How Massive Can Galaxies Get in ΛCDM?

Empirical estimates from Extreme Value Statistics

Miguel Enríquez-Vargas

with A. Rodríguez-Puebla, A. Manuwal, L. Y. A. Yung, V. Avila-Reese & C. Cannarozzo

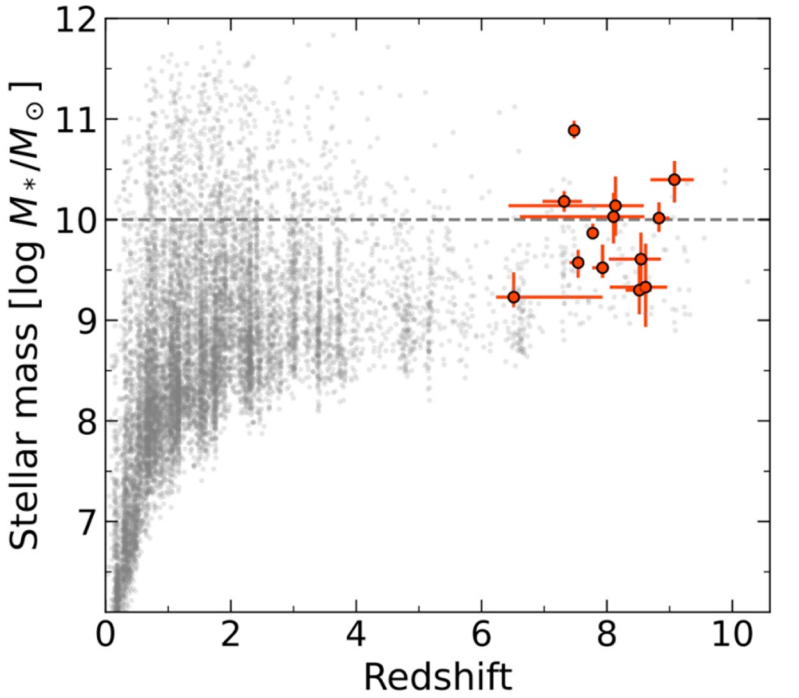
Instituto de Astronomía, UNAM

Galaxies That “Shouldn't Exist”

Labbé et al. (2023) identified six candidate galaxies in JWST/CEERS imaging at redshifts 7.4–9.1 (only 500–700 Myr after the Big Bang!) with stellar masses inferred from SED fitting above 10^{10} $M_{\odot}$, one approaching 10^{11}.

Their combined stellar-mass density far exceeded pre-JWST expectations, and the two most massive would require converting almost all available baryons into stars.

The crucial caveat: these are photometric, SED-derived masses, not spectroscopic measurements.



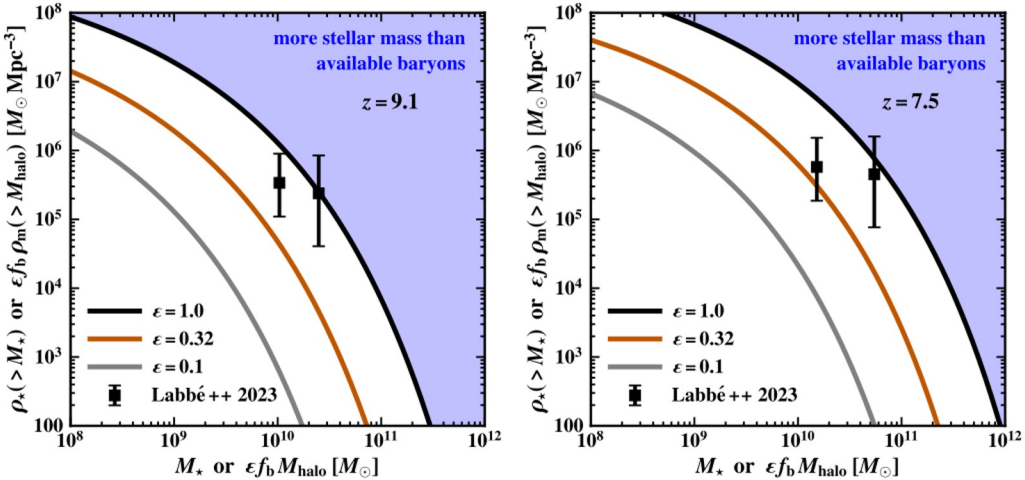
Labbé et al. (2023)

The Hard Ceiling

The panels show the comoving stellar-mass density locked in galaxies more massive than M_* at $z \approx 9.1$ and $z \approx 7.5$, compared with the maximum ΛCDM allows for three conversion efficiencies, $\varepsilon = 0.1$, 0.32 and 1.0, the last being the absolute baryonic limit, $M_* = f_b \cdot M_{halo}$.

At $z \approx 9$ the measured density brushes the $\varepsilon = 1$ line: essentially every baryon would have to become a star. For context, even the most efficient galaxies we know, in Milky Way-like halos, convert only ~20–30%.

Something has to give: the measurements, the galaxy physics, or the cosmology.



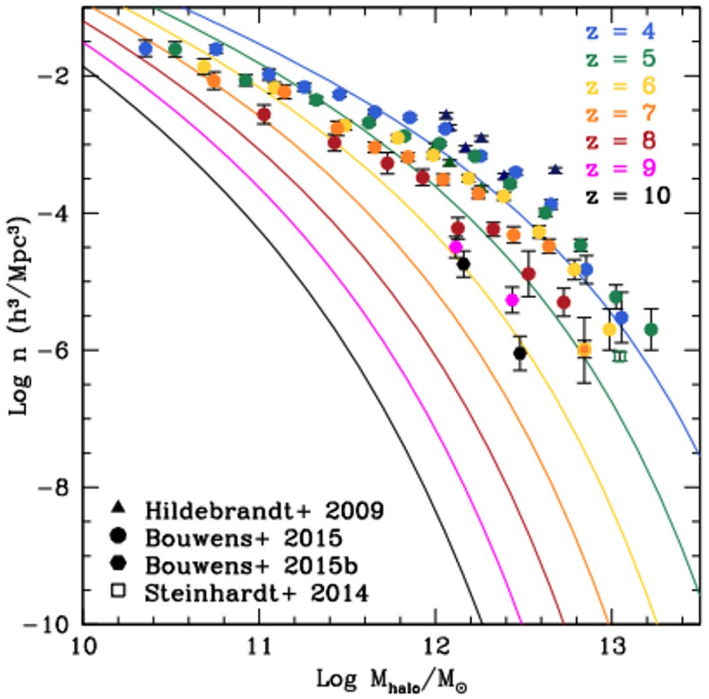
Boylan-Kolchin (2023)

An Old Problem: “Impossibly Early” Galaxies

This tension is not new: Steinhardt et al. (2016) named “the impossibly early galaxy problem” well before JWST. Their figure shows observed high-z galaxies implying far more massive halos than ΛCDM predicts.

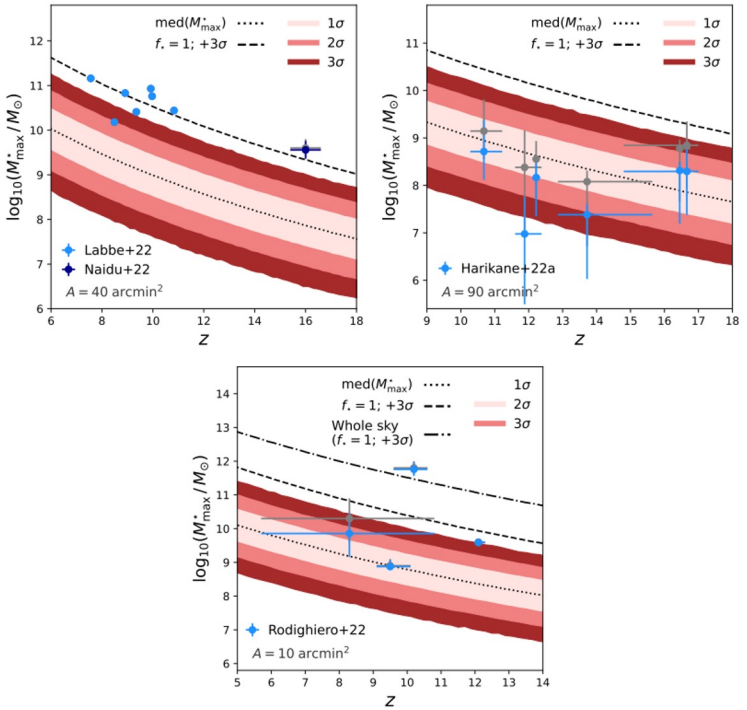
Since then, explanations have followed two main lines. Some invoke new physics (Early Dark Energy, or a higher fluctuation amplitude) to grow structure faster. Others argue we must first account properly for how galaxy masses are measured and counted: random errors, AGN contamination, low-redshift interlopers, dust, the IMF, cosmic variance.

This work follows the second line.



Steinhardt et al. (2016)

A Previous Approach



Lovell et al. (2023)

This problem has already been attacked with the very tool we will use. Lovell et al. (2023) applied Extreme Value Statistics to the ΛCDM halo mass function, converted halo masses into a maximum stellar mass ($f_b = 0.16$ and, conservatively, $f_* = 1$: 100% of baryons turned into stars), and asked how massive a galaxy can be in a given JWST survey volume.

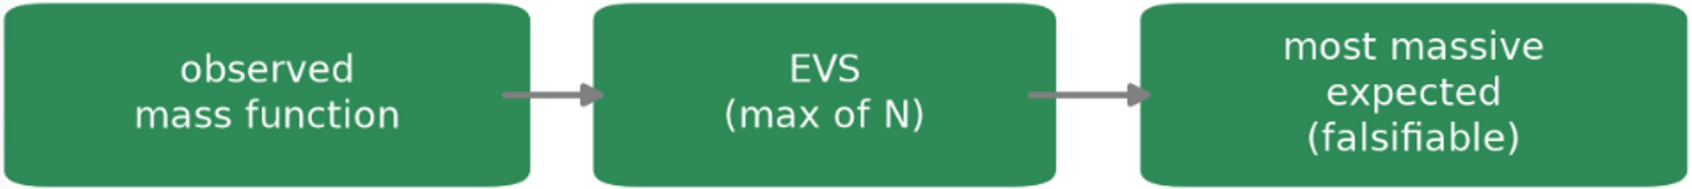
The shaded bands are the resulting confidence intervals (up to 5σ) for three early-release areas (40, 90, 10 arcmin^2); the points are JWST candidates (Labbé+22, Harikane+22, Donnan+22, Naidu+22, Rodighiero+22).

Reframing the question

✗ **Model-based (model-dependent):**

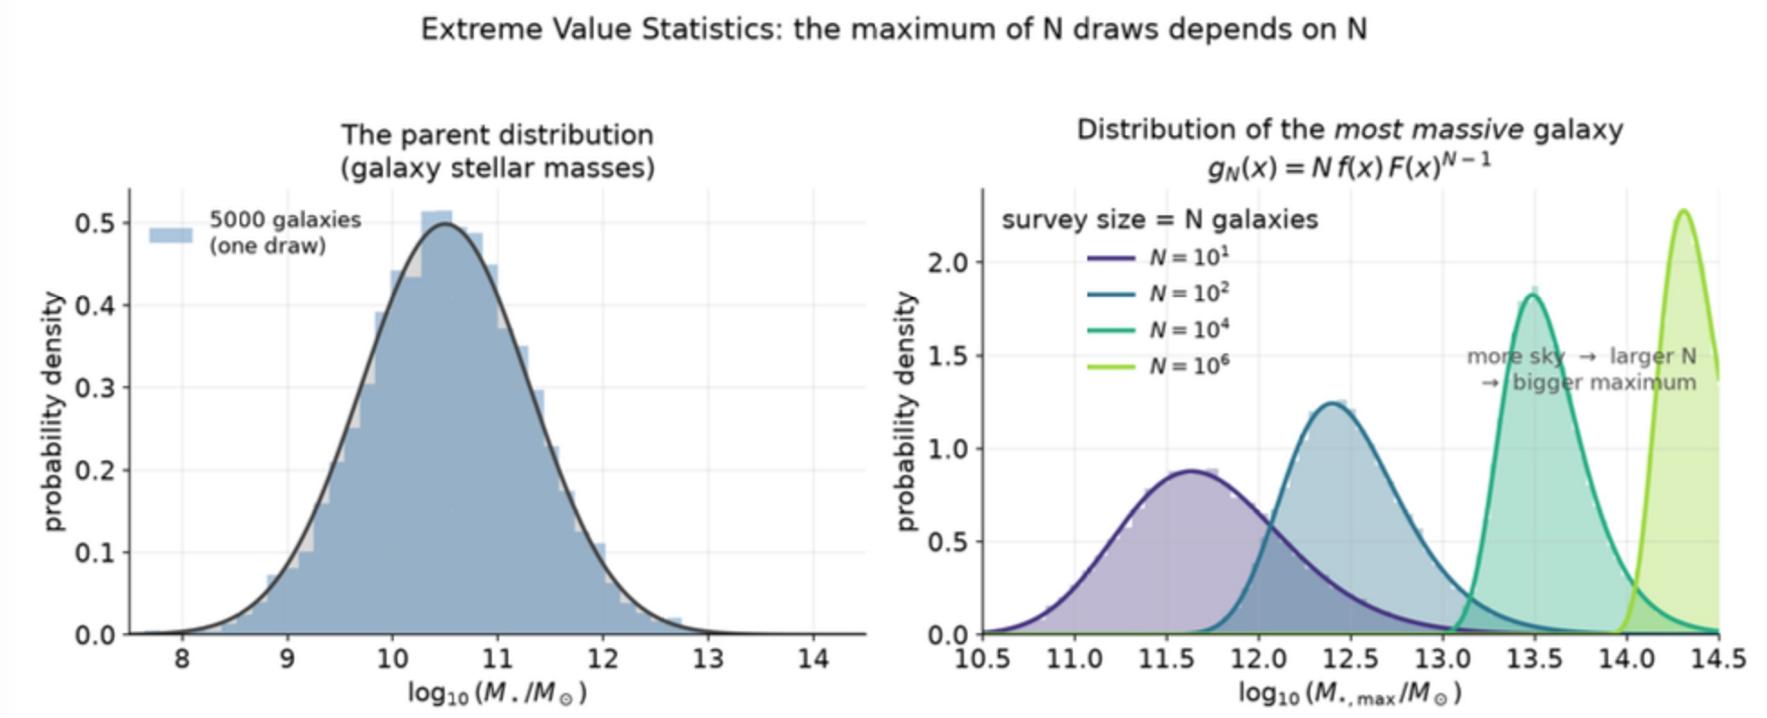


✓ **Empirical (this work):**



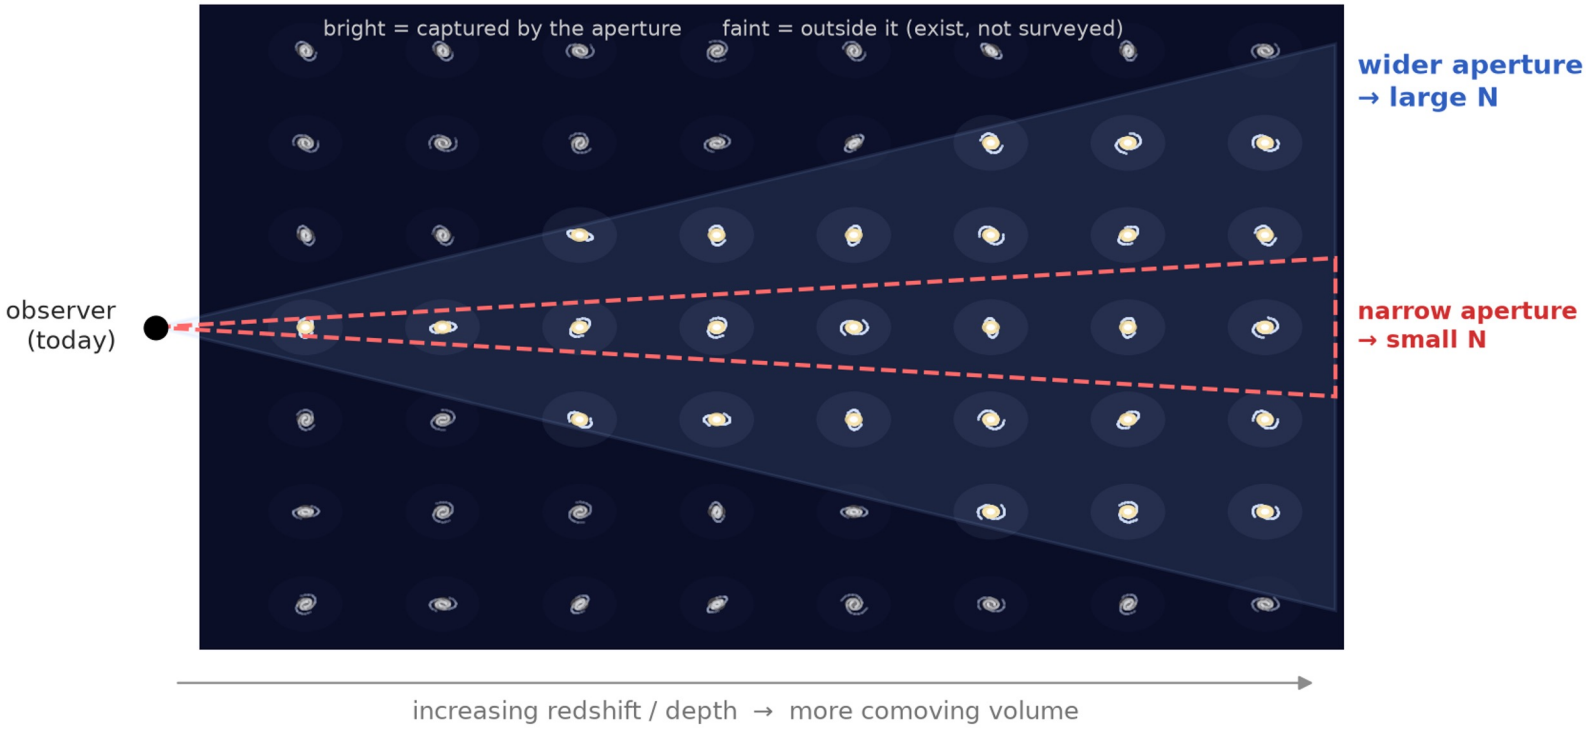
METHOD

Extreme Value Statistics



Survey area sets N — the light-cone volume

the survey aperture selects which galaxies you capture — widen it and N grows



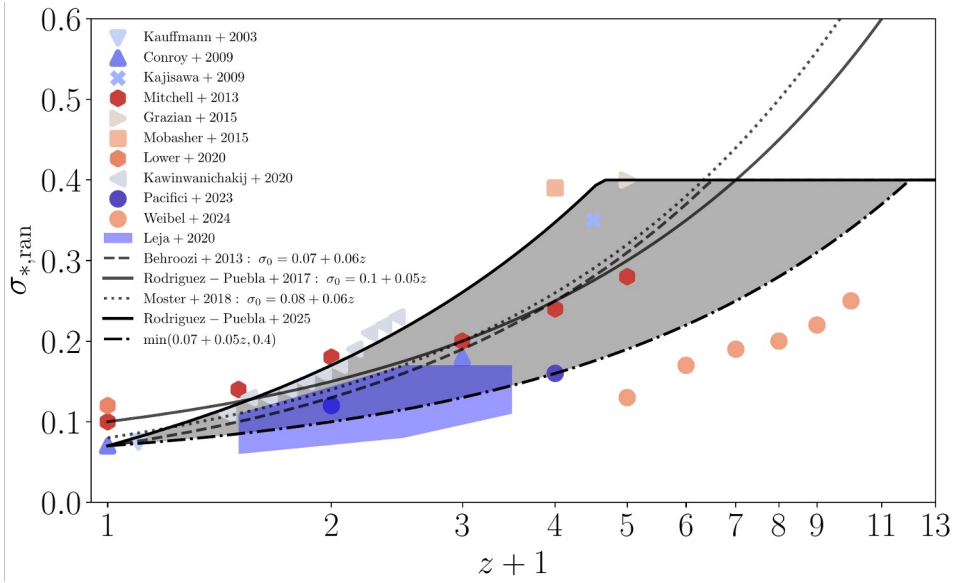
METHOD

Why Stellar Masses Are Uncertain

Stellar masses are not measured but inferred by fitting model spectra to photometry, and the result depends on the assumed star-formation history, dust, nebular emission, and the “outshining” of old stars by young ones. So every mass carries a random error, and it grows with redshift.

The points are observational estimates of that error from individual studies. Some recent studies (Leja+20, Pacifici+23, Weibel+24) find smaller errors, so we adopt the shaded range, taking its upper bound (Rodríguez-Puebla+25) as our fiducial model.

The curves show literature estimates of the 1σ random error rising from ~0.1 dex locally to ~0.4 dex by z ~ 5 and ~0.6 dex by z ~ 7–8.



Enríquez-Vargas et al. (2026)

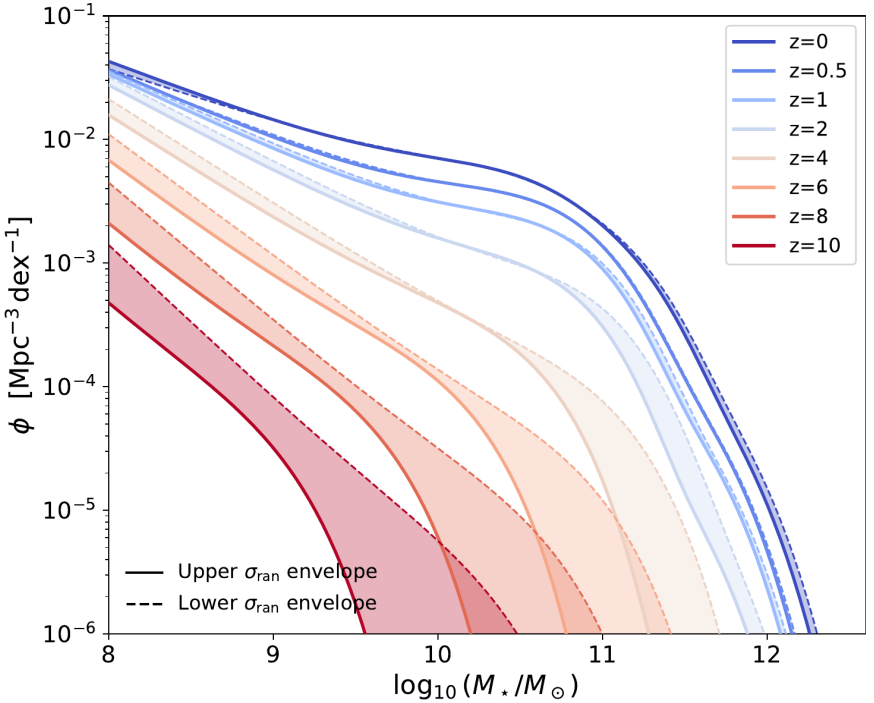
METHOD

Eddington Bias: Errors Inflate the Massive End

These errors are symmetric for each galaxy. But the mass function is not symmetric: above the knee there are far fewer galaxies than below it, so many more low-mass galaxies scatter up into the massive bins than massive galaxies scatter down (Eddington 1913, 1940).

The figure shows our mass functions from z = 0 to 10, corrected for random errors assuming the upper (solid) and lower (dashed) bounds of the errors. The correction is negligible at low redshift, but at the massive end and at high redshift it is large and depends strongly on the size of the errors.

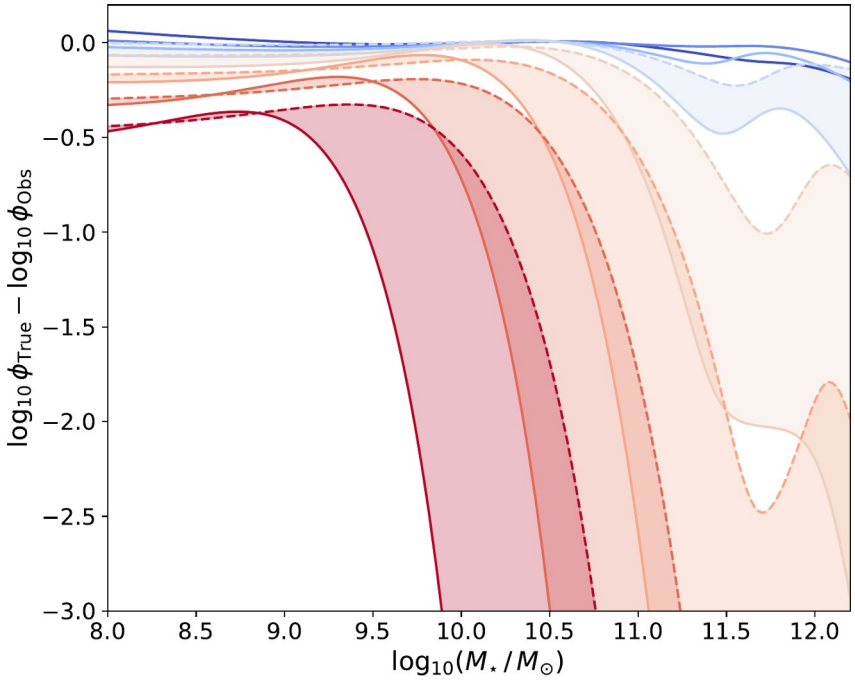
The observed massive end, $\phi_{obs} = G * \phi_{true}$, must be deconvolved before comparing with ΛCDM.



Enríquez-Vargas et al. (2026)

METHOD

The Massive End



Enríquez-Vargas et al. (2026)

The curves show the logarithmic difference between the observed mass function and the “true” one recovered by deconvolving the random errors, at several redshifts.

The difference is concentrated at the massive end (as the slope argument predicts) and grows dramatically with redshift, reaching up to two orders of magnitude in number density at the high-mass end by z ~ 4–6.

A large part of the observed abundance of the most massive galaxies is, therefore, an artifact of measurement noise.

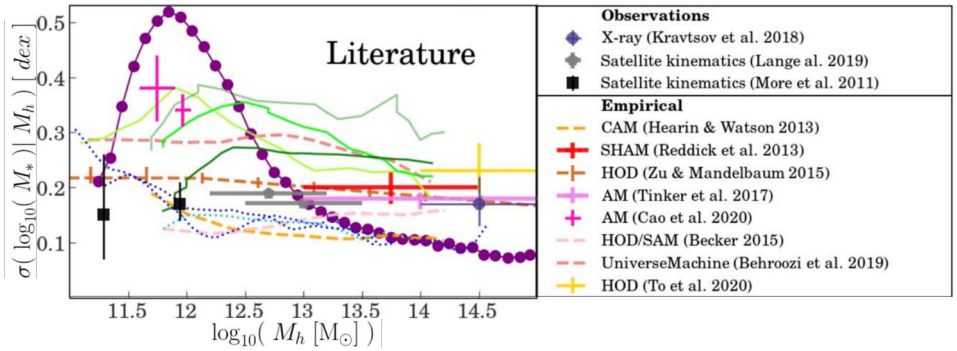
METHOD

A Second Bias: Scatter in the Galaxy–Halo Connection

Abundance matching ties galaxies to halos through the stellar-to-halo mass relation, but that relation is not exact: at fixed halo mass, central galaxies span a range of stellar masses.

The figure plots this scatter, $\sigma(\log M_* \mid M_{halo})$, against halo mass. Observations, empirical models and IllustrisTNG cluster near 0.15–0.2 dex around the 10^{12} $M_\odot$ pivot (some semi-analytic models predict more).

Because the halo mass function is steep, this scatter promotes lower-mass halos into high-stellar-mass bins: the same Eddington-type inflation, a second time. **We adopt 0.15 dex and deconvolve it to recover the intrinsic mass function.**



Porras-Valverde et al. (2024)

RESULTS I

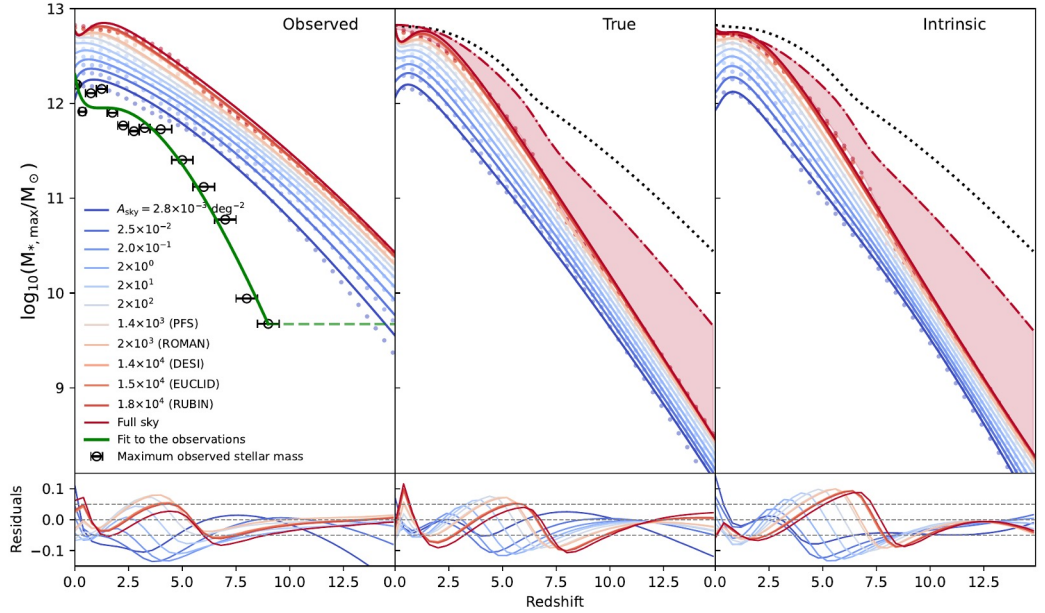
Eddington Bias and the EVS of the Most Massive Galaxies

Observed → True → Intrinsic

The first main result, in three panels: observed, true, intrinsic. For the full sky the most massive galaxy expected is ~7×10^{12} $M_\odot$ at z ~ 0, falling to ~10^{12} by z ~ 7.5 and ~10^{10} by z ~ 16.

Two trends matter: smaller survey areas give systematically smaller maxima (the gap to full sky exceeds an order of magnitude by $z \gtrsim 12$), and removing the biases steepens the fall-off, with the observed-to-true difference exceeding ~1 dex by z ~ 10.

The most massive galaxy at high redshift is much smaller than the raw numbers suggest.



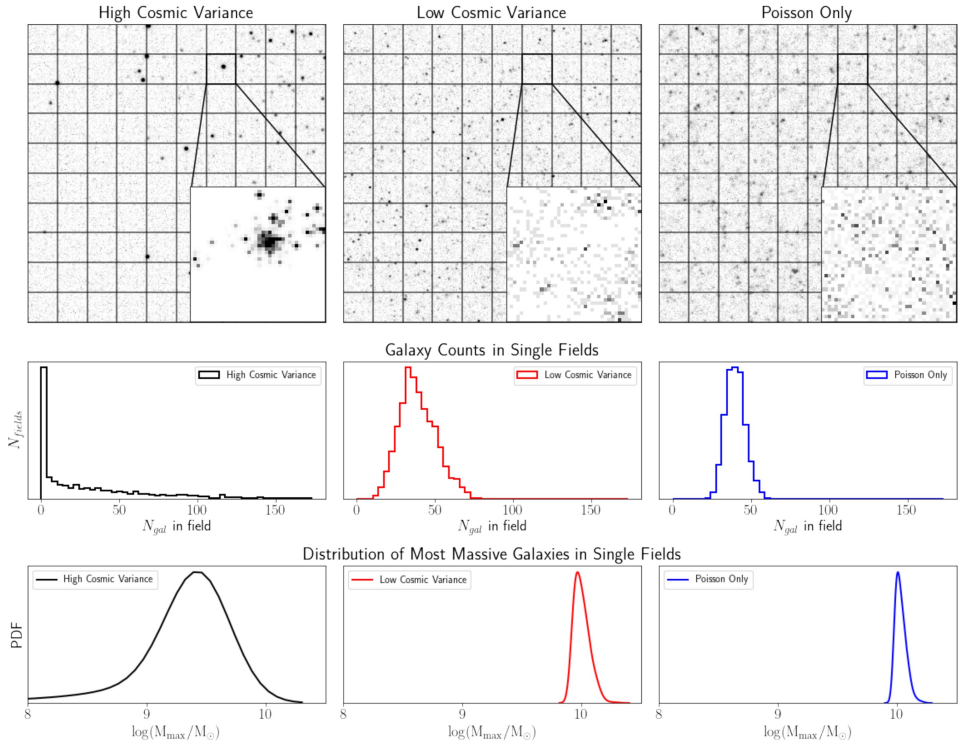
Enríquez-Vargas et al. (2026)

Clustering Breaks the Independence Assumption

The EVS formula assumes the N galaxies in a field are mutually independent. Clustering breaks that: galaxies pile up in the same few fields.

Three mock surveys above the same mass limit. With high cosmic variance (left) the counts per field go from Poisson-like to a skewed distribution with a long tail and a high chance of finding nothing at all, and the most massive galaxy becomes much wider and shifts to lower masses.

A field that hosts a massive galaxy likely hosts several, but only the most massive one is counted, while most fields are left with none.



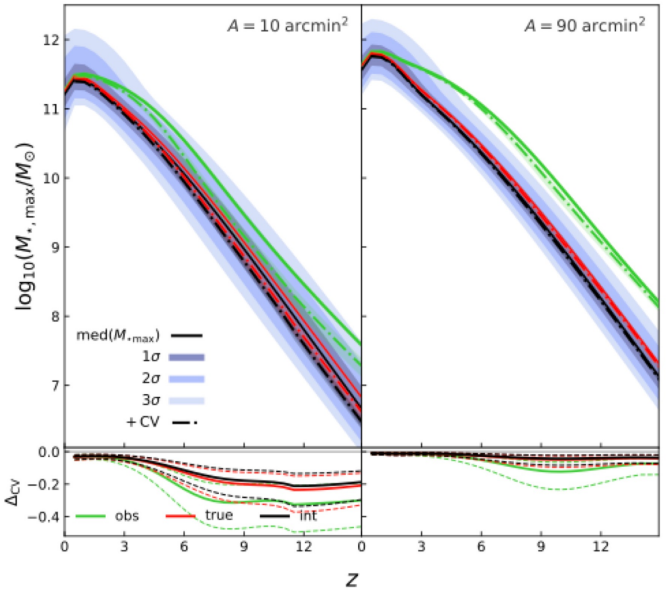
Jespersen, Steinhardt, Somerville & Lovell (2025), ApJ 982, 23 — their Figure 1

How Much It Moves Our Prediction

Following their procedure, we calibrate the cosmic variance on eight UniverseMachine GOODS-S lightcones and recompute the EVS for our two smallest apertures.

Every contour moves to lower masses. At z ≳ 8 the median drops by 0.30–0.33 dex for 10 arcmin² and 0.07–0.12 dex for 90 arcmin²; once the biases are removed (intrinsic), only 0.17–0.21 and 0.04 dex.

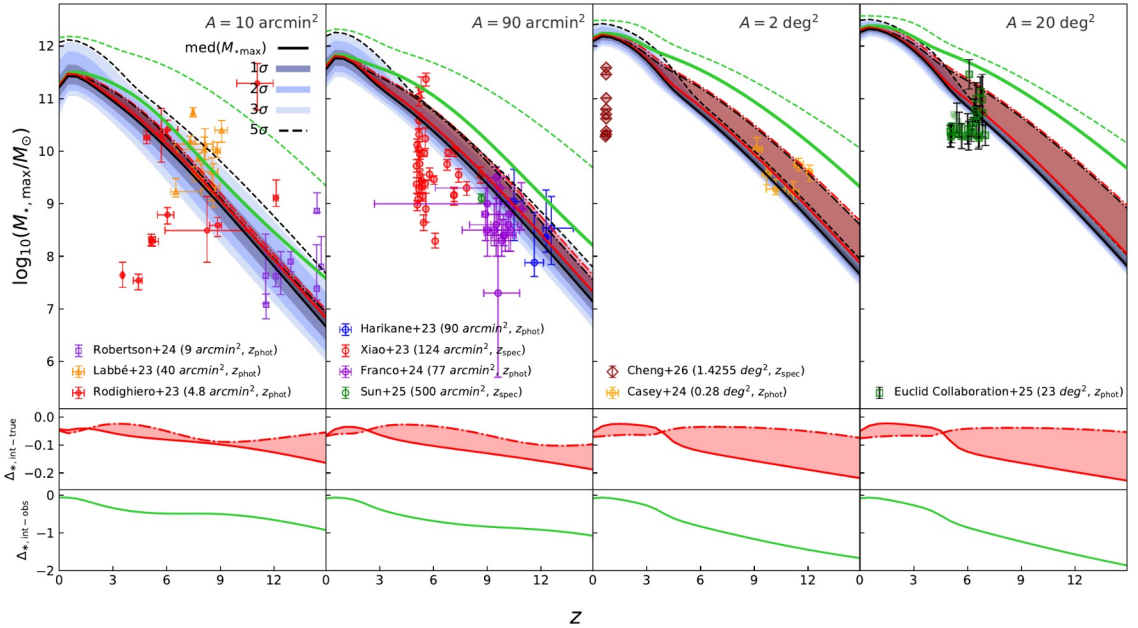
The +5σ limits do not move, so our conclusions are unaffected. The medians in the small-area panels are averages over many fields, not predictions for a single one.



Enríquez-Vargas et al. (2026), Appendix D

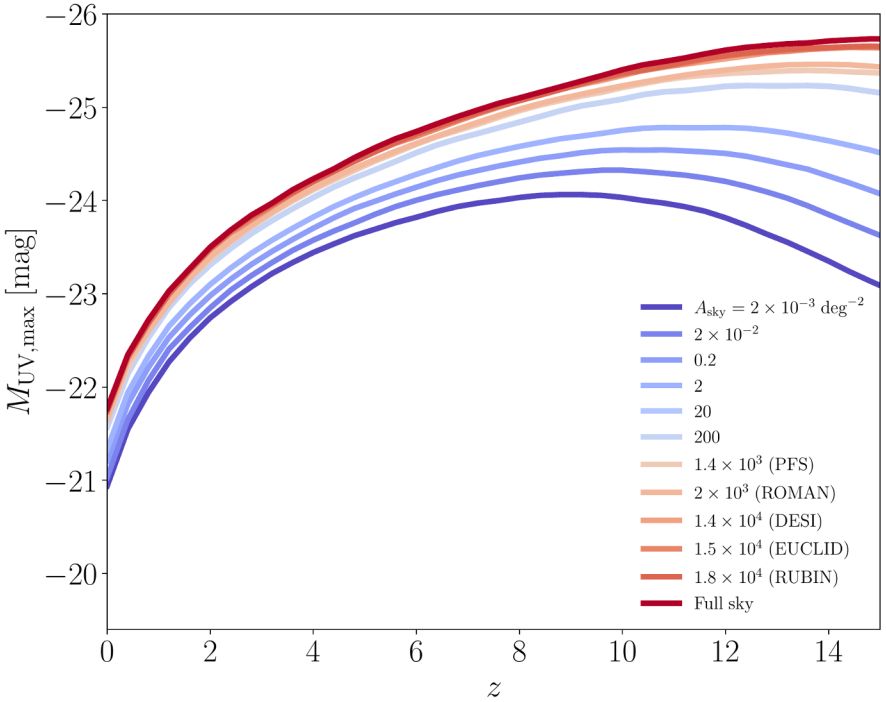
JWST Meets the Prediction

Across four survey areas, the extreme galaxies reported by JWST and Euclid teams (homogenized to a common IMF) fall essentially on the extreme-value prediction from the observed mass function: **given their survey areas, they are not anomalous.** Some very bright photometric candidates at z ≳ 10 (Sun & Yan 2025) lie far above even the 5σ limit and are not shown: they are likely low-redshift interlopers, since a dusty z ~ 3–4 galaxy can mimic z ~ 10. Whether they fit inside ΛCDM halos comes next.



Enríquez-Vargas et al. (2026)

The Most Luminous Galaxy



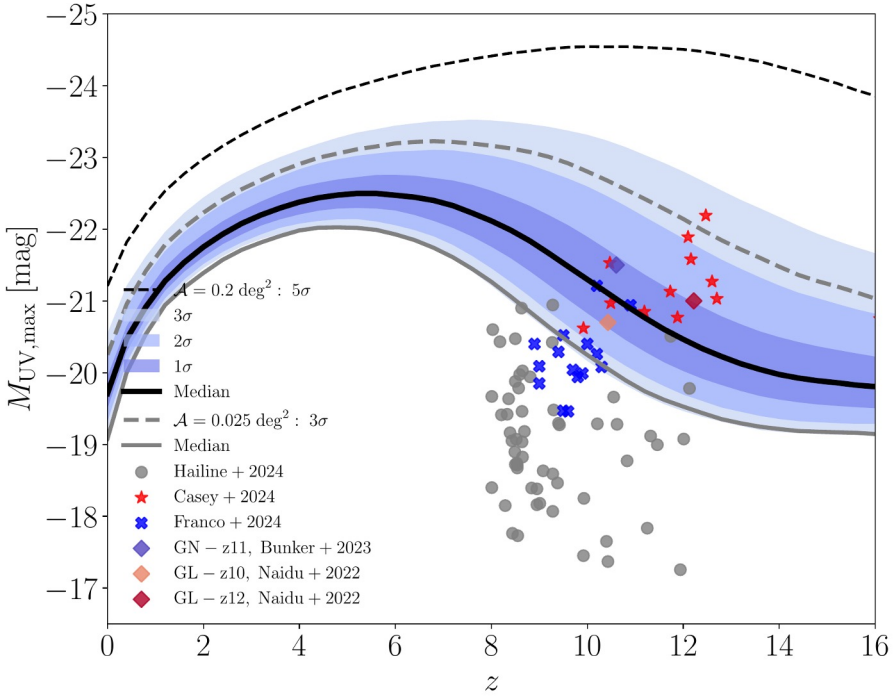
Enríquez-Vargas et al. (2026)

The same machinery runs on the observed UV luminosity function to predict the brightest galaxy.

Unlike mass, which accumulates monotonically, UV brightness is bell-shaped in redshift: it peaks and then declines toward low z, because the most massive galaxies quench and stop forming the young stars that produce ultraviolet light.

For the full sky the brightest galaxy reaches M_{UV} ~ –25.7 near z ~ 15, a star-formation rate of ~470 $M_{\odot}$ yr^{-1} which, being uncorrected for dust, is a lower limit.

The Brightest Galaxies, Captured



Enríquez-Vargas et al. (2026)

The same check for ultraviolet brightness, for a COSMOS-Web-sized area. The brightest JWST candidates (Casey's and Franco's luminous samples, GN-z11, the GLASS objects) all fall within the 1–3σ band of the extreme-value prediction.

The framework neither under- nor over-predicts the brightest galaxies we actually know about.

Masses and luminosities tell the same story: the observed extremes sit comfortably within ΛCDM once the survey volume is respected.

RESULTS II

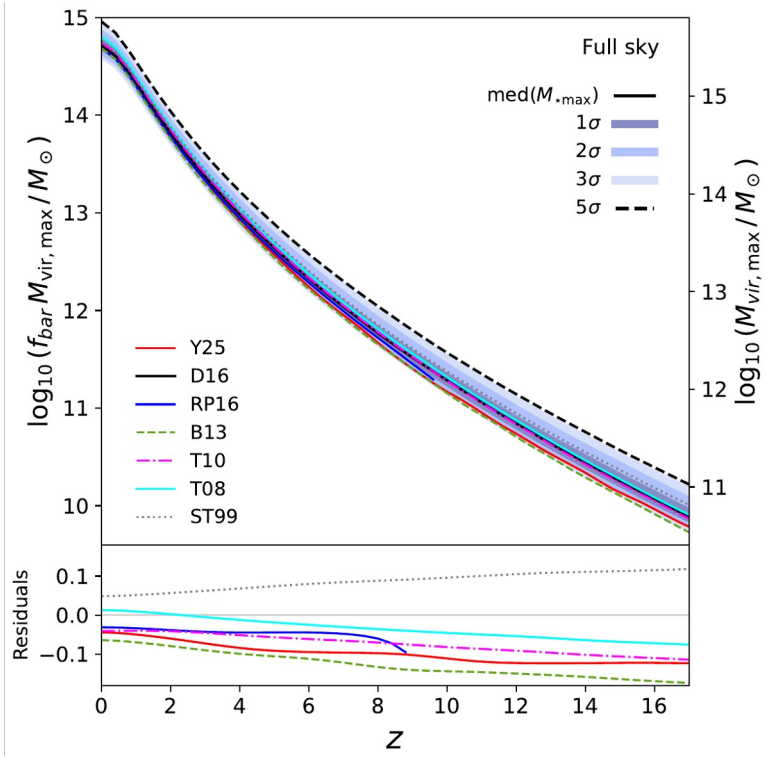
EVS and Dark Matter Halos

The Most Massive Halo & Its Baryons

Feeding the halo mass function through the same extreme-value machinery gives the most massive halo, and its baryonic budget $f_b \cdot M_{vir}$ is the absolute ceiling on stellar mass.

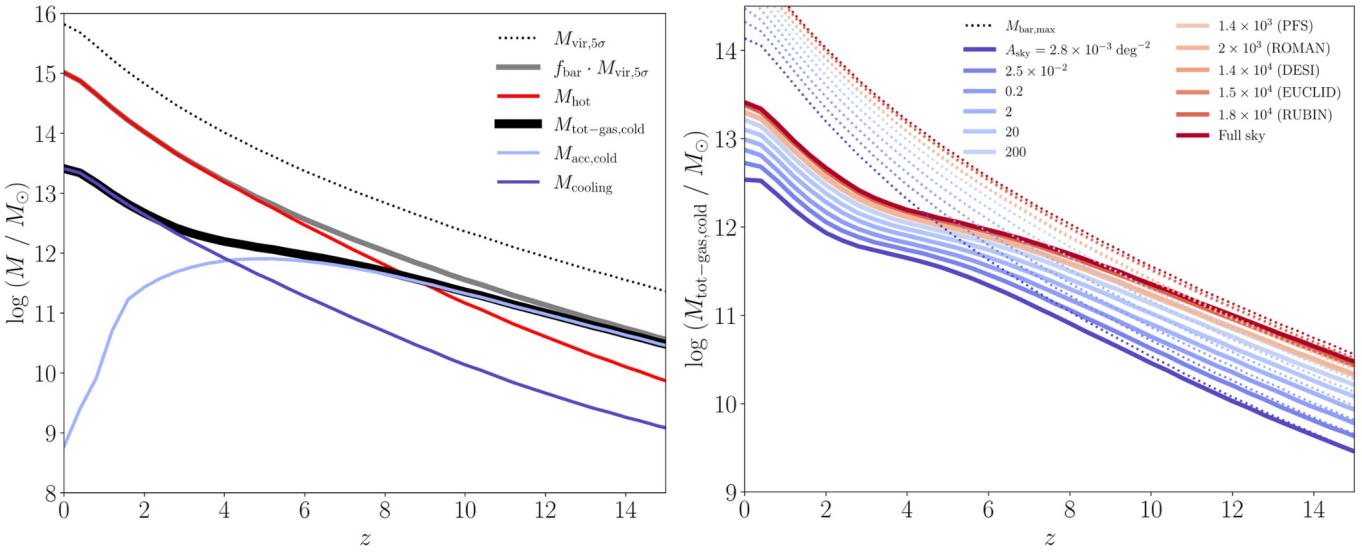
Seven different halo mass functions from the literature agree to ~0.1 dex at low redshift and ~0.2 dex at high redshift, so the cosmological side of the problem is robust.

For the full sky at z ~ 0 the most massive halo is 5.2×10^{15} $M_\odot$, with a baryonic budget of 8×10^{14} $M_\odot$.



Enríquez-Vargas et al. (2026)

The Cold-Gas Reservoir



Enríquez-Vargas et al. (2026)

The maximum stellar mass is set not by the full baryon budget but by the cold gas that can actually form stars. Above a shock-heating mass of ~$10^{11.8}$ $M_\odot$ gas should heat and stall, yet at high redshift the most massive halos are fed by narrow cold streams that penetrate the hot halo; below z ~ 1.2 these give way to cooling flows. **Modeling both channels (calibrated on simulations) gives the total cold-gas reservoir, $M_{cold} \leq f_b \cdot M_{vir}$.**

RESULTS III

The Galaxy–Halo Connection in the EVS Formalism

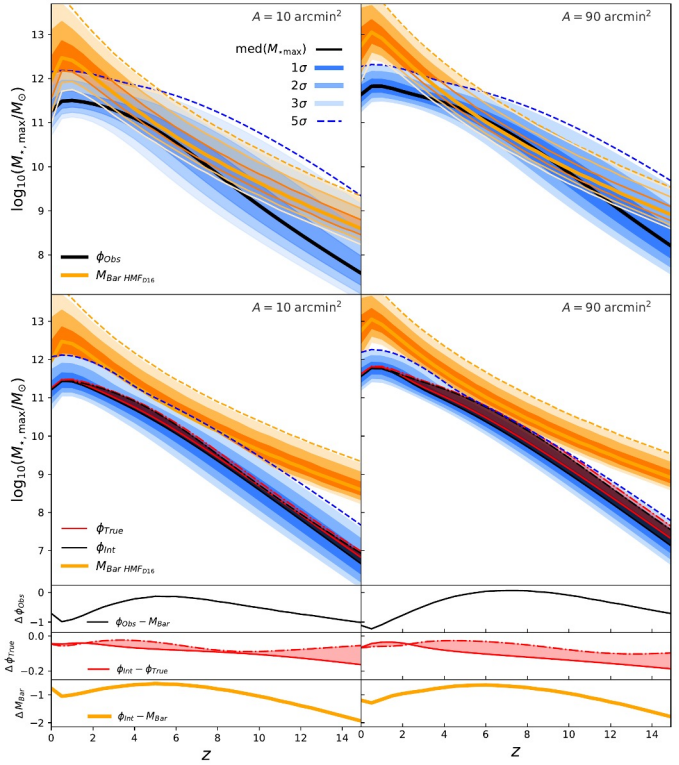
How Efficient Can Galaxies Be?

Comparing the most massive galaxy directly with the most massive halo's baryon budget, for the small areas JWST surveys.

With the observed stellar masses the two overlap across z ~ 3–6, which would imply halos converting essentially all their baryons into stars: the ~100% efficiency that triggered the alarm.

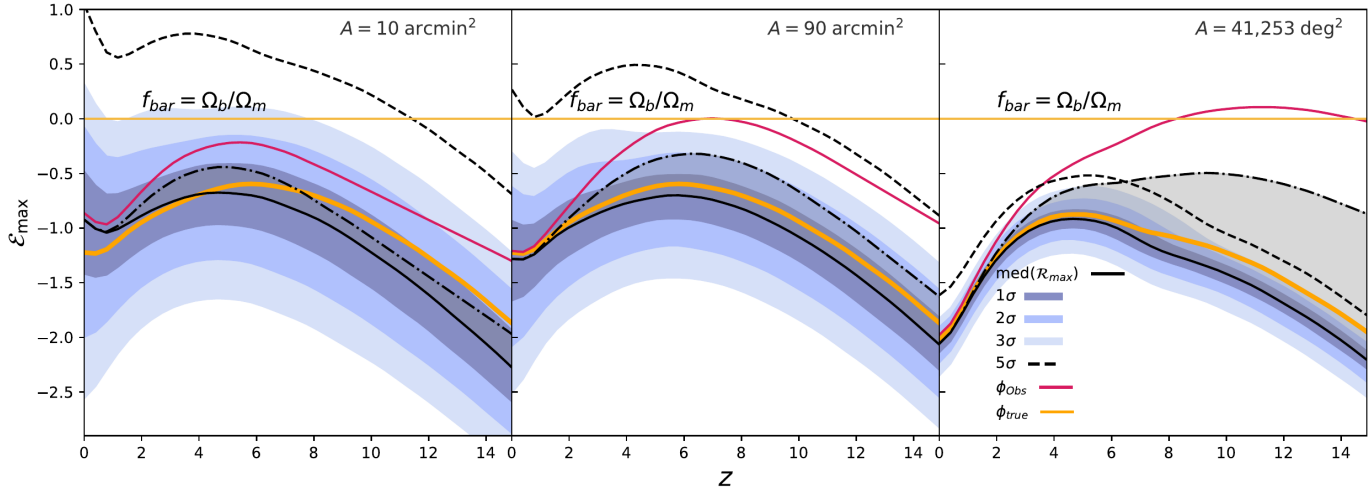
With the intrinsic masses (removing both Eddington bias and the galaxy–halo scatter) the stellar masses drop below the ceiling, and far below it at z ≳ 8.

The apparent near-100% efficiency was largely an artifact of the bias.



Enríquez-Vargas et al. (2026)

Pairing Galaxies and Halos: Three Options



Enríquez-Vargas et al. (2026)

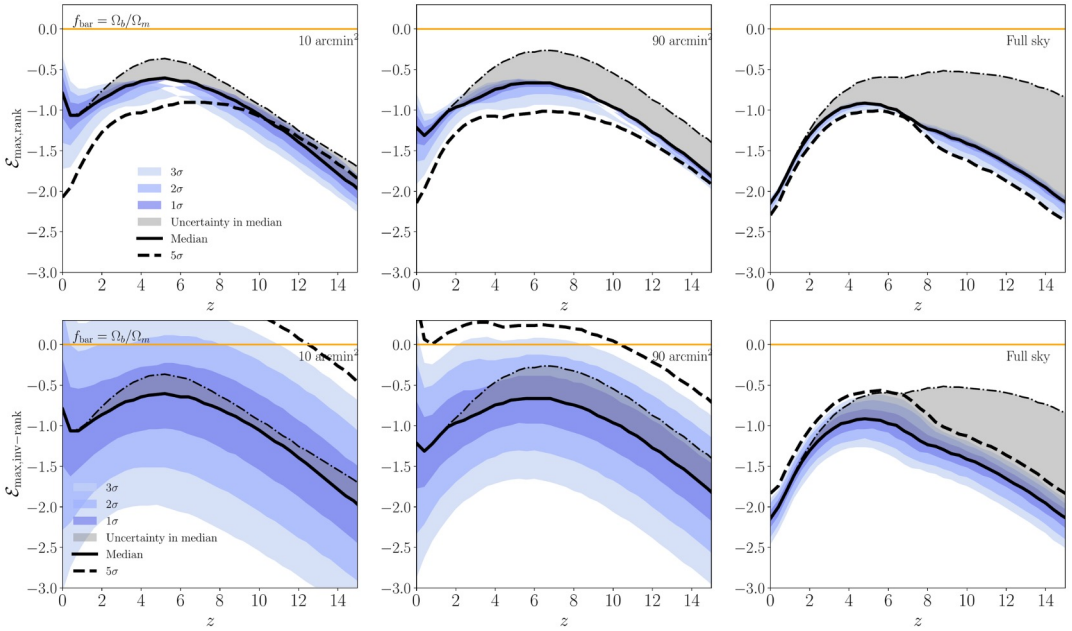
EVS gives the most massive galaxy and the most massive halo separately; how to pair them is an extra assumption. **Stochastic:** independent variables, so ε_{max} follows from a convolution (this figure). **Rank-order:** the most massive halo hosts the most massive galaxy. **Inverted rank:** it hosts the least massive one.

In the stochastic case a sizeable tail reaches $\varepsilon > 1$ in small fields, even for the intrinsic GSMF: **the galaxy–halo connection in these halos is unlikely to be stochastic, so we adopt rank-order.**

Efficiency, Corrected

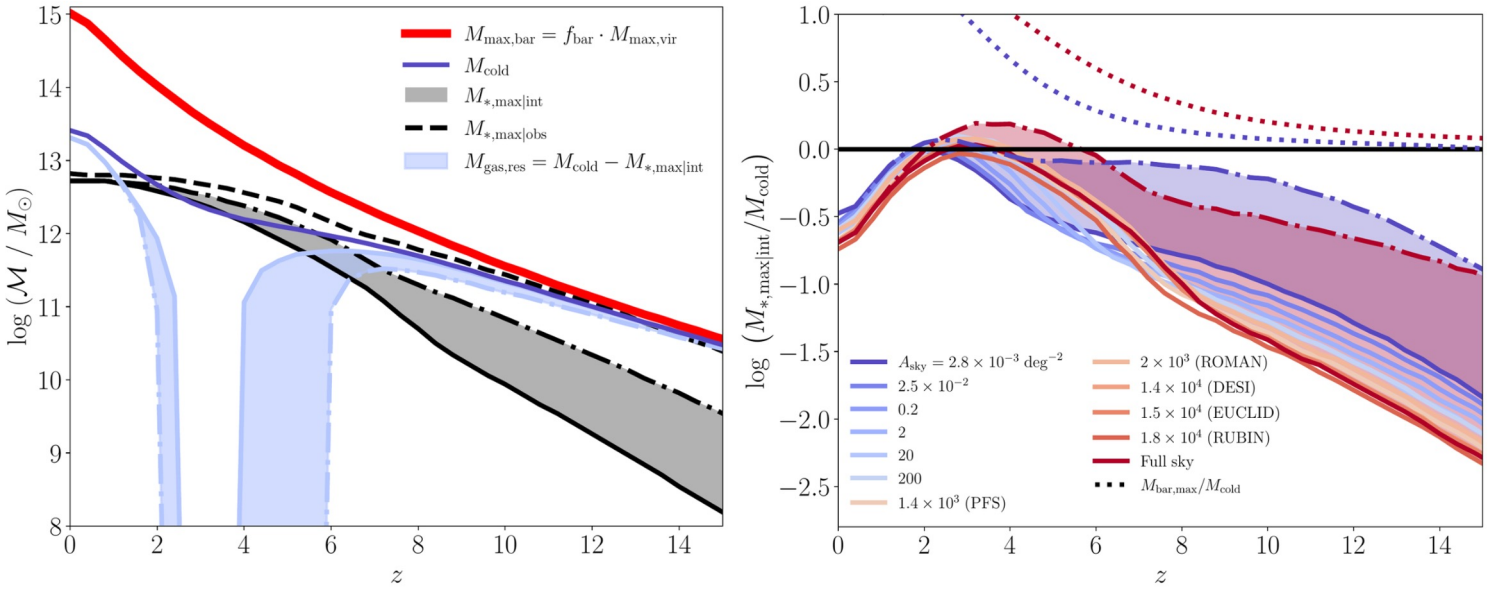
With the realistic rank-ordered pairing and the intrinsic masses, the cumulative star-formation efficiency stays below one everywhere: the unphysical region is excluded even in the smallest fields.

The most extreme halos reach efficiencies peaking at ~0.2–0.4 near z ~ 5, and for the full sky at low redshift they fall to about 1%, because larger volumes sample more massive, less efficient halos.



Enríquez-Vargas et al. (2026)

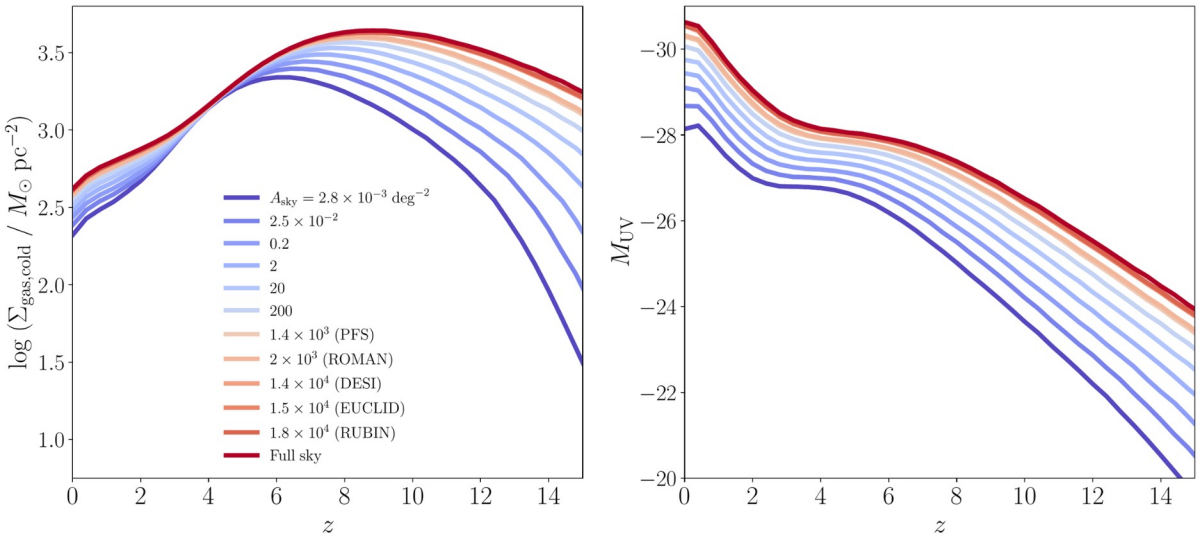
Near-Maximal Efficiency at z ~ 2–6



Enríquez-Vargas et al. (2026)

Across z ~ 2–6 the most massive galaxies have stellar masses comparable to their entire cold-gas reservoir (near-maximal efficiency), and this window coincides with the epoch of dusty, submillimeter starbursts that contribute about half of all cosmic star formation. **The cosmological "crisis" at z > 8 dissolves once the biases are removed, but a real episode of near-maximal star formation emerges at cosmic noon.**

Physical Conditions in the Most Massive Halo



Enríquez-Vargas et al. (2026)

Left: the cold-gas surface density reaches starburst-like 10^2–10^3 $M_\odot$ pc^{-2}, peaking at z ~ 6–9 and nearing the cloud-disruption threshold Σ_{crit} ~ 2200 $M_\odot$ pc^{-2}; small areas drop to normal-galaxy values by $z \gtrsim 12$.

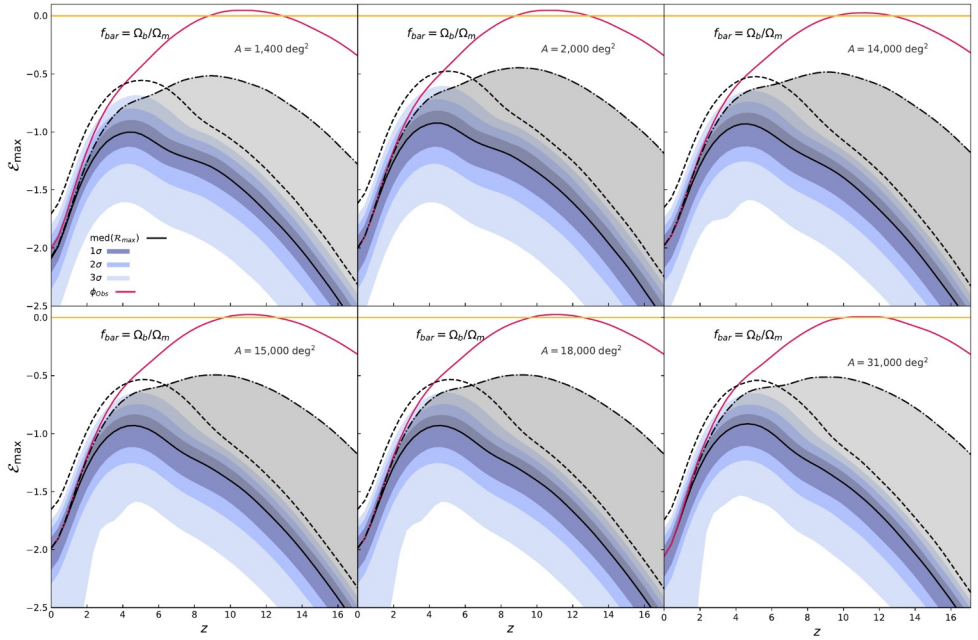
Right: turning that gas into stars makes the halo extremely UV-bright (M_{UV} ~ −27 to −23 at $z \gtrsim 8$, full sky), with a z ~ 2–7 plateau marking the cold→hot accretion transition. Without dust and without subtracting formed stars, the low-z values (down to M_{UV} ~ −30.5) are upper limits.

Survey Area Is the Lever

Because survey area sets N, wide surveys reach furthest into the rare tail and provide the sharpest test of the ΛCDM ceiling.

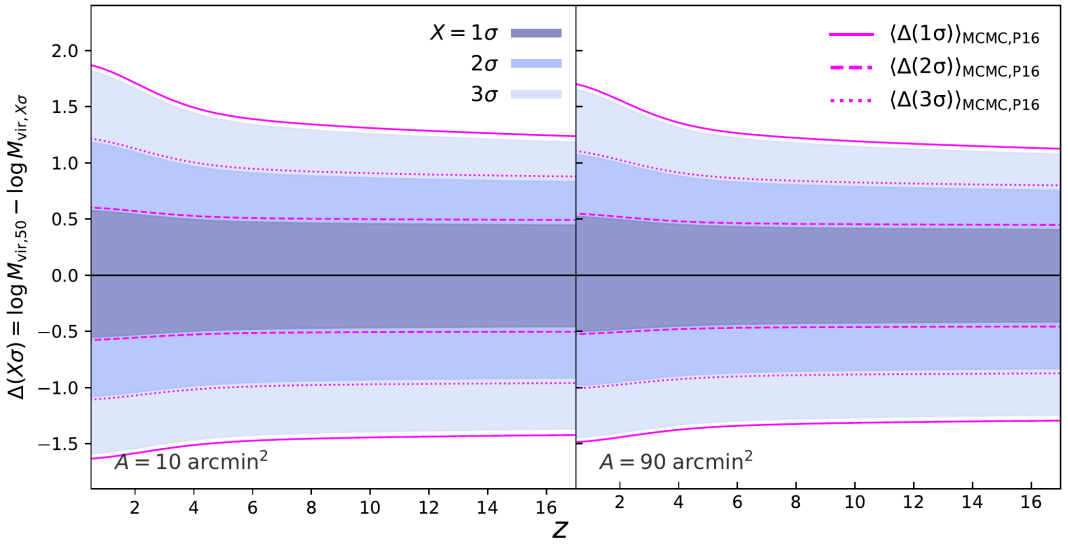
The panels show the predicted extreme-value efficiency for PFS, Roman, DESI, Euclid and Rubin; Rubin (18,000 deg^2) and Euclid (15,000 deg^2) probe the largest volumes. Each survey gets an area-matched prediction for the most massive and most luminous galaxy it should find.

A clean, spectroscopically confirmed galaxy above that ceiling would be the signal to revisit the physics.



Enríquez-Vargas et al. (2026)

Cosmological Parameters Barely Matter



Enríquez-Vargas et al. (2026)

Spread of the most-massive-halo EVS around its median (1σ, 2σ, 3σ) for the fiducial Planck cosmology (shaded), compared with the same quantity averaged over the full Planck MCMC chains (magenta), for 10 and 90 arcmin^2.

The two are nearly identical: uncertainties in the ΛCDM parameters cannot explain a tension with the most massive galaxies or halos.

SUMMARY

What We Found

1. **Empirical EVS limits** on the most massive & most luminous galaxies vs. redshift and survey area, read straight off the observed GSMF and UV LF, with no galaxy-formation model assumed.

2. **Eddington bias + SHMR scatter** inflate the observed massive end; de-biasing lowers the maximum M_* by up to ~1 dex at $z \gtrsim 10$, so corrected galaxies sit below the baryon ceiling ($M_* < f_b \cdot M_{vir}$): the JWST "impossible-galaxy" tension largely dissolves within ΛCDM.

3. **Galaxy and halo extremes are correlated** (rank-order), not stochastic: the cumulative star-formation efficiency stays physical ($\varepsilon < 1$, peaking at ~0.2–0.4).

4. **The real astrophysics:** at z ~ 2–6 the most massive galaxies are near-maximally efficient ($M_* \approx M_{cold}$), the dusty submillimetre starburst epoch.

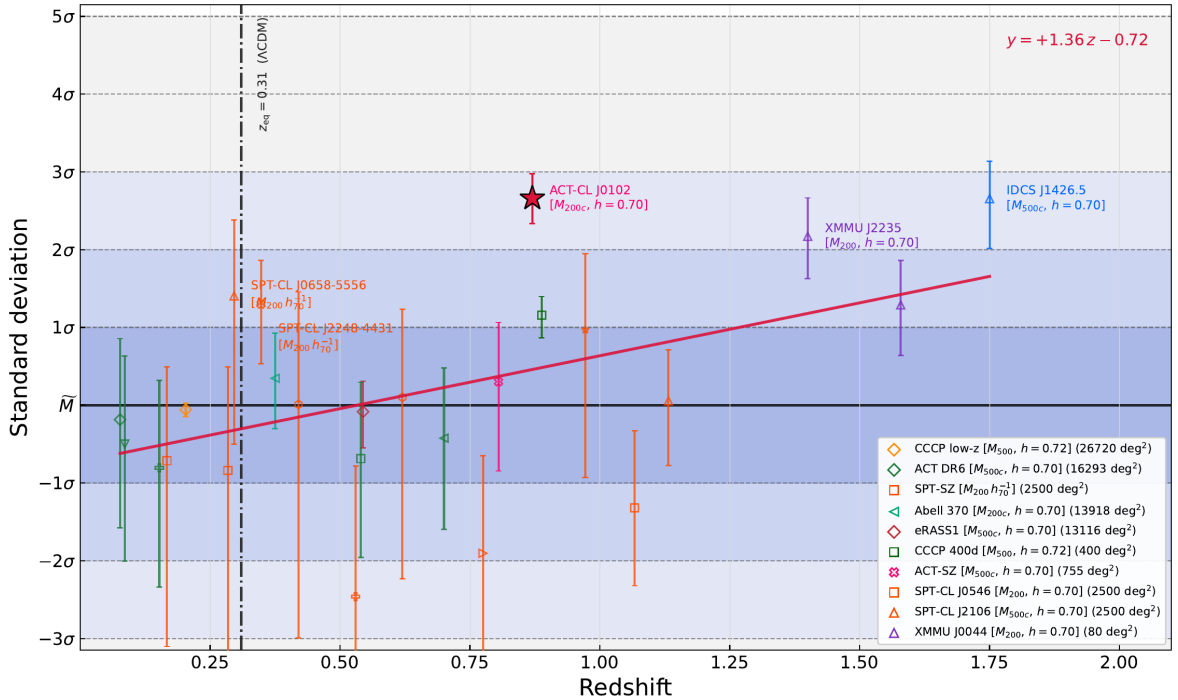
5. **The halo side is robust:** seven halo mass functions agree to ~0.1–0.2 dex and Planck parameter uncertainties barely matter, so any remaining tension lies in the baryons.

CURRENT WORK

Are We in a New Crisis in ΛCDM?



The Most Extreme Clusters under Planck ΛCDM



Enríquez-Vargas et al. (in prep.)

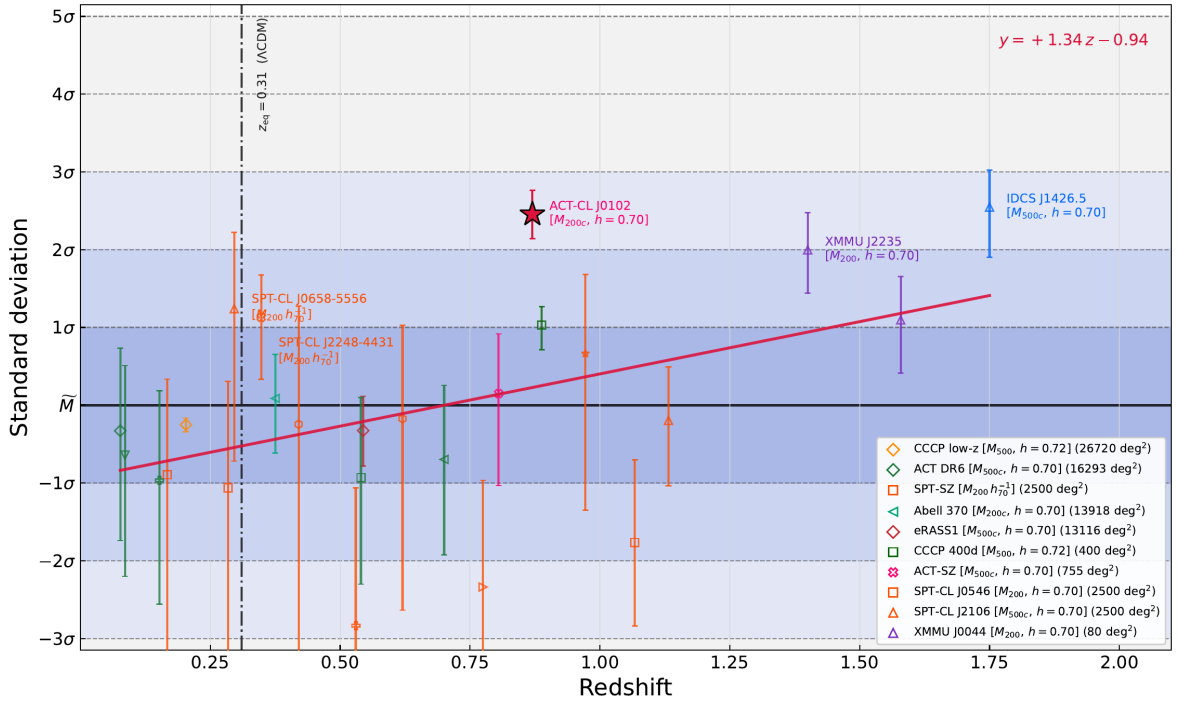
We extend EVS to the most massive known clusters, from SZ (ACT, SPT) and X-ray (eRASS1, CCCP) surveys plus individual high-z systems.

Each cluster is placed by how many σ its mass lies from the EVS median for its own survey area and redshift.

No cluster exceeds 3σ: El Gordo (ACT-CL J0102) sits at ≈2.6σ.

Yet the fit (red), σ ≈ 1.36 z − 0.72, shows clusters becoming more extreme with redshift.

Adding Primordial Non-Gaussianity (f_{NL} = 100)



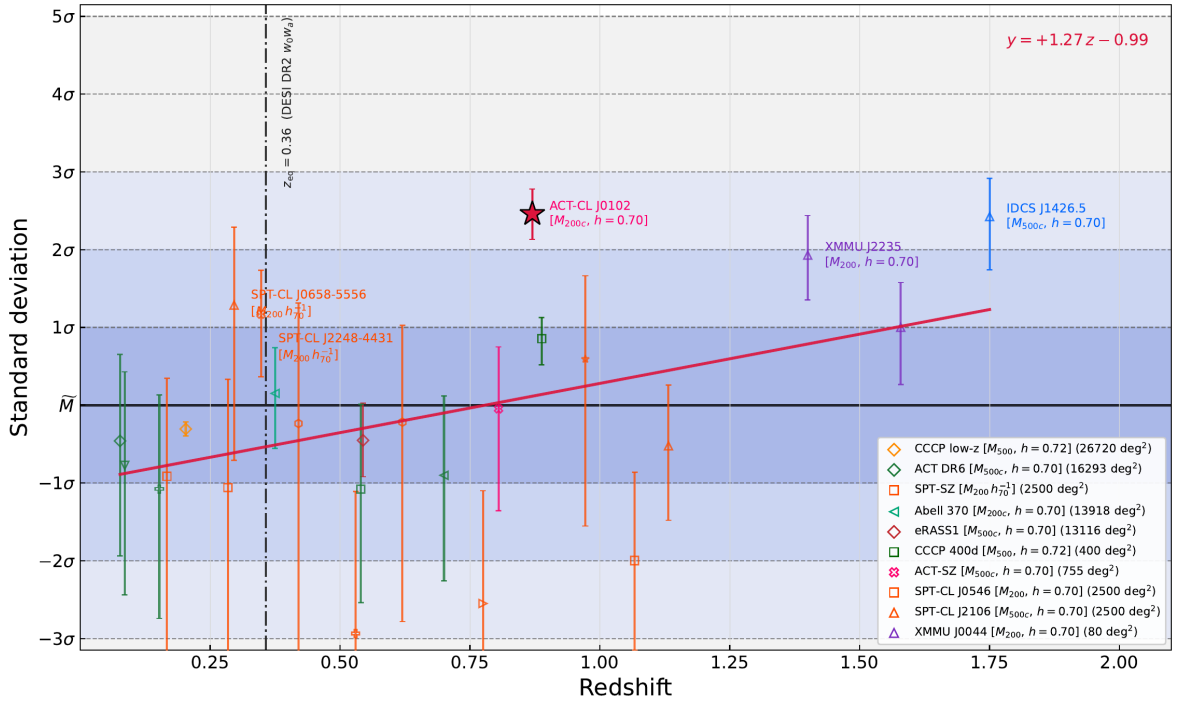
Enríquez-Vargas et al. (in prep.)

Same clusters, now with primordial non-Gaussianity, f_{NL} = 100, which boosts the abundance of the rarest halos.

All clusters move down by ≈0.2σ; El Gordo drops to ≈2.4σ.

The redshift trend survives almost unchanged: σ ≈ 1.34 z – 0.94.

The Same Clusters under DESI DR2 Dark Energy



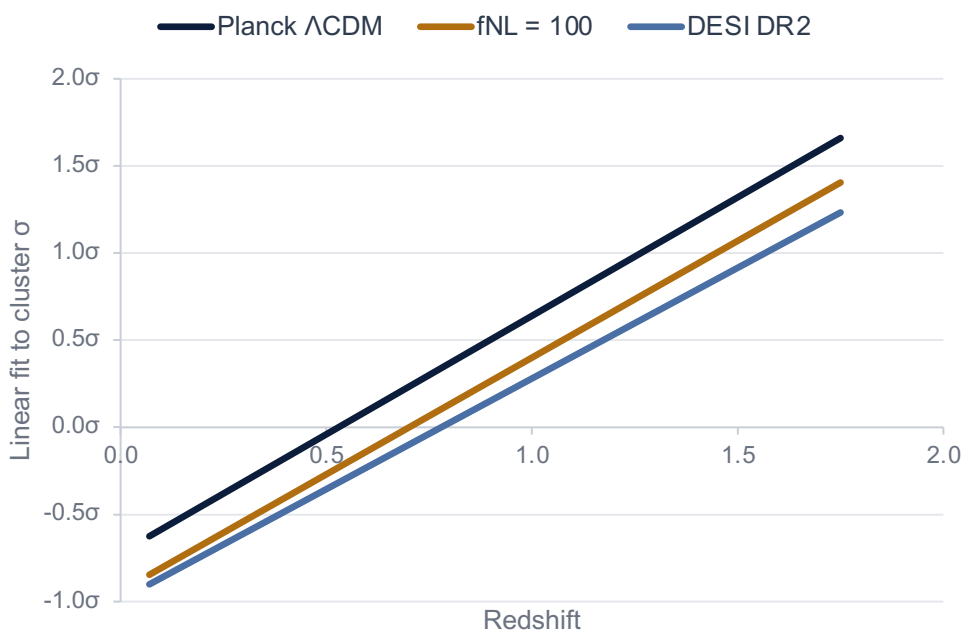
Enríquez-Vargas et al. (in prep.)

With the w_0w_aCDM cosmology favoured by DESI DR2, matter–dark-energy equality shifts from z_{eq} = 0.31 to 0.36.

Clusters move down by ≈0.2–0.3σ; El Gordo drops to ≈2.4σ.

The slope flattens only slightly: σ ≈ 1.27 z – 0.99.

What Changes Between Cosmologies



Fits from Enríquez-Vargas et al. (in prep.)

Model	Slope	Intercept	El Gordo
Planck ΛCDM	1.36	−0.72	≈2.6σ
f_{NL} = 100	1.34	−0.94	≈2.4σ
DESI DR2	1.27	−0.99	≈2.4σ

Both alternatives mainly **shift all clusters down** by ≈0.2–0.3σ; the slope, ≈1.3σ per unit redshift, barely changes.

No cluster exceeds 3σ in any of the three cosmologies.

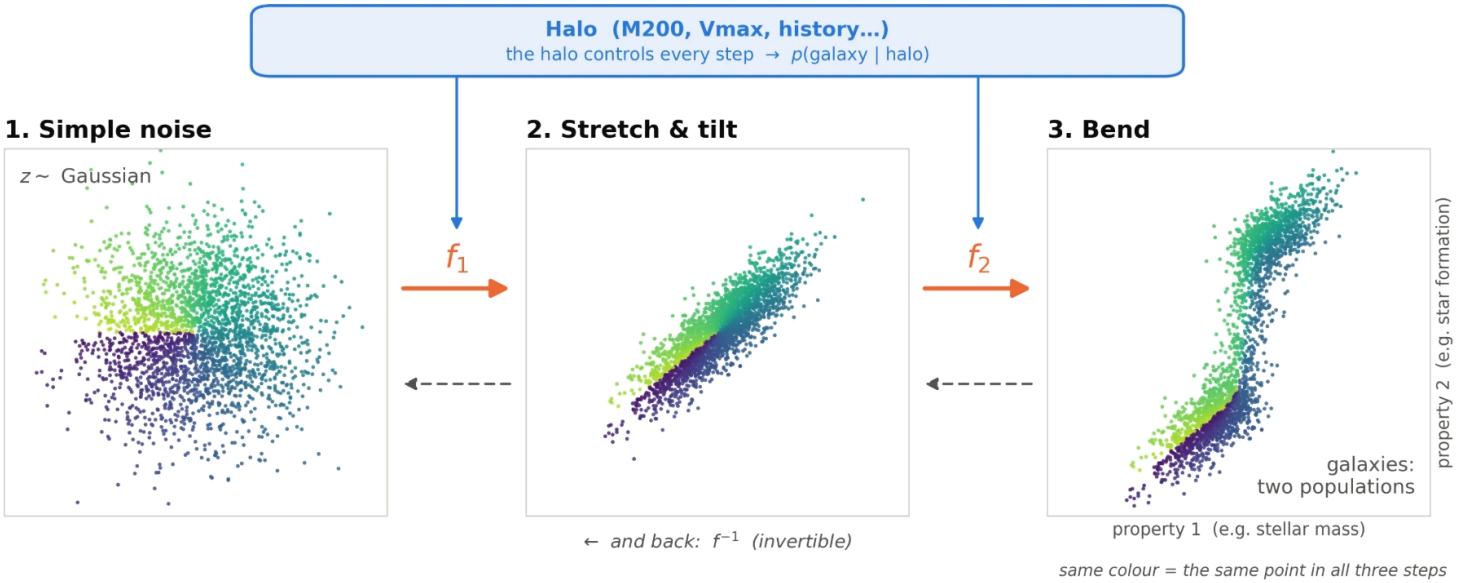
SIDE PROJECT · CIEMAT & UAM

Populating Dark-Matter Simulations with Normalizing Flows

A separate project, briefly: what I am working on in Madrid

What Is a Normalizing Flow?

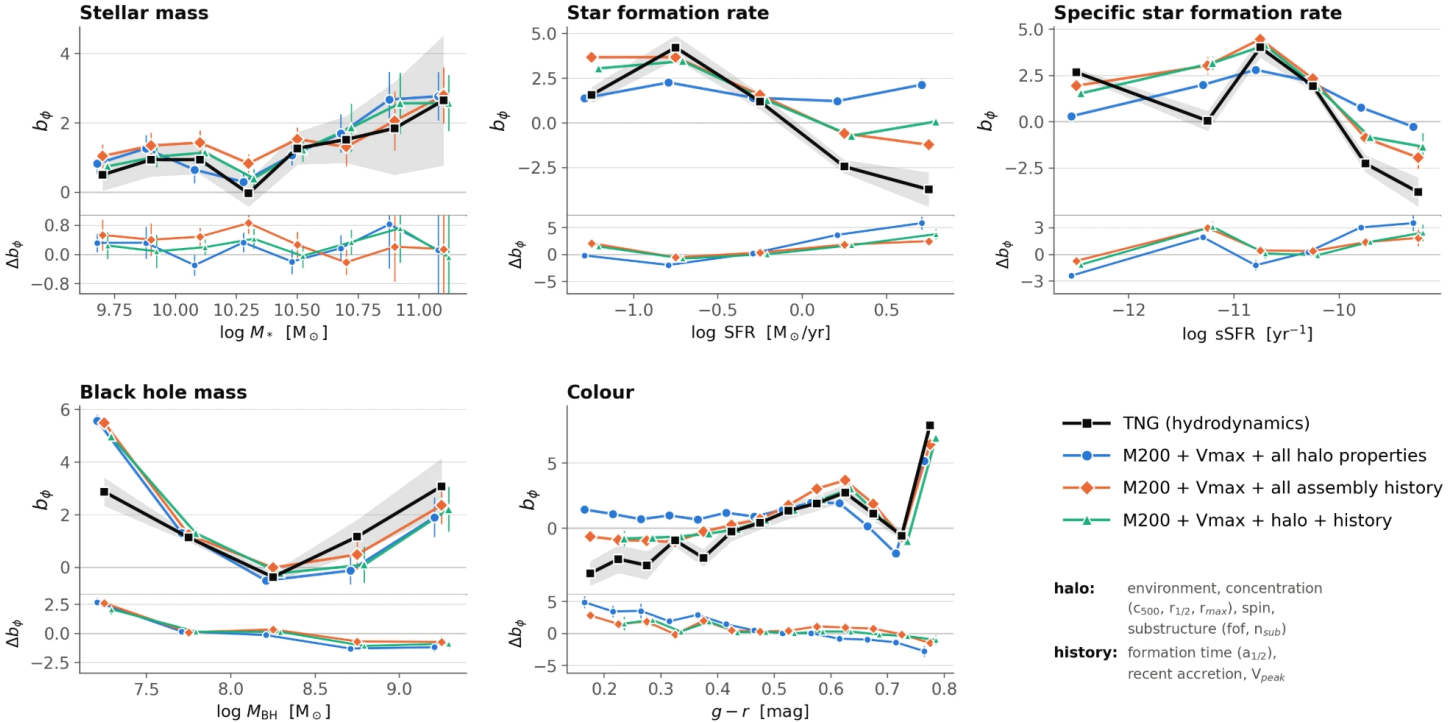
We learn p(galaxy | halo) from a hydrodynamical simulation, then use it to **paint galaxies onto halos of dark-matter-only simulations**.



Generate	draw noise z → transform → a new galaxy for that halo
Train	map each real galaxy back to the Gaussian → exact probability → maximize it

What Halo Information Does the Flow Need?

Flows trained on 27 CAMELS-TNG boxes populate 5 boxes with σ_8 = 0.6–1.0 (separate universe; centrals only). b_ϕ: response of galaxy abundance to σ_8. Lower panels: flow − TNG.



THANK YOU

Questions?

Outlook: no known cluster exceeds 3σ under Planck ΛCDM, f_{NL} = 100 or DESI DR2 dark energy, but in all three, clusters become more extreme with redshift.

Miguel Enríquez-Vargas · mvargas@astro.unam.mx
Instituto de Astronomía, UNAM