Extracting cosmological parameters from multi-wavelength cross-correlations on small scales

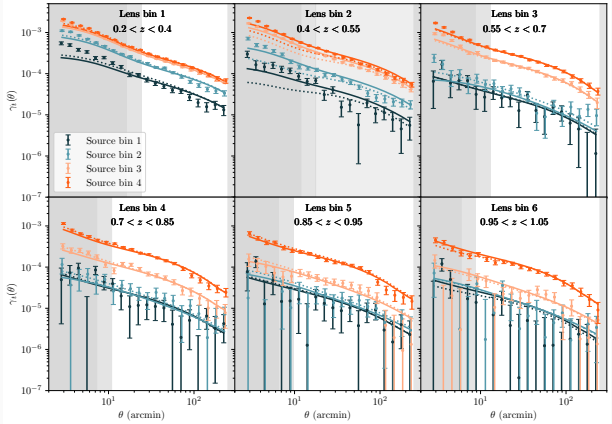
A differentiable cosmological halo model for galaxies, AGN, gas and matter

Johan Comparat

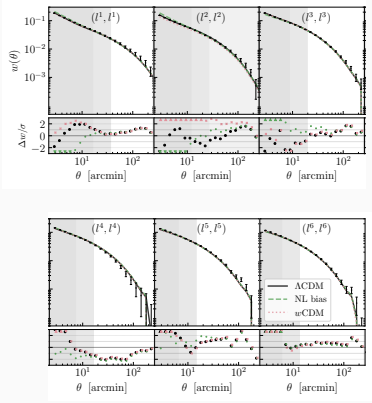
`ggah_mod`, `emu_pk`, `emu_hmf`

Université Grenoble Alpes, LPSC–IN2P3 (CNRS)

DES Y6 3×2pt: S/N = 302 on all scales, 95 after the scale cuts



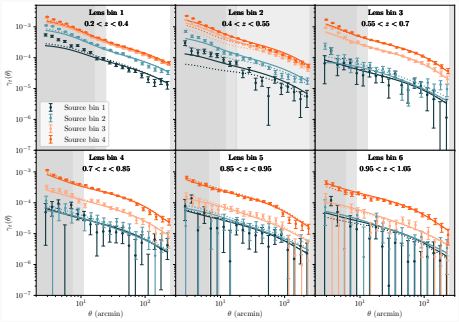
Giannini et al. 2026, arXiv:2601.15175, Fig. 3



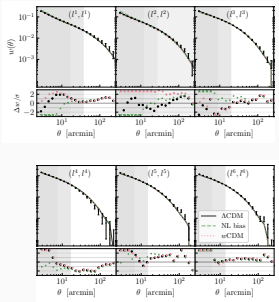
DES Collaboration, Abbott et al. 2026,

arXiv:2601.14559, Fig. 3

DES Y6 3×2pt: S/N = 302 on all scales, 95 after the scale cuts



Giannini et al. 2026, arXiv:2601.15175, Fig. 3



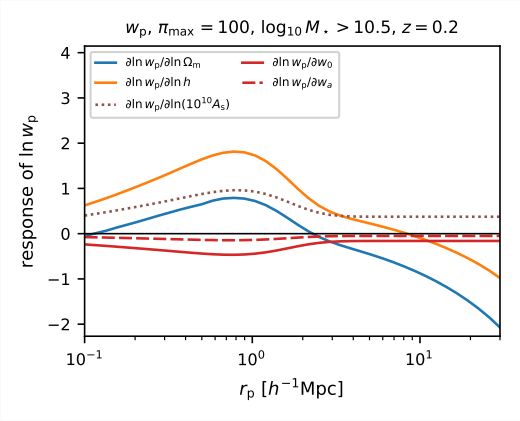
DES Collaboration, Abbott et al. 2026,

arXiv:2601.14559, Fig. 3

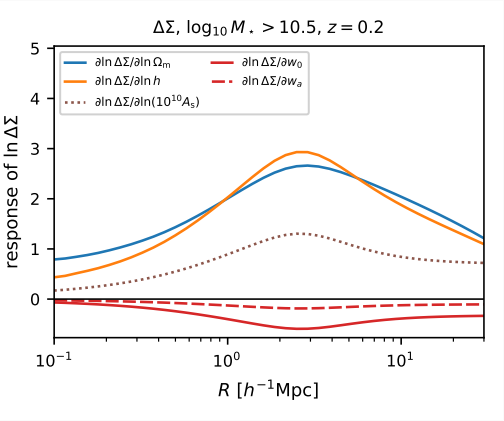
Is there cosmological information on small scales?

Can we retrieve it? Or is it buried in astrophysics?

How $w_{\rm p}$ and $\Delta\Sigma$ respond to the cosmology ?

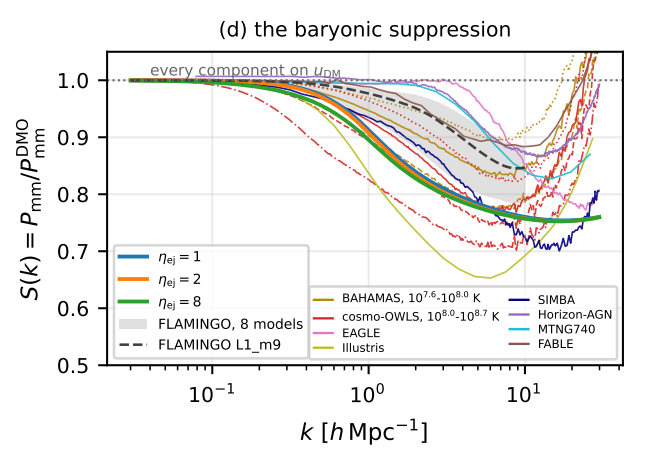


A 1% change in $\Omega_{\rm m}$ changes $w_{\rm p}(r_{\rm p})$ by $+0.67\%$, $+0.14\%$ and -0.86% at 0.5, 2 and 10 h^{-1}Mpc.



A 1% change in $\Omega_{\rm m}$ changes $\Delta\Sigma(R)$ by $+1.43\%$, $+2.58\%$ and $+2.05\%$ at 0.5, 2 and 10 h^{-1}Mpc.

Small scales carry cosmological information ... and the baryons



FLAMINGO at $z = 0$ (Schaller et al. 2025): spread of eight models and L1_m9. BAHAMAS (McCarthy et al. 2017), cosmo-OWLS (Le Brun et al. 2014), EAGLE (Schaye et al. 2015), Illustris (Vogelsberger et al. 2014), SIMBA (Davé et al. 2019) and Horizon-AGN (Dubois et al. 2014; Chisari et al. 2018), from the power-spectrum library of van Daalen et al. (2020); MTNG740 from Fig. 4 of Pakmor et al. (2023); FABLE (Henden et al. 2018) from Fig. 2 of Martin-Alvarez et al. (2025).

JAX 101 — autodifferentiation & JIT

JAX: NumPy-style array programming + composable transformations, same code on CPU / GPU.

Autodifferentiation — exact derivatives

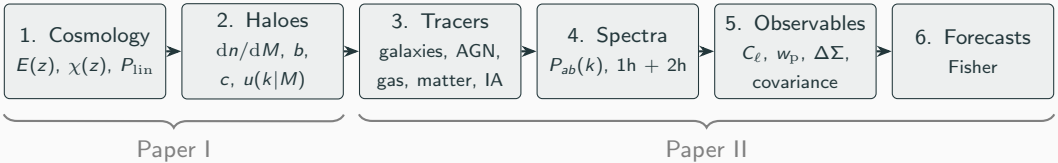
- a program is a chain of primitives (+, ×, `exp`, `erf`, `interp`, ...), each with a known derivative
- JAX differentiates the *program*, not the formula: the local derivatives are multiplied through the chain rule, step by step through the code ⇒ exact at machine precision
- no finite-difference step ϵ, no symbolic formula: only numbers flow, cost ≈ a few forward evaluations
- `J = jax.jacfwd(model)(theta)`: all directions in one call
- the price: every step must be JAX-traceable — CAMB is not, hence the emu_pk differentiable emulator for the linear power spectrum

JIT — compile once, run fast

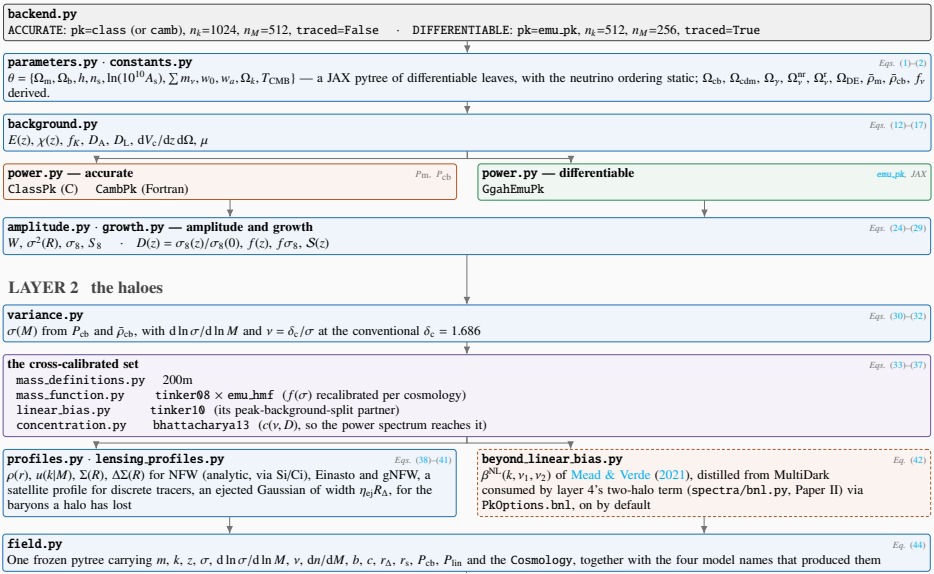
- `jax.jit` traces the function once, XLA compiles it into one fused kernel; later calls bypass Python entirely
- `vmap` batches loops: the many multipole Limber integral becomes a single XLA call ($\sim \times\mathbf{10}$)
- XLA releases the Python GIL ⇒ thread-parallel over redshift ($\sim \times\mathbf{5}$)
- $w_p(r_p)$ evaluates in $< 1\,\mathbf{s}$ post-JIT

Acronyms — **JAX**: Autograd + XLA, Google's differentiable array framework (Bradbury et al. 2018) • **XLA**: Accelerated Linear Algebra (its compiler) • **JIT**: Just-In-Time compilation • **GIL**: Python's Global Interpreter Lock • timings: hod_mod docs `timing_joint_model`.

One halo model in six layers



- Written in JAX: every prediction is differentiable in the cosmological and astrophysical parameters, so Fisher matrices come from exact Jacobians.
- Two flavours of the cosmology layer: accurate (CLASS, CAMB) for fitting, differentiable (`emu_pk`) for forecasting.
- A complete population of haloes with one baryon budget.


LAYER 1 the cosmology
backend.py
ACCURATE: pk=class (or camb), n_k=1024, n_M=512, traced=False · DIFFERENTIABLE: pk=emu_pk, n_k=512, n_M=256, traced=True
parameters.py · constants.py
Eqs. (1)–(2)
$\theta = \{\Omega_m, \Omega_b, h, n_s, \ln(10^{10}A_s), \sum m_\nu, w_0, w_a, \Omega_k, T_{CMB}\}$ — a JAX pytree of differentiable leaves, with the neutrino ordering static; Ω_{cb}, Ω_{cdm}, Ω_γ, Ω_ν^{nr}, Ω_ν^r, Ω_{DE}, $\bar{\rho}_m$, $\bar{\rho}_{cb}$, f_ν derived.
background.py
Eqs. (12)–(17)
$E(z)$, $\chi(z)$, f_K, D_A, D_L, $dV_c/dz\,d\Omega$, μ
power.py — accurate
P_m, P_{cb}
ClassPk (C) CambPk (Fortran)
power.py — differentiable
emu_pk, JAX
GgahEmuPk
amplitude.py · growth.py — amplitude and growth
Eqs. (24)–(29)
W, $\sigma^2(R)$, σ_8, S_8 · $D(z) = \sigma_8(z)/\sigma_8(0)$, $f(z)$, $f\sigma_8$, $\mathcal{S}(z)$
LAYER 2 the haloes
variance.py
Eqs. (30)–(32)
$\sigma(M)$ from P_{cb} and $\bar{\rho}_{cb}$, with $d\ln\sigma/d\ln M$ and $\nu = \delta_c/\sigma$ at the conventional $\delta_c = 1.686$
the cross-calibrated set
Eqs. (33)–(37)
mass_definitions.py 200m
mass_function.py tinker08 × emu_hmf ($f(\sigma)$ recalibrated per cosmology)
linear_bias.py tinker10 (its peak-background-split partner)
concentration.py bhattacharya13 ($c(\nu, D)$, so the power spectrum reaches it)
profiles.py · lensing_profiles.py
Eqs. (38)–(41)
$\rho(r)$, $u(k|M)$, $\Sigma(R)$, $\Delta\Sigma(R)$ for NFW (analytic, via Si/Ci), Einasto and gNFW, a satellite profile for discrete tracers, an ejected Gaussian of width $\eta_{ej}R_\Delta$, for the baryons a halo has lost
beyond_linear_bias.py
Eq. (42)
$\beta^{NL}(k, \nu_1, \nu_2)$ of Mead & Verde (2021), distilled from MultiDark consumed by layer 4's two-halo term (spectra/bnl.py, Paper II) via PkOptions.bnl, on by default
field.py
Eq. (44)
One frozen pytree carrying m, k, z, σ, $d\ln\sigma/d\ln M$, ν, dn/dM, b, c, r_Δ, r_s, P_{cb}, P_{lin} and the Cosmology, together with the four model names that produced them

emu_pk: a differentiable linear power spectrum

parameter	range
$\omega_{\rm b}$	0.0170 … 0.0280
$\omega_{\rm cdm}$	0.0500 … 0.3000
h	0.5500 … 0.8500
$n_{\rm s}$	0.8400 … 1.1000
$\ln(10^{10}A_{\rm s})$	1.6100 … 4.0000
$\sum m_\nu$ [eV]	0.0000 … 0.6000
w_0	−1.5000 … −0.5000
w_a	−1.0000 … 0.6000
Ω_k	−0.1500 … 0.1500
r_1	0.0000 … 0.3333
r_2	0.0000 … 0.5000
k [$h\,\mathrm{Mpc}^{-1}$]	10^{-4} … 300
z	0 … 5

$m_i = r_i \sum m_\nu$, $r_3 = 1 - r_1 - r_2$

Error in $P_{\rm m}$ against converged CAMB, 32 held-out cosmologies, $z = 0$

	median	90th	max
$k \in [10^{-3}, 10]\,h\,\mathrm{Mpc}^{-1}$	**0.063%**	0.130%	0.190%
$k \in [10, 300]\,h\,\mathrm{Mpc}^{-1}$	**0.044%**	0.069%	0.086%

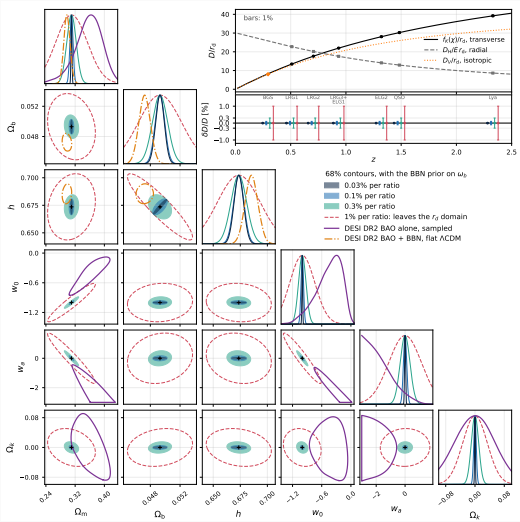
Error in $\partial \ln P / \partial\theta$ (AD against finite differences), $z = 0$

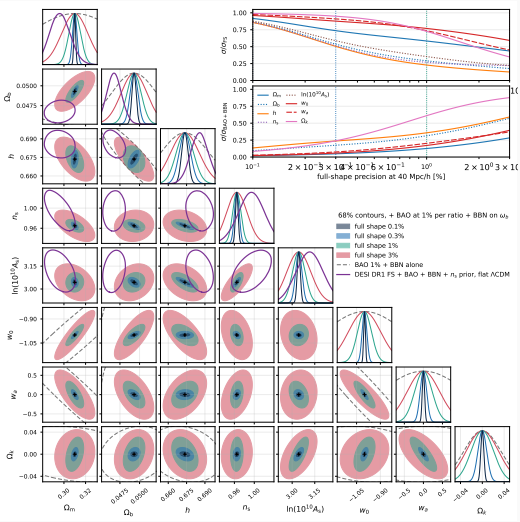
parameter	error	floor	parameter	error	floor
$\ln(10^{10}A_{\rm s})$	**exact**	–	$\sum m_\nu$	0.234%	0.010%
$n_{\rm s}$	**exact**	–	w_a	0.295%	0.040%
$\omega_{\rm cdm}$	0.048%	0.038%	Ω_k	0.168%	0.002%
h	0.064%	0.005%	r_1	3.93%	0.248%
w_0	0.135%	0.036%	r_2	4.46%	0.179%
$\omega_{\rm b}$	0.124%	0.006%			

Network architecture after CosmoPower (Spurio Mancini et al. 2022, MNRAS 511, 1771).

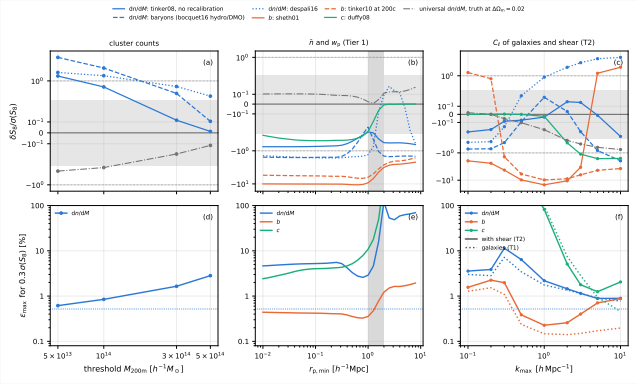
Documentation: `emu-pk.readthedocs.io` Install: `pip install emu_pk`

Fisher forecasts: BAO and Full-shape





Accuracy required of the halo layer to model number density, clustering and lensing

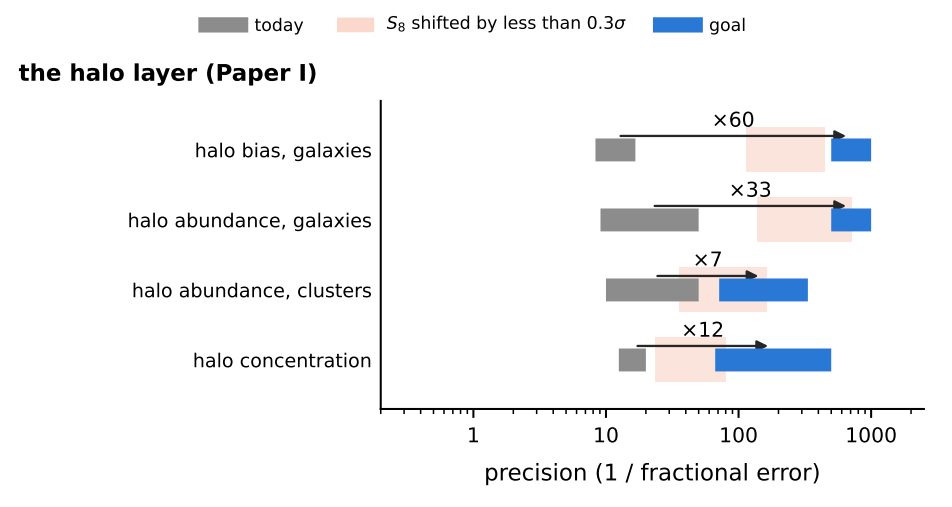


The error on each ingredient that shifts S_8 by 0.3σ (Fisher bias):

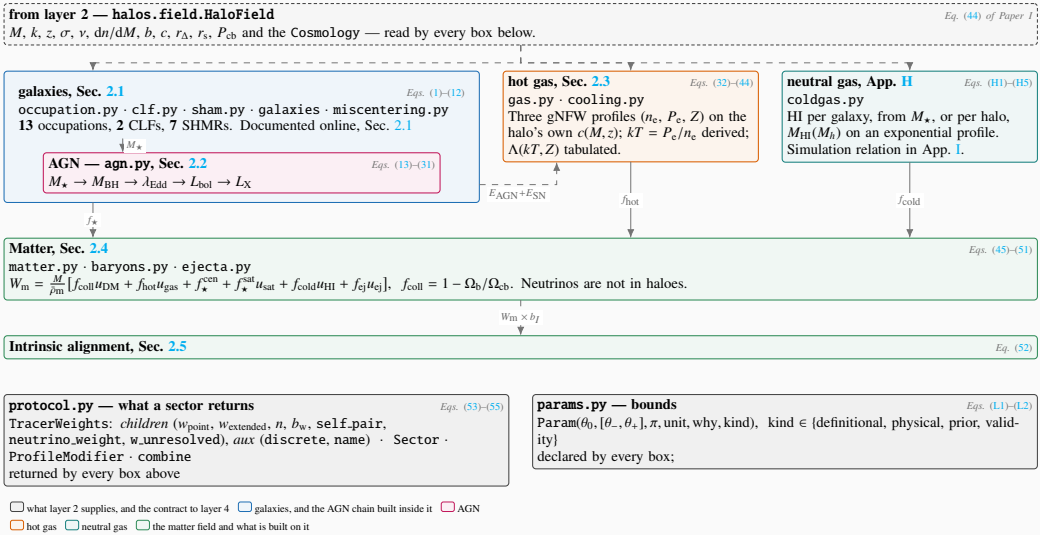
- $\mathrm{d}n/\mathrm{d}M$ to 0.62–2.8% for the counts;
- $b(M)$ to 0.33–0.44% for $\bar{n}$ and w_p, 0.23–0.88% for the angular spectra;
- the published fits differ by up to 12% in b and 2–8% in $\mathrm{d}n/\mathrm{d}M$.

The halo layer, not the data, limits these constraints.

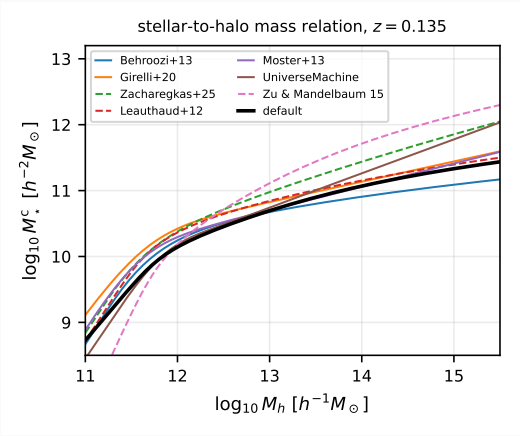
Requirements on the model

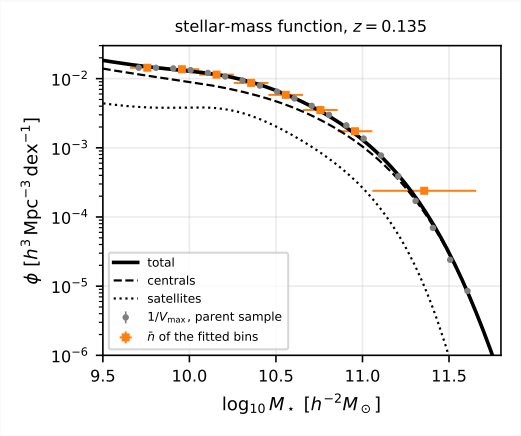


Gas, galaxies, AGN

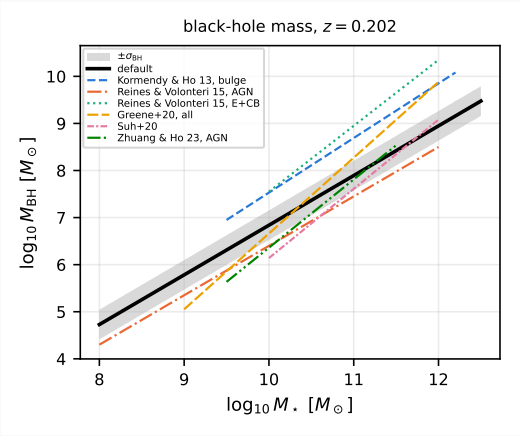


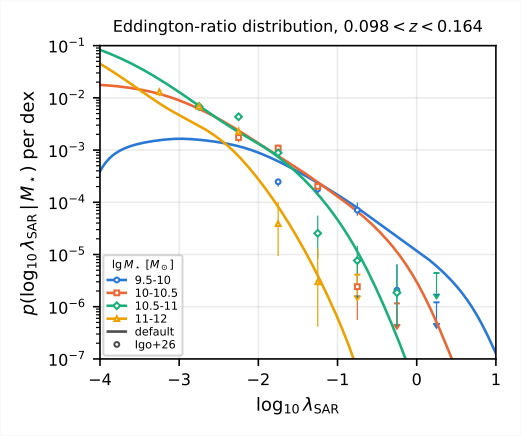
The default galaxy model: Zu & Mandelbaum (2015, 2016)



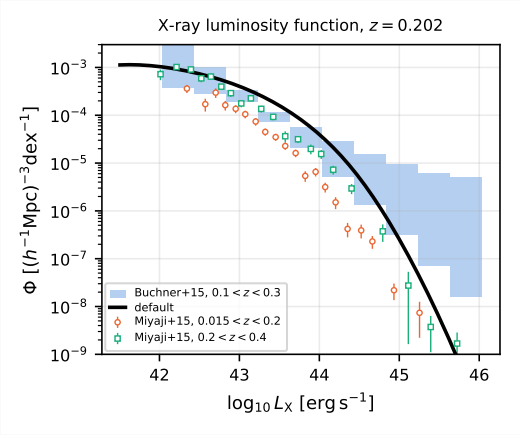


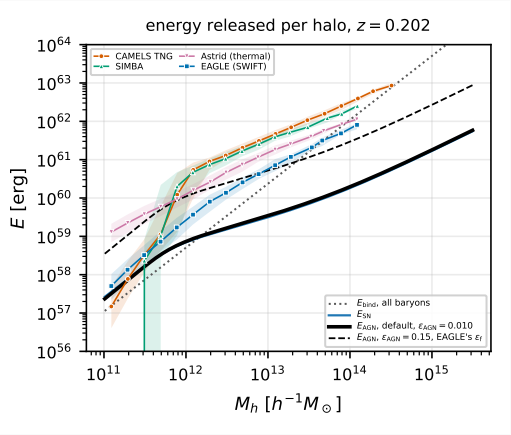
The AGN model



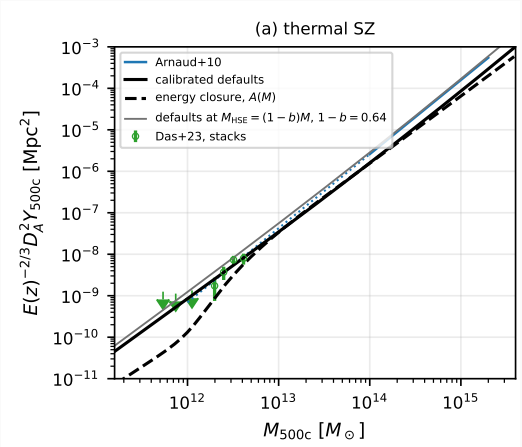


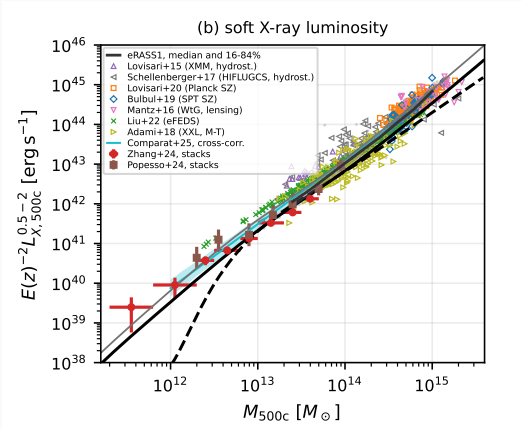
The AGN model



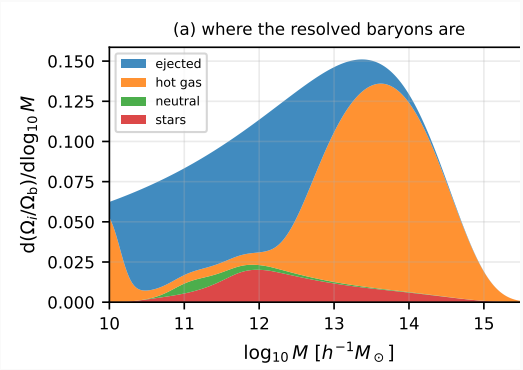


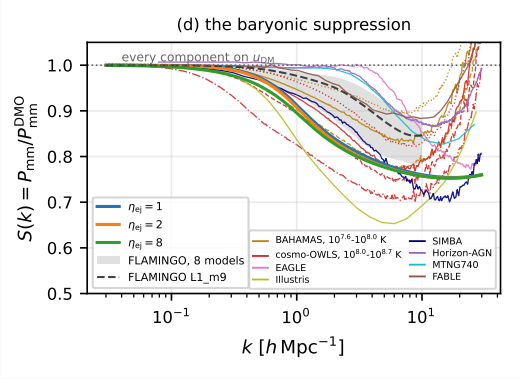
The gas model



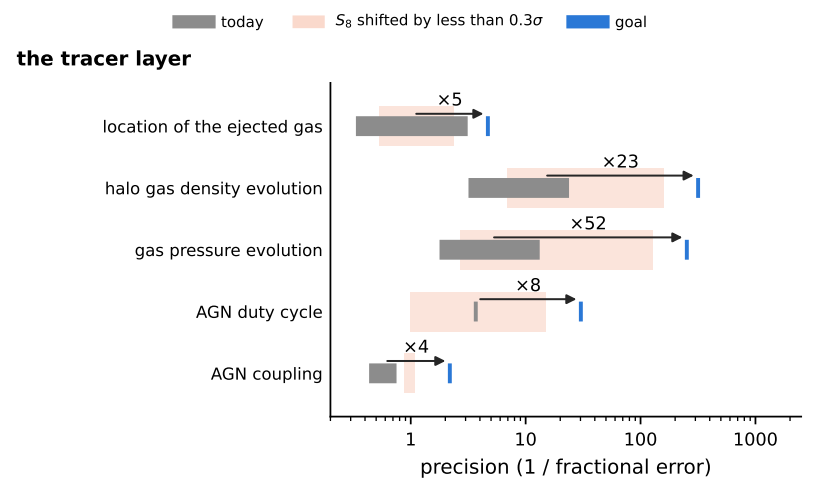


The gas model

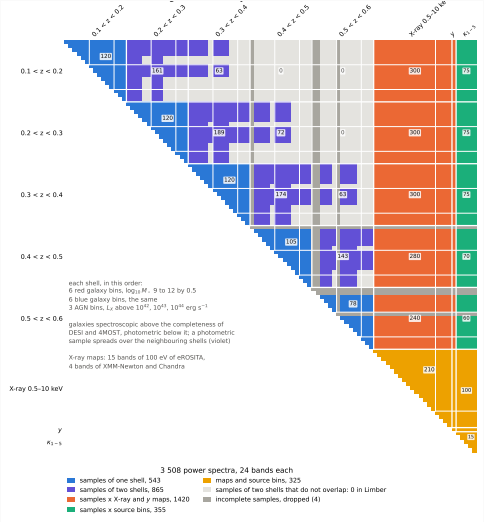




Requirements on the model

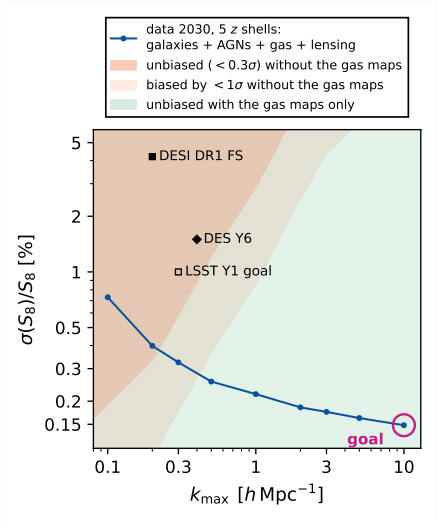


A data vector of 2031



- Five shells, $0.1 < z < 0.6$: red and blue galaxies in six $\lg M_\star$ bins, AGN in three L_X bins; spectroscopic above the completeness of DESI and 4MOST, photometric below.
- X-ray maps in 19 bands (eROSITA, XMM-Newton, Chandra), the y map of the Simons Observatory, five source bins of LSST + Euclid.
- 100 tracers, 3 508 power spectra in 24 bands of $20 < \ell < 10^4$; 212 parameters, all galaxy and AGN parameters free in each shell.
- Gaussian and super-sample covariance.

Forecast precision on S_8

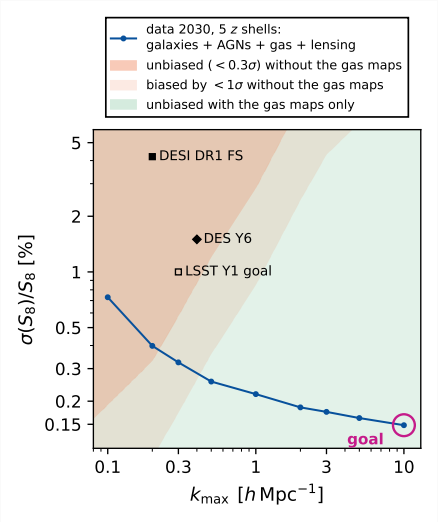


$10^3\,\sigma(S_8)/S_8$ at $k_{\rm max} = 10\,h\,{\rm Mpc}^{-1}$:

data of 2031	one shell	five shells
collaborations	2.06	1.48
public	2.39	1.83

- At the scale cuts of today (0.2–$0.4\,h\,{\rm Mpc}^{-1}$): 3.98–2.83.
- Without the maps, the gas recipe biases S_8 by 0.3σ from $k_{\rm max} = 0.28\,h\,{\rm Mpc}^{-1}$; beyond it, the maps cut the error by 2.4.

Summary



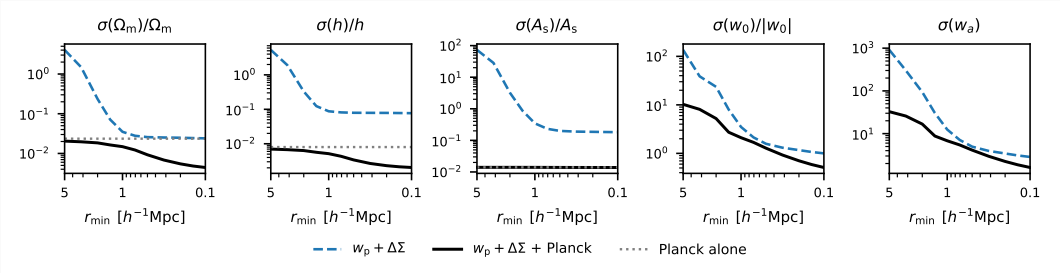
- One differentiable halo model for galaxies, AGN, hot and neutral gas and matter, with one baryon budget; every auto- and cross-spectrum and its covariance.
- The gain comes from constraining feedback with the data instead of marginalising over it.

`github.com/JohanComparat/ggah_mod`

`pip install ggah-mod emu_pk emu_hmf`

Backup slides

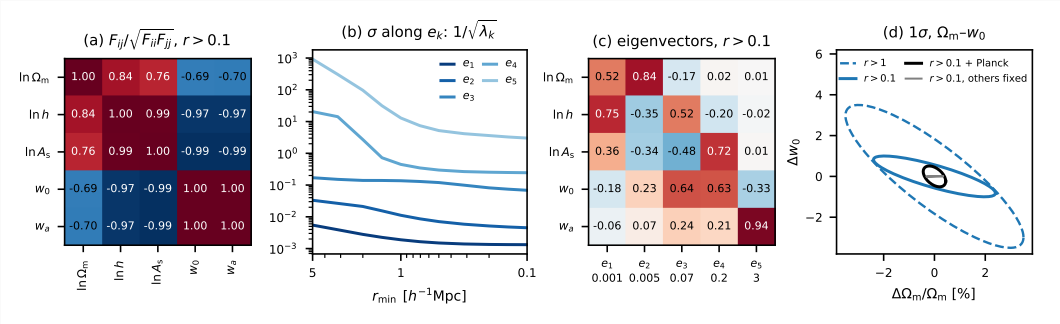
What the small scales add to cosmological parameters determination



Fisher forecast of the joint fit, 1σ against the scale cut $r > r_{\rm min}$:

- $w_{\rm p}$ ($\pi_{\rm max} = 100\,h^{-1}{\rm Mpc}$) and $\Delta\Sigma$ at $z = 0.2$, 15 bins from 0.1 to $30\,h^{-1}{\rm Mpc}$, the same cut on both.
- Lenses: LS10, $\log_{10} M_\star > 10.5$, $0.05 < z < 0.26$, its $\bar{n}$; sources DES Y3-like ($\sigma_e = 0.26$, $n_{\rm eff} = 5.6\,{\rm arcmin}^{-2}$, $z_s = 0.65$) on $5000\,{\rm deg}^2$.
- Model: `ggah_mod` defaults without the beyond-linear bias, differentiable backend; galaxy and gas parameters fixed.
- Gaussian covariance with the $w_{\rm p} \times \Delta\Sigma$ block; no trispectrum, no super-sample term.
- Planck 2018 TT,TE,EE+lowE+lensing on $\Omega_{\rm m}$, h and $A_{\rm s}$; no prior on w_0, w_a.

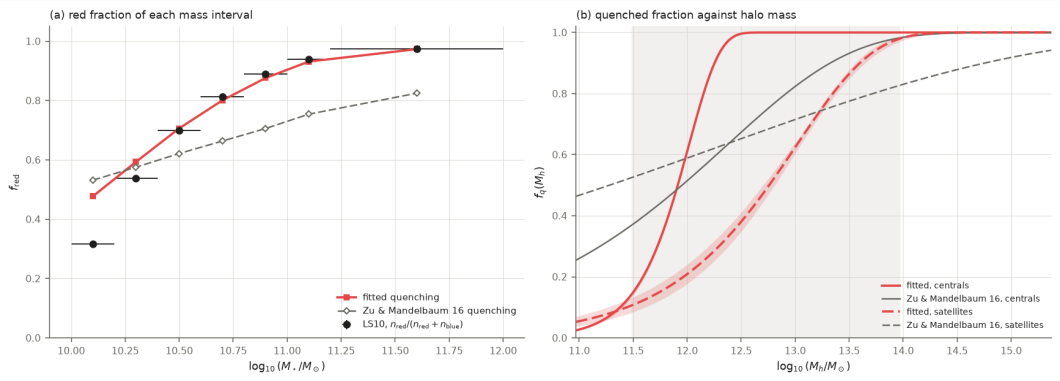
Why the small scales plateau: the Fisher matrix



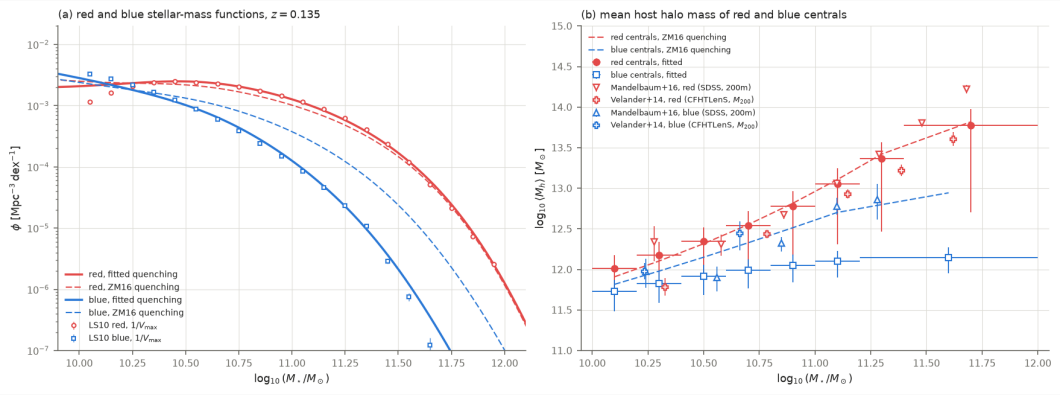
The joint fit of the previous slide, in $(\ln\Omega_{\rm m}, \ln h, \ln A_{\rm s}, w_0, w_a)$, galaxy parameters fixed:

- (a) At $r > 0.1\,h^{-1}\mathrm{Mpc}$ the five derivatives are nearly parallel: $|F_{ij}|/\sqrt{F_{ii}F_{jj}} \geq 0.97$ between h, $A_{\rm s}$, w_0 and w_a.
- (b) Lowering $r_{\rm min}$ keeps shrinking the error along the two best directions; the three weak ones flatten below $1\,h^{-1}\mathrm{Mpc}$.
- (c) e_1 is an amplitude, $\Omega_{\rm m}^{0.5}h^{0.75}A_{\rm s}^{0.36}$, the abundance of the host halos; e_2 is mostly $\Omega_{\rm m}$; e_5 is w_a alone.
- (d) With the others fixed, $\sigma(\Omega_{\rm m})/\Omega_{\rm m} = 0.2\%$; marginalised, 2.4%. Planck pins the weak directions and turns the small-scale amplitude into a gain on each parameter.

Quenching and red fractions



Quenching and red fractions



Energy from AGN

