched fraction for the different model variations, in particular at high z

Principal component analysis and LOO validation: BH vs stellar mass

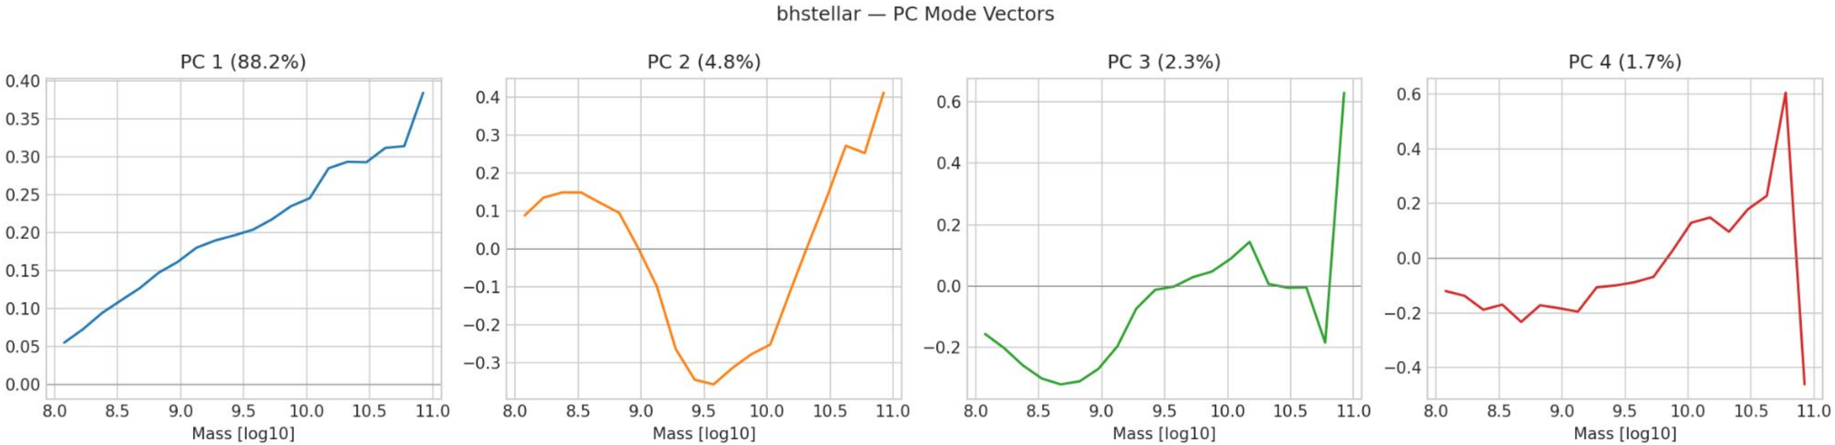

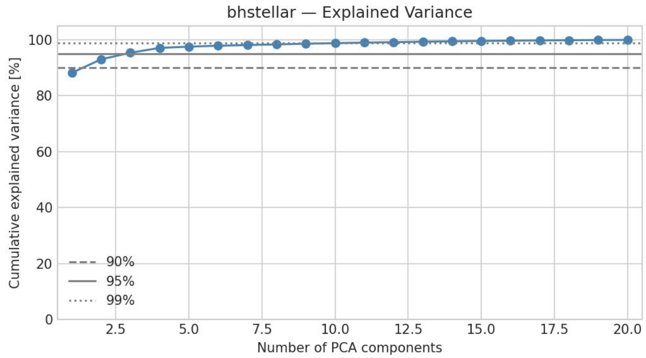

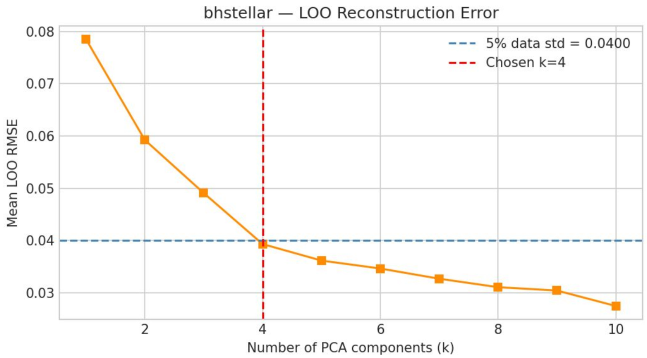

Principal component analysis and LOO validation: Stellar-to-halo mass relation

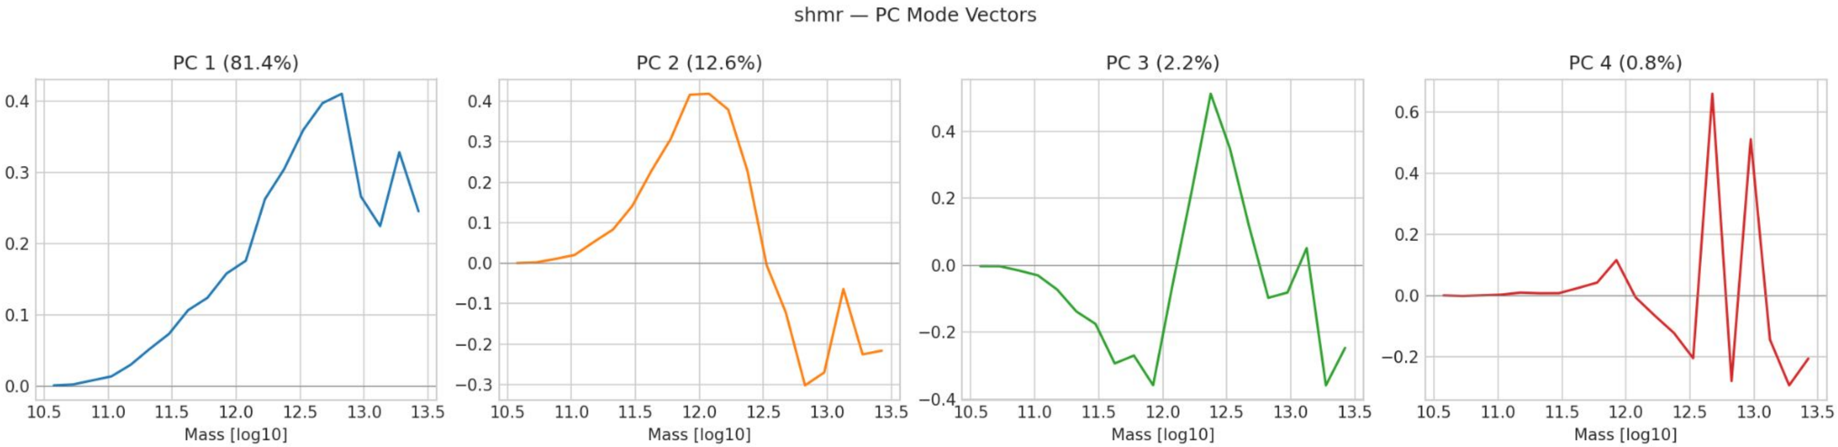

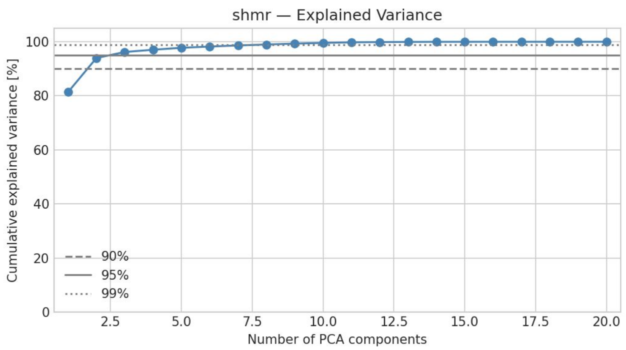

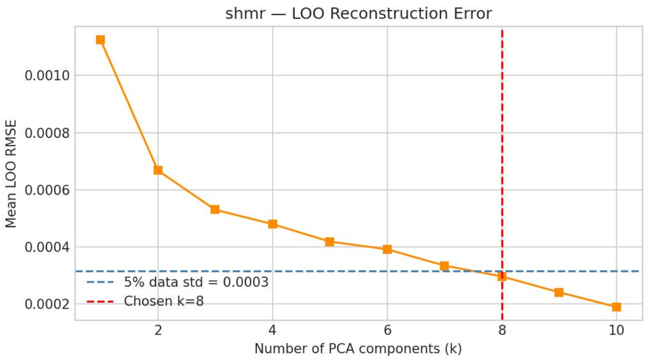

Gaussian process emulator: examples

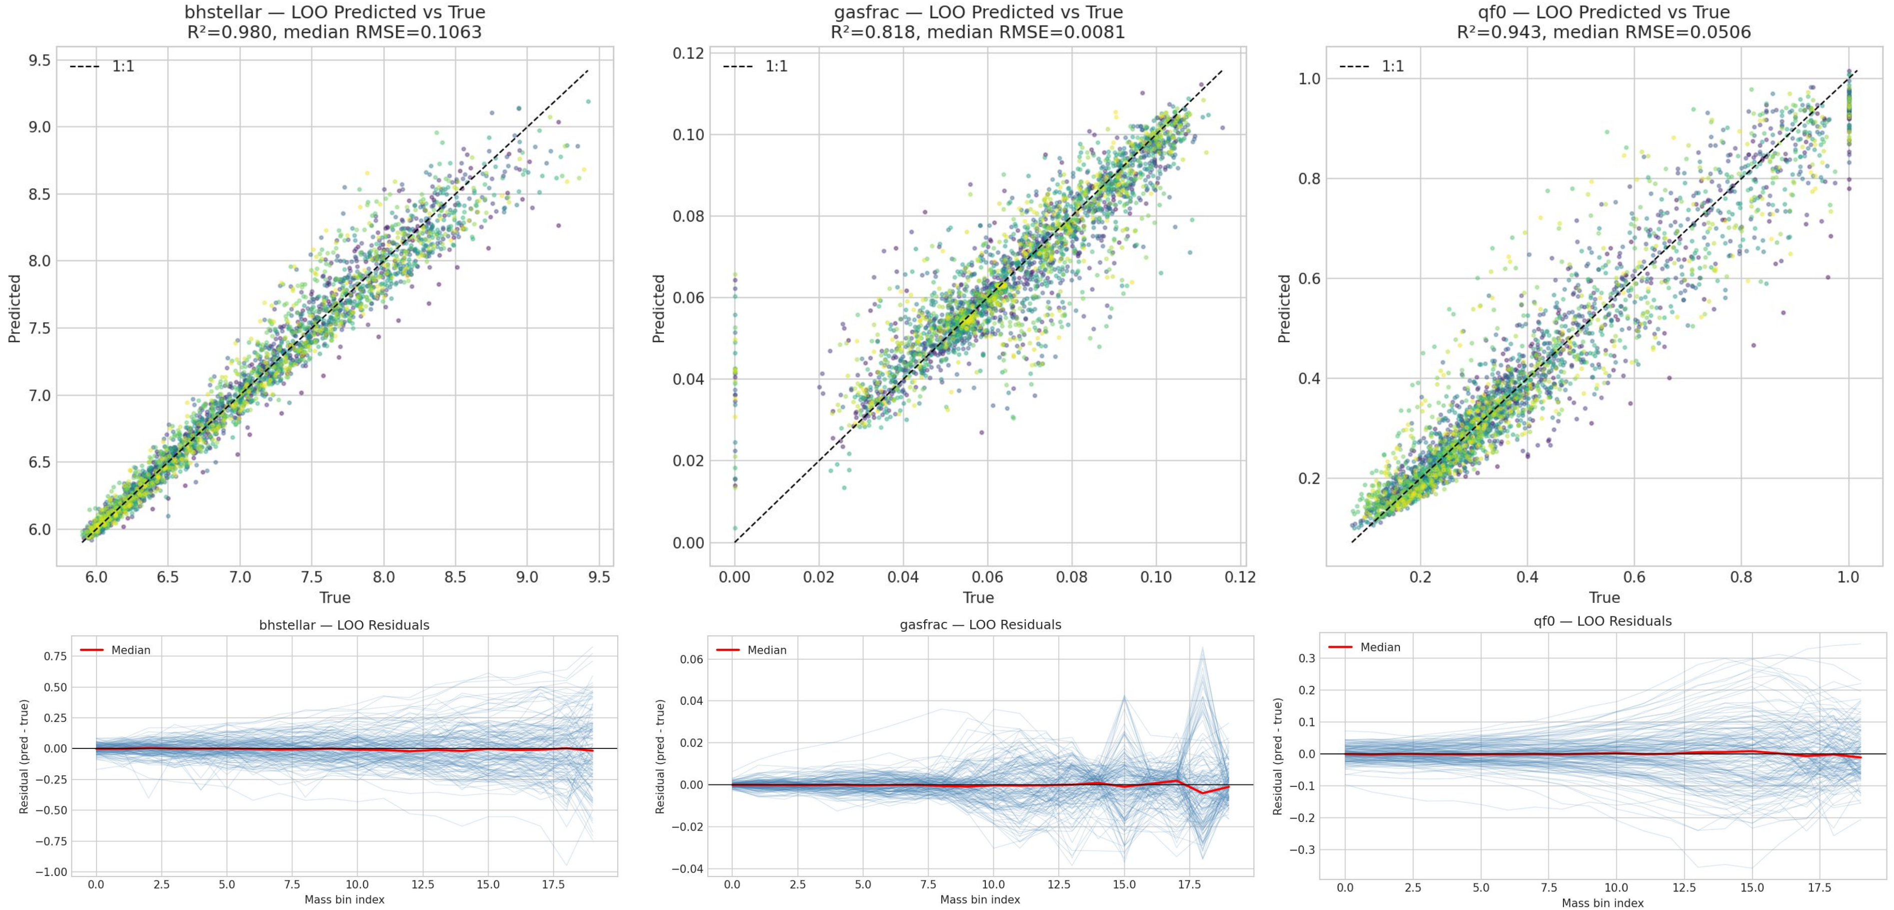

Sensitivity studies (Bondi rate)

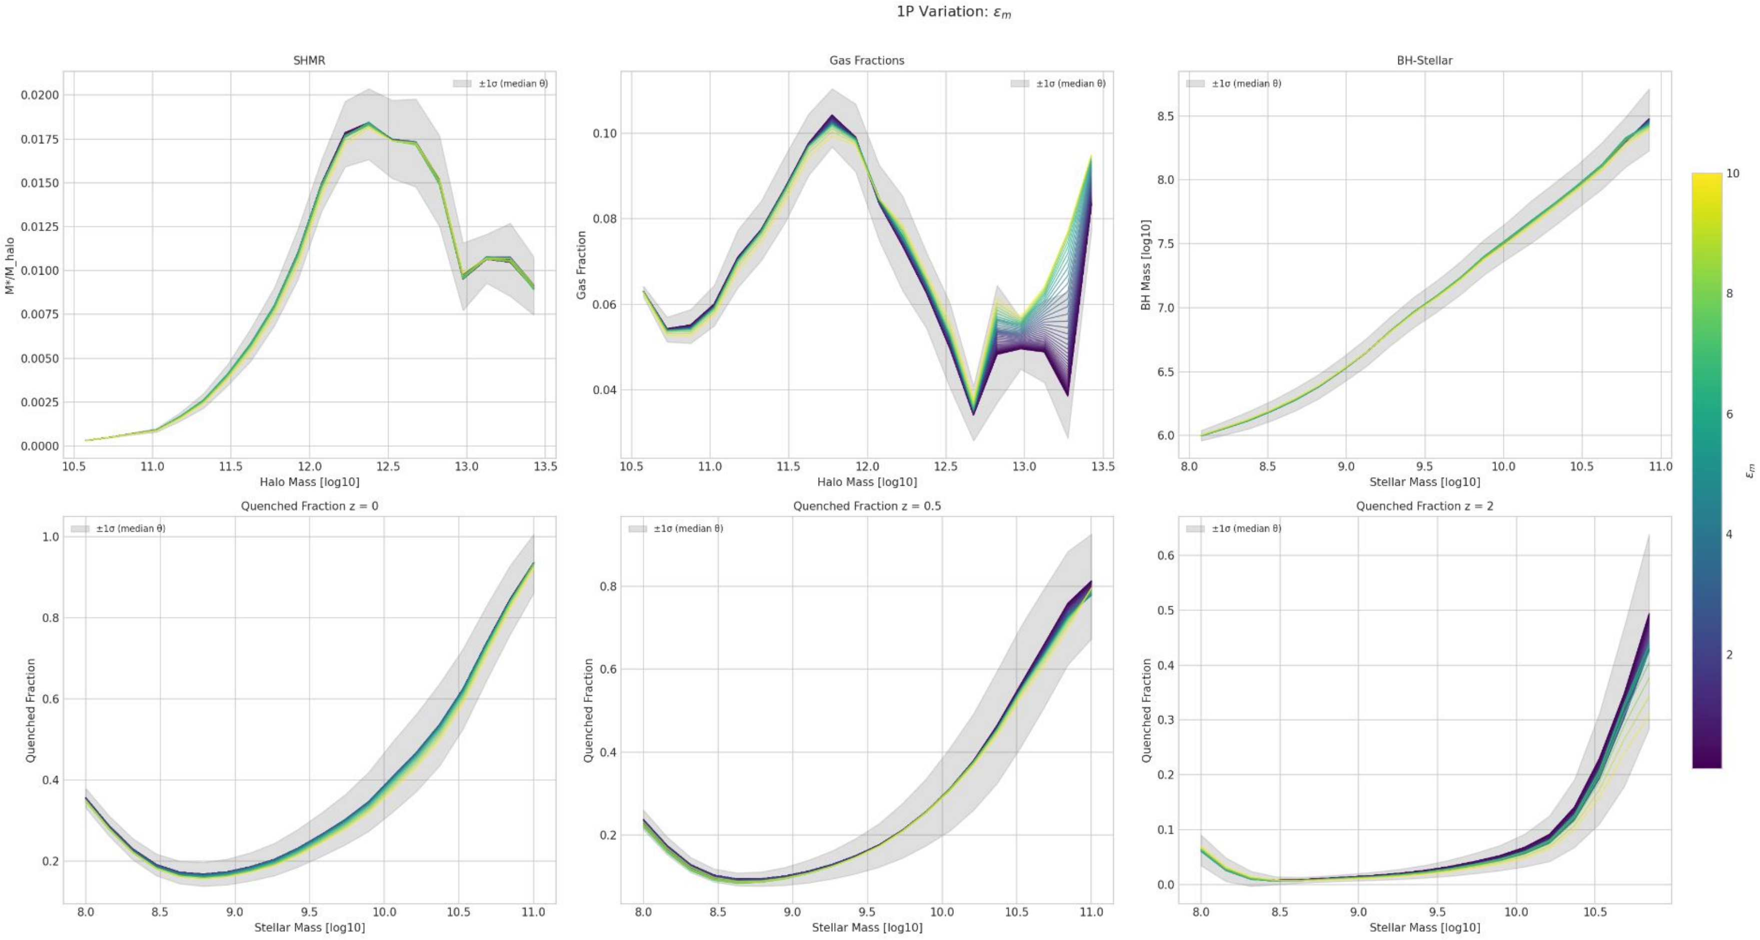

Sensitivity studies (TLA rate)

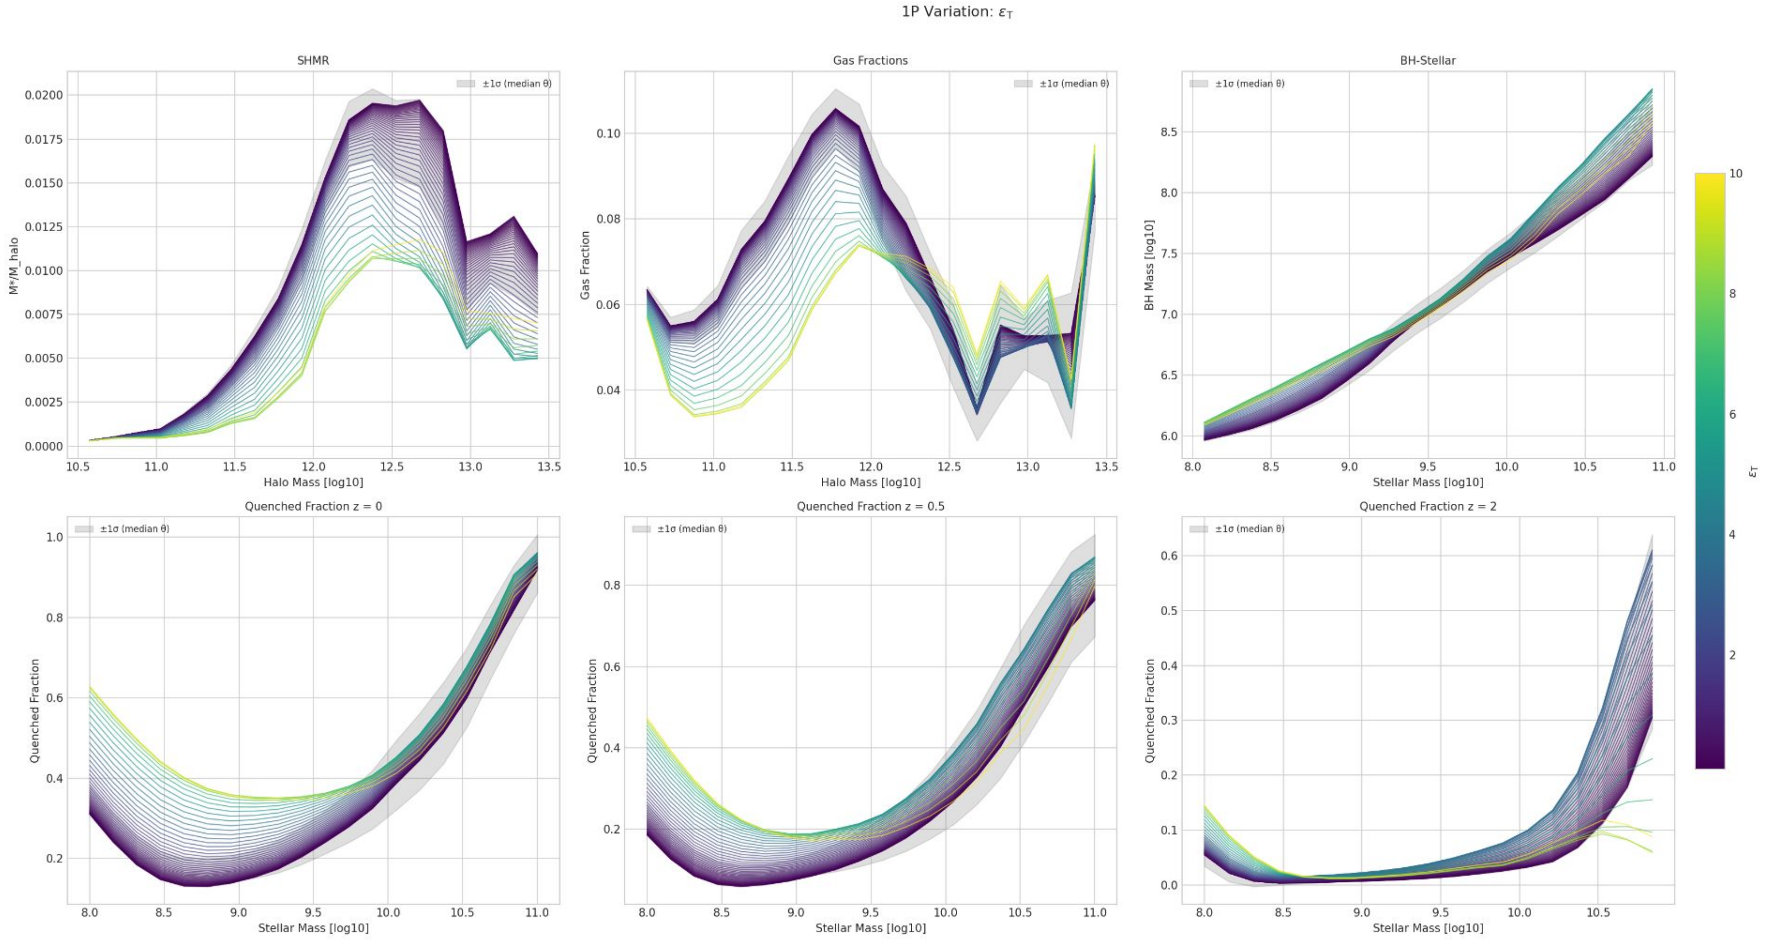

Sensitivity studies (Quasarthreshold)

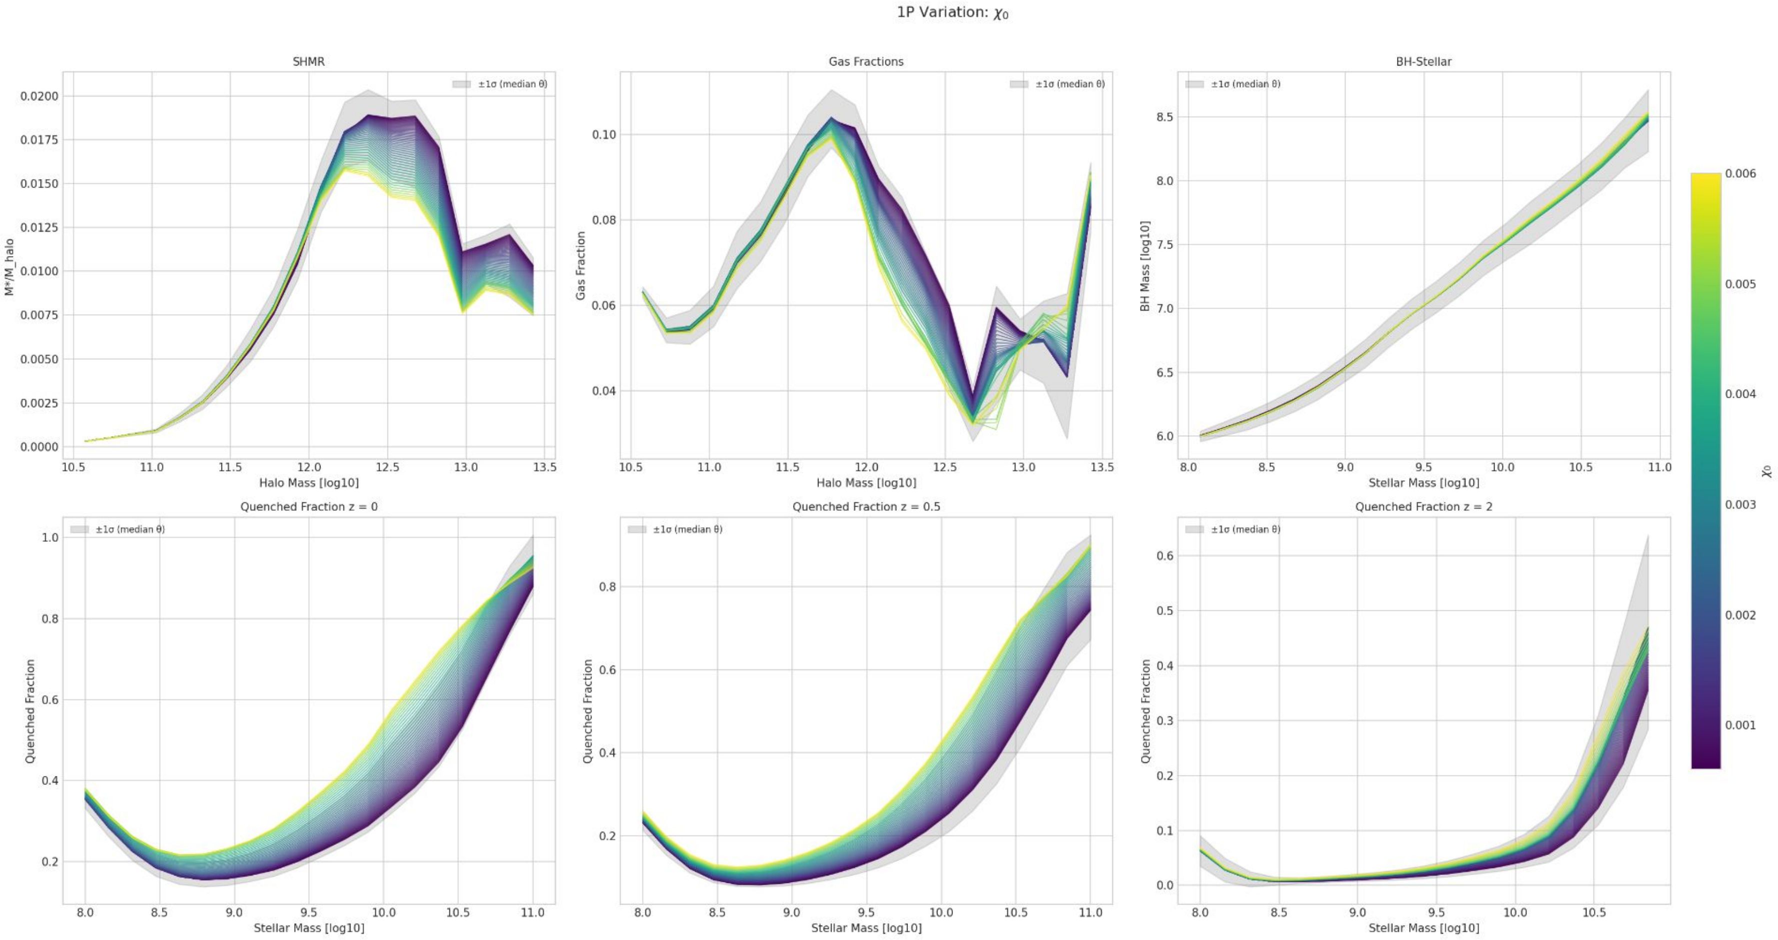

Sensitivity studies (Feedback factor)

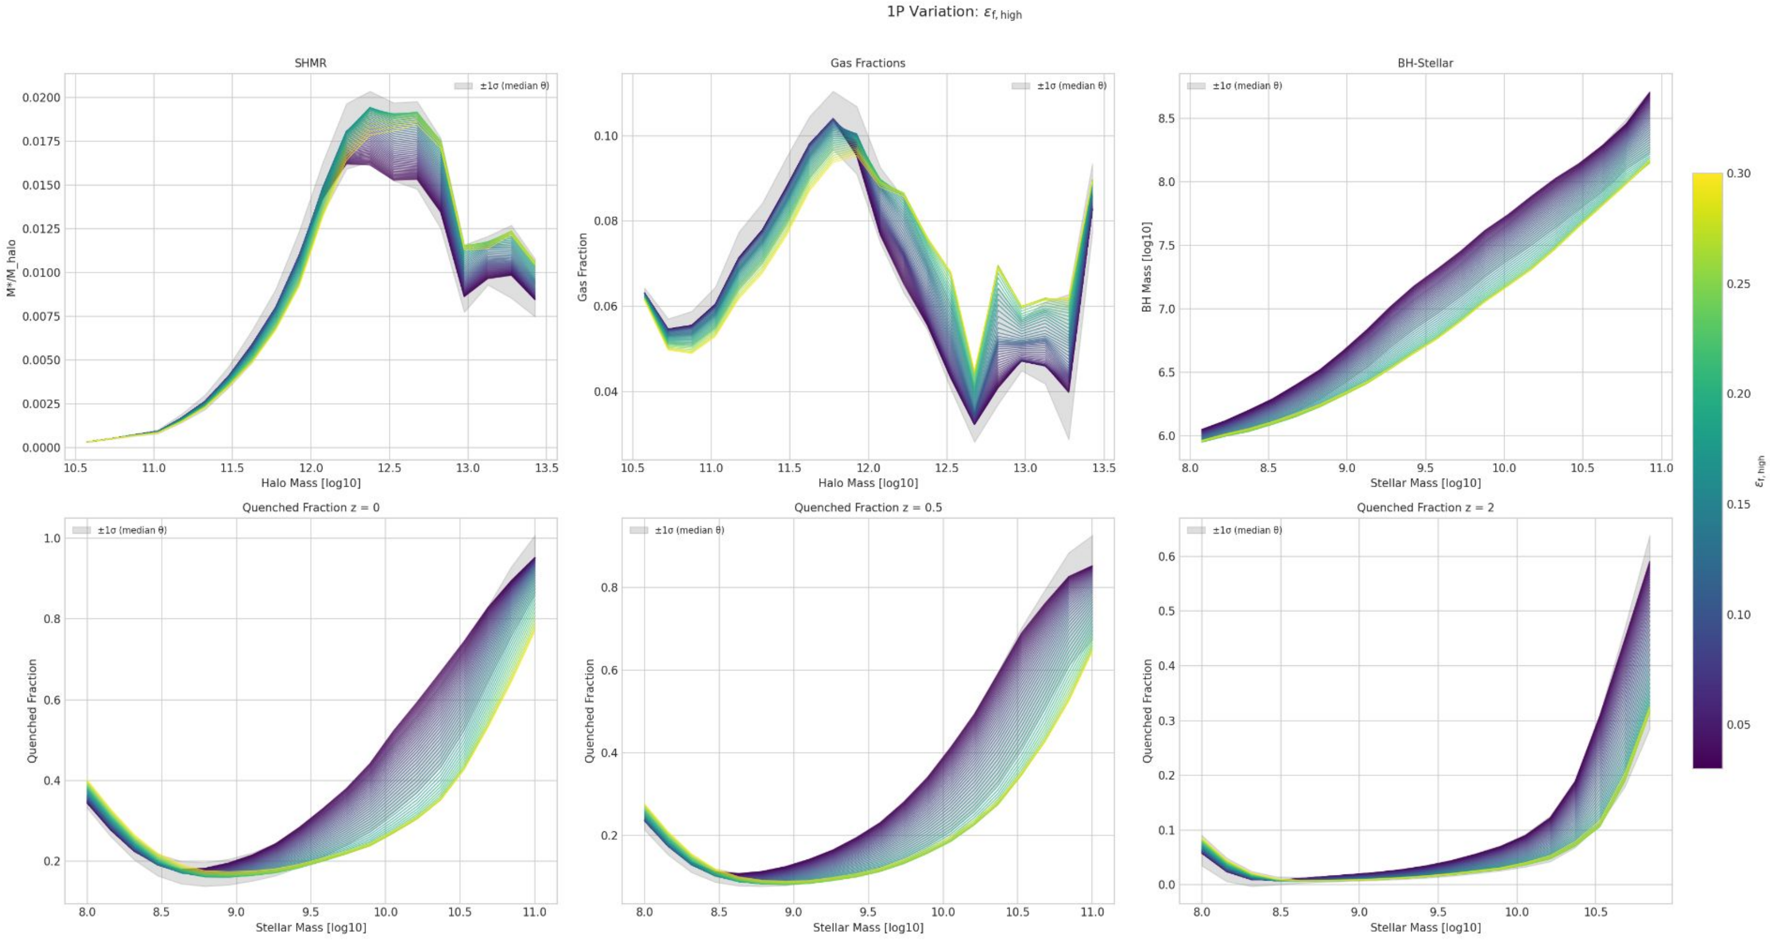

Sensitivity studies (Radio efficiency)

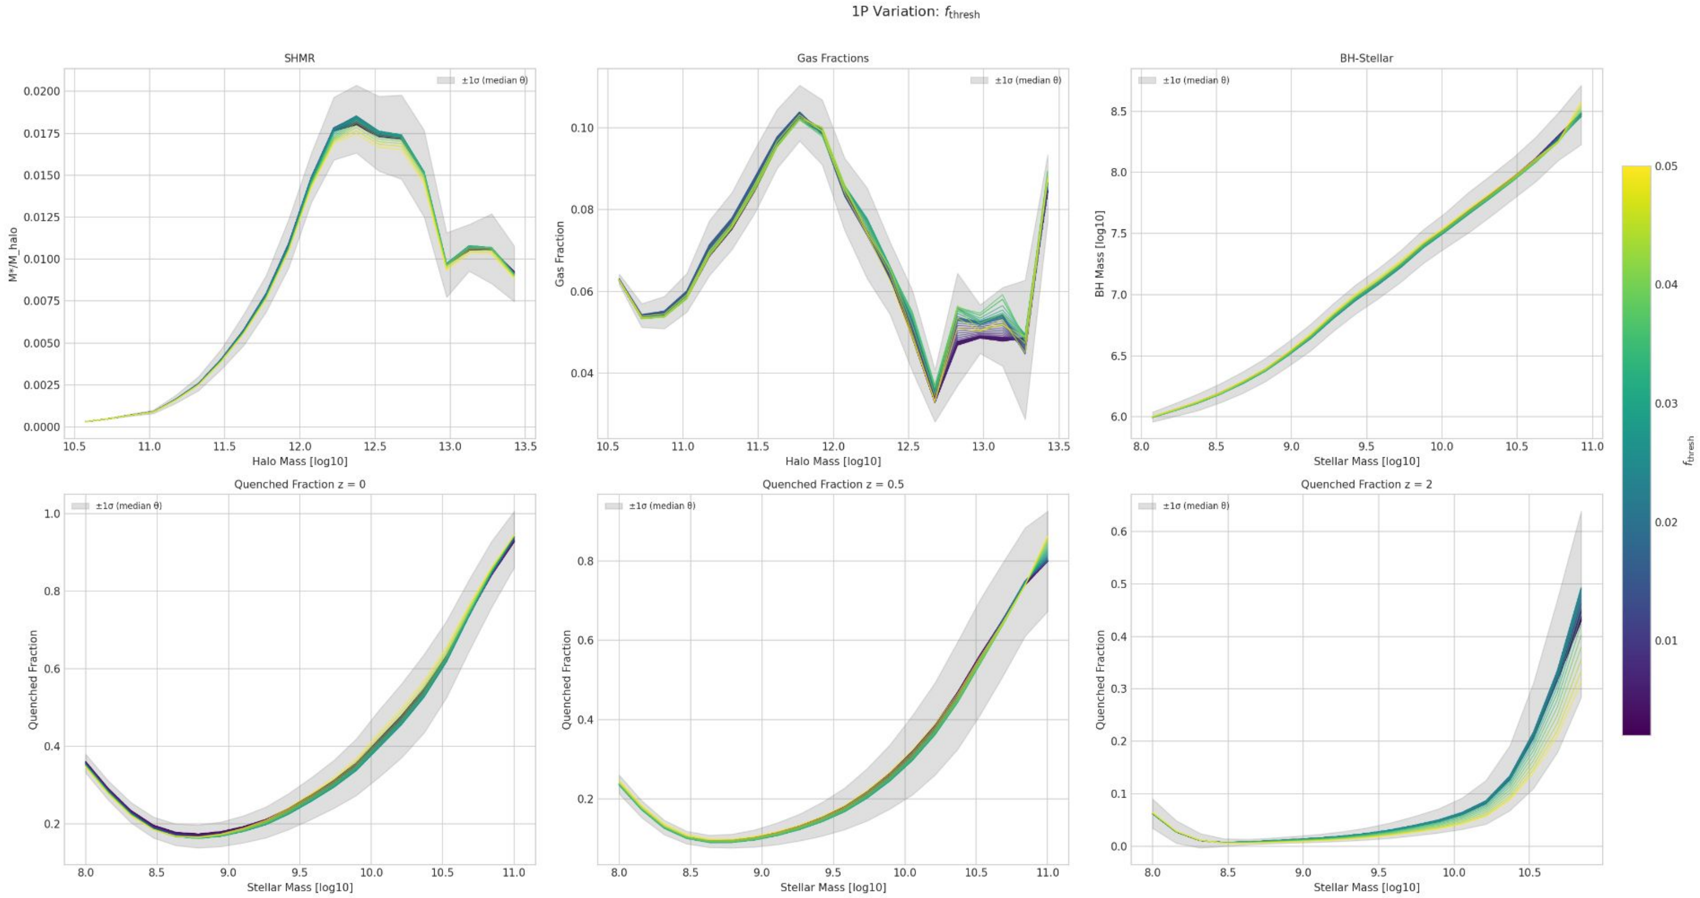

MCMC results: median curve

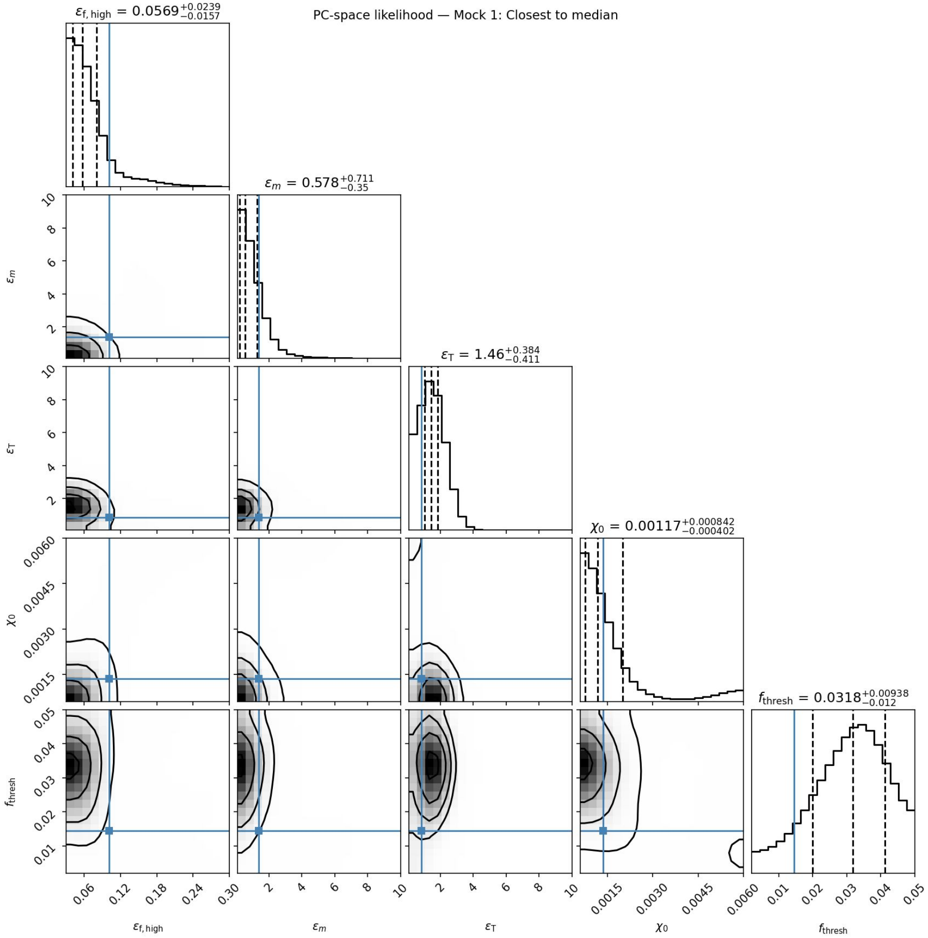

- Constraining parameters are generally well recovered
- Good constraints on the feedback factor and the accretion rates
- Slightly overconfident in the Bondi accretion rate, as this parameter hardly affects relations
- The posterior for the minimum density factor for radio feedback is very broad, as expected

MCMC results: smallest L2 from min

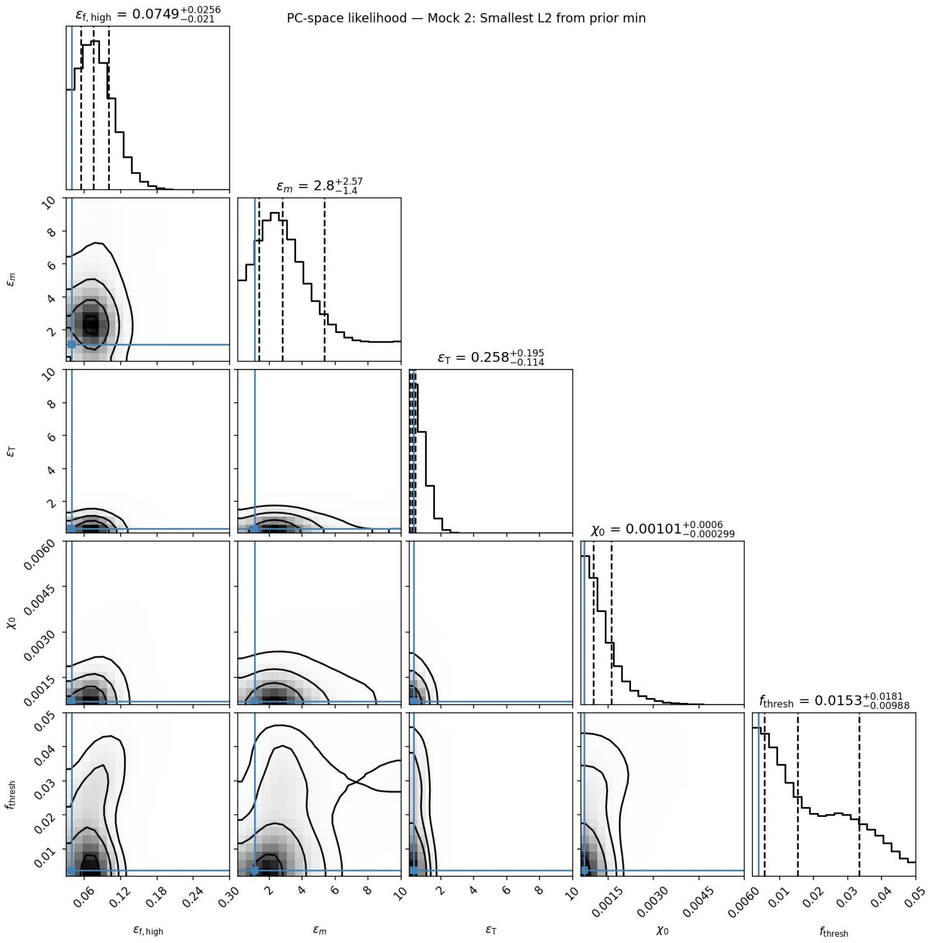

- All three crucial parameters are well recovered
- Broad posteriors for the two parameters that hardly affect the scaling relations, as expected
- Here, inference pipeline works remarkably well

MCMC results: largest L2 from min

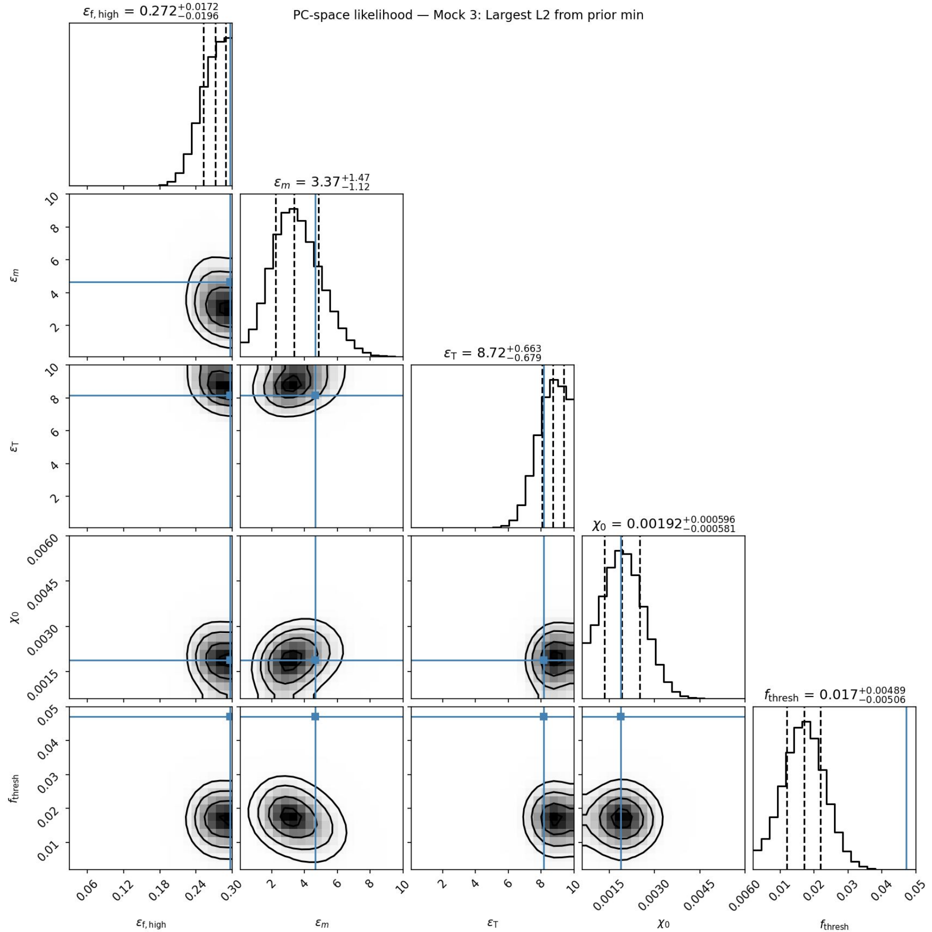

- All three crucial parameters are well recovered
- Broad posterior for the Bondi accretion rate
- Clear mismatch between actual and inferred parameter for the minimum density factor for radio feedback
    - There is an issue with the inference pipeline for what is essentially a nuisance parameter

Conclusions and future work

- We have built a simple inference pipeline that reliably recovers 3 out of 5 model parameters, based on 150 training simulations
- Parameters that do not affect the scaling relations used to train our emulator are not recovered
- To calibrate the full LtU model, we will use a DREAMS suite of zoom simulations
- Our results here raise some key questions:
    - How do we reliably identify inconsequential parameters?
    - How do we pick values for such parameters?
    - How robust do we think our inference methods are and can we address the remaining mismatches?