

Stress testing Λ CDM with high-redshift galaxy candidates

Received: 4 January 2023

Accepted: 7 March 2023

Published online: 13 April 2023

 Check for updates
Michael Boylan-Kolchin  

Early data from the James Webb Space Telescope (JWST) have revealed a bevy of high-redshift galaxy candidates with unexpectedly high stellar masses. An immediate concern is the consistency of these candidates with galaxy formation in the standard Λ CDM cosmological model, wherein the stellar mass (M_*) of a galaxy is limited by the available baryonic reservoir of its host dark matter halo. The mass function of dark matter haloes therefore imposes an absolute upper limit on the number density $n(>M_*, z)$ and stellar mass density $\rho_*(>M_*, z)$ of galaxies more massive than M_* at any epoch z . Here I show that the most massive galaxy candidates in JWST observations at $z \approx 7$ –10 lie at the very edge of these limits, indicating an important unresolved issue with the properties of galaxies derived from the observations, how galaxies form at early times in Λ CDM or within this standard cosmology itself.

A cold dark matter model (Λ CDM)-like cosmological models share a similar basic assumption: baryons and dark matter are well mixed at very early times, and as baryons collapse into dark matter haloes, the maximum amount of baryonic material within a halo will be equal to $M_b = f_b M_{\text{halo}}$, where $f_b \equiv \Omega_b/\Omega_m$ is the cosmic baryon fraction. This, in turn, bounds the total stellar content of a dark matter halo: $M_*(M_{\text{halo}}) \leq M_b(M_{\text{halo}})$. I show how this simple relation can be used as a stringent test of either cosmological models with minimal assumptions about galaxy formation, or the reliability of photometric selection and physical characterization of high-redshift galaxy candidates. My analysis is in many ways similar to that of Behroozi and Silk¹, who connected cumulative number densities of dark matter haloes to high-redshift galaxy stellar mass functions (see also Steinhardt et al.²), although I also consider the maximal cumulative stellar mass density allowed in Λ CDM. The question of the consistency of stellar mass functions and the underlying cosmological dark matter halo mass functions has become considerably more urgent with the release of the first data from the James Webb Space Telescope (JWST), and with it, a swarm of high-redshift galaxy candidates^{3–11}.

Assumptions

I adopt the base Λ CDM model of the Planck Collaboration¹², which assumes no spatial curvature and initial conditions that are Gaussian and adiabatic, as the standard cosmological model. I use best-fit values for cosmological parameters based on the plik TT,TE,EE + lowE + lensing

likelihood applied to the full mission data. The relevant parameters and values for this work are the present-day Hubble constant, $H_0 = 67.32 \text{ km s}^{-1} \text{ Mpc}^{-1}$; the $z = 0$ density parameter for matter, $\Omega_m = 0.3158$ (which includes baryons, dark matter and non-relativistic neutrinos); the slope of the primordial power spectrum of density fluctuations, $n_s = 0.96605$; the root mean square (r.m.s.) amplitude of the linear matter power spectrum at $z = 0$ as measured in spheres of radius $8 \text{ h}^{-1} \text{ Mpc}$, $\sigma_8 = 0.8120$; and the cosmic baryon fraction, $f_b \equiv \Omega_b/\Omega_m = 0.156$ (ref. 12).

With these values, the linear matter power spectrum is specified at all times relevant for structure formation. The non-linear density field, home to the dark matter haloes that host galaxies, must be computed numerically. However, a long line of research starting with Press and Schechter¹³ has been devoted to connecting the abundance of dark matter haloes as a function of redshift and mass to the underlying linear matter power spectrum. In what follows, I use the Sheth and Tormen¹⁴ dark matter halo mass function $dN(M, z)/dM$ —the number of dark matter haloes of mass M per unit mass per unit co-moving volume at redshift z —to compute the co-moving number density of haloes above a given halo mass threshold,

$$n(> M_{\text{halo}}, z) = \int_{M_{\text{halo}}}^{\infty} dM \frac{dN(M, z)}{dM} \quad (1)$$

and the co-moving mass density in haloes more massive than M_{halo} ,




Fig. 1 | Limits on the abundance of galaxies as a function of redshift. Curves show the relationship between M_* and z at fixed cumulative halo abundance (left) and fixed $\rho_b (>M_{\text{halo}})$, or equivalently fixed peak height v (right). The most extreme L23 galaxy candidates are shown as blue stars, with uncertainties indicating 68% intervals (symmetric about the median) of the posterior probability distribution. The existence of a galaxy with M_* at redshift z requires that such galaxies have a cumulative co-moving number density that is, at most, the number density shown in the left panel, as those galaxies must reside in host halo of mass $M_{\text{halo}} = M_*/(f_b \epsilon)$.

The cumulative co-moving number density corresponding to an observed M_* will probably be (much) smaller than is indicated here, as the curves are placed on the plot by assuming the physically maximal $\epsilon = 1.0$. For smaller values of ϵ , the curves in each panel move down relative to the points by a factor of ϵ (as indicated by the black downward-facing arrows). The right panel demonstrates that even for the most conservative assumption of $\epsilon = 1.0$, the data points correspond to very rare peaks in the density field, implying a limited baryonic reservoir that is in tension with the measured stellar masses of the galaxies.

$$\rho_m(> M_{\text{halo}}, z) = \int_{M_{\text{halo}}}^{\infty} dM M \frac{dn(M, z)}{dM}. \quad (2)$$

These translate directly to upper limits on the statistics of galaxies through the straightforward assumption that the largest stellar content a halo can have, given its cosmic allotment of baryons, is $M_{*,\text{max}} = f_b M_{\text{halo}}$. More generally, we may write $M_* = \epsilon f_b M_{\text{halo}}$, with $\epsilon \leq 1.0$ being the efficiency of converting baryons into stars.

The cumulative co-moving number density of dark matter haloes more massive than M_{halo} thus sets an upper limit on the co-moving number density of galaxies more massive than M_* ,

$$n_{\text{gal}}(> M_*) \leq n_{\text{halo}}(> M_*/f_b). \quad (3)$$

Similarly, the cumulative co-moving density of collapsed mass sets an upper limit on the density of collapsed baryons, $\rho_b(>M_{\text{halo}}) = f_b \rho_m(>M_{\text{halo}})$, which in turn strictly bounds the co-moving mass density of stars contained in haloes more massive than M_{halo} ,

$$\rho_*(> M_{\text{halo}}) \leq f_b \rho_m(> M_{\text{halo}}), \quad (4)$$

and the density of stars contained in galaxies above a given M_* ,

$$\rho_*(> M_*) \leq f_b \rho_m(> M_*/f_b). \quad (5)$$

Results

The left panel of Fig. 1 shows the relationship between the maximal inferred stellar mass for a given M_{halo} , $M_* = f_b M_{\text{halo}}$ (that is, assuming the maximal $\epsilon = 1.0$) and redshift z for fixed cumulative co-moving halo number densities ranging from $10^{-10} \text{ Mpc}^{-3}$ (light grey) to 10^{-2} Mpc^{-3} (yellow). The curves evolve rapidly with redshift, with the maximal stellar mass corresponding to a fixed cumulative co-moving halo number

density increasing by three orders of magnitude from $z = 20$ to $z = 5$. This rapid rise indicates that the mass reservoir available for the most massive galaxies increases quickly with redshift at fixed halo number density. The two most massive high-redshift galaxy candidates from the Labbé et al.⁸ (hereafter L23) sample, at $z \approx 7.5$ ($M_* \approx 10^{11} M_\odot$) and $z \approx 9.1$ ($M_* \approx 10^{10.5} M_\odot$), are shown as blue stars. These objects are unexpectedly massive, with stellar content reflective of haloes that have cumulative co-moving number densities no higher than $\approx 10^{-5.2} \text{ Mpc}^{-3}$ (if $\epsilon = 1.0$); for $\epsilon = 0.32$ (0.10), the implied number density is $\approx 10^{-7}$ ($10^{-9.3}$) Mpc^{-3} . By comparison, the candidates were found in a survey of 38 arcmin^2 , a volume of $V \approx 10^5 \text{ Mpc}^3$ at each of the redshift bins $-7 < z < 8.5$ and $8.5 < z < 10$ —considered by L23.

The right panel of Fig. 1 recasts the issue in terms of the scarcity of systems as measured by cumulative mass density. In extended Press–Schechter models, the peak height $v(M_{\text{halo}}, z) = \delta_c / \sigma(M_{\text{halo}}, z)$ of an object—where $\delta_c \approx 1.7$ is the linear collapse threshold and $\sigma^2(M_{\text{halo}}, z)$ is the variance of the linear density field at redshift z smoothed on a scale containing an average mass of M_{halo} —is a measure of the fraction of mass in the Universe contained in virialized objects more massive than M_{halo} at redshift z . Typical haloes at z have $v = 1$, which corresponds to 24% of the mass in the Universe residing in haloes at least that massive; larger values of v indicate increasingly massive and therefore rare peaks in the density field at that epoch. The co-moving baryon density for each peak height in the figure is given in the legend; multiplying this number by the volume of a survey gives the total amount of baryons contained above the mass corresponding to that peak height and redshift. The L23 galaxies have peak heights of at least $v = 4.5$ (assuming $\epsilon = 1.0$), meaning that, at most, a fraction 6.2×10^{-5} of the baryons in the Universe are contained in haloes massive enough to host these galaxies. For reference, $v = 4.5$ at $z = 0$ corresponds to $M_{\text{halo}} \approx 5 \times 10^{15} M_\odot$. Adopting more reasonable efficiencies of $\epsilon = 0.32$ or 0.10 results in rarer peaks with $v \approx 5.4$ or 6.4 .

Figure 2 shows the cumulative stellar mass density reported by L23 at $z \approx 9$ (left) and $z \approx 7.5$ (right). The data, which come from individual

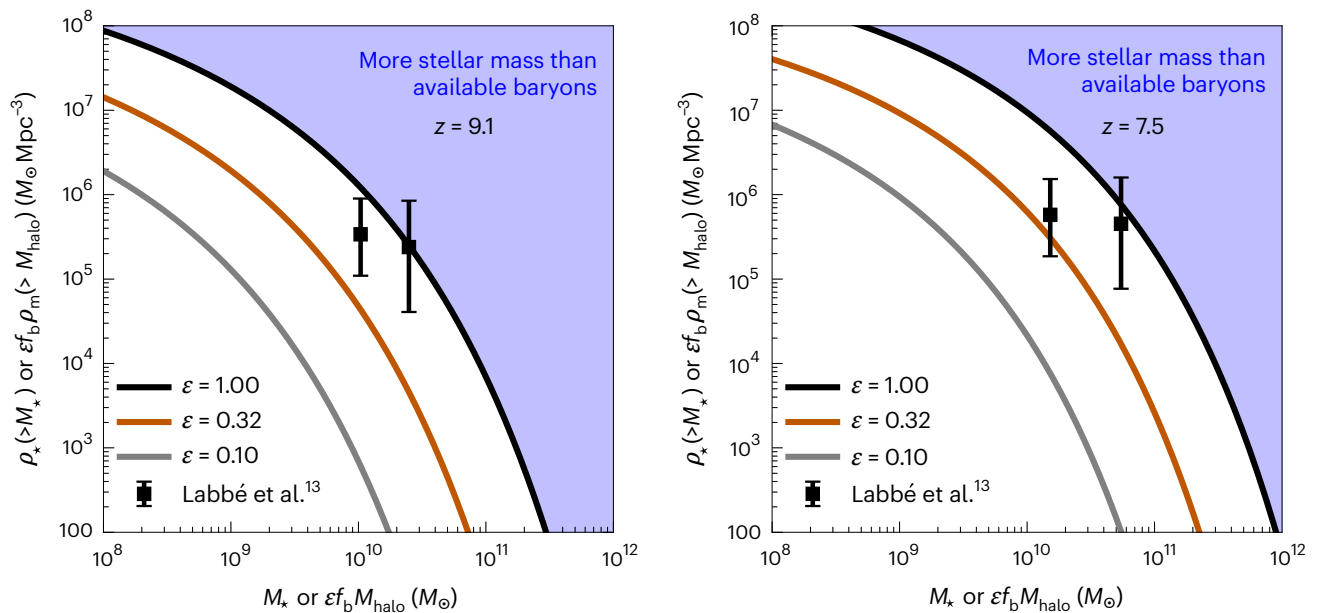


Fig. 2 | Stellar mass density limits. The co-moving stellar mass density contained within galaxies more massive than M_* at $z \approx 9.1$ (left) and $z \approx 7.5$ (right) for three values of the assumed conversion efficiency ϵ of a halo's cosmic allotment of baryons into stars. Only if all available baryons in all haloes with enough baryons to form the galaxies reported by L23 have indeed been converted into stars by that point—an unrealistic limit—is it possible to produce the stellar mass density in the highest M_* bin at $z \approx 9$ measured by L23 in a

typical volume of a Λ CDM Universe with the Planck 2020 cosmology. Results are similar at $z \approx 7.5$. For more realistic values of ϵ , the required baryon reservoir is substantially larger than the theoretical maximum in this cosmology. When considering 1σ shot noise and sample variance errors added in quadrature (which comprise the uncertainties on the L23 data points in each panel), the measurements are consistent with the base Λ CDM model if $\epsilon > 0.57$, which would still imply incredibly efficient star formation in the high-redshift Universe.

massive objects, lie at the extreme of Λ CDM expectations, even in the most optimistic scenario: at both redshifts, the measurements lie at the theoretical limit of $\rho_*(> M_*) = f_b \rho_m(> M_*/f_b)$, implying physically implausible values of $\epsilon(z \approx 9) = 0.99$ and $\epsilon(z \approx 7.5) = 0.84$. When considering the 1σ error (which incorporates uncertainties from Poisson fluctuations and sample variance added in quadrature), the data become marginally consistent with the available baryon reservoirs for an efficiency of $\epsilon(z \approx 9) \geq 0.57$, which is probably an unrealistically high value. Assuming a more plausible value of $\epsilon = 0.10$ or 0.32 yields a strong discrepancy with Λ CDM expectations at both redshifts, even when considering observational uncertainties.

Discussion

The first glimpse of high-redshift galaxy formation with JWST has revealed surprisingly massive galaxy candidates at early cosmic times. These systems provide a way to test a bedrock property of the Λ CDM model (or alternately, assumptions in derivations of stellar masses or the viability of high-redshift galaxy candidates): the stellar content of haloes should not exceed the available baryonic material in those haloes. This requirement does not rely on assumptions such as abundance matching, but rather is simply a statement about the distribution of virialized mass in the Universe as a function of redshift and the baryonic reservoirs associated with those virialized haloes: galaxies of mass M_* can only form if haloes of mass $M_*/(\epsilon f_b)$ have formed. It is also more stringent than the requirement that the observed galaxy ultraviolet luminosity function not exceed the theoretical maximum coming from a nearly instantaneous (10 Myr) conversion of a halo's full baryonic reservoir into stars¹⁵, as it is an integral constraint as opposed to a differential one. The massive, high-redshift galaxy candidates catalogued in L23 lie near or at the stellar mass density constraint in Λ CDM.

There are several sources of observational uncertainty that enter these results. The flux calibration of NIRCcam is continually being updated; L23 use calibrations that take into account updated detector offsets that are not yet part of the official JWST reduction pipeline

(see, for example, Boyer et al.¹⁶ for examples of this effect and Nardiello et al.¹⁷ for related discussions of empirical point spread function modelling for JWST). With NIRCcam photometry, a Balmer α break at $z \approx 5$ can be mistaken for a Lyman α break at $z \geq 12$ (ref. 18); the L23 sample was selected to contain both Lyman and Balmer breaks, however, and is at low enough redshift (relative to $z \approx 15$ sources) that NIRCcam filters can typically exclude $z \approx 5$ photometric solutions. The resulting photometric redshift estimates have single, narrow ($\sigma_z \approx 0.25$) peaks. The masses of the galaxies are computed using the median of four methods for fitting the photometry (see L23 for details) and assume a Salpeter¹⁹ initial mass function. Different assumptions about the photometry (in particular, properties of nebular emission lines or the presence of an accreting supermassive black hole) or initial mass function could affect the derived stellar masses. The mass of the candidate at $z \approx 7.5$ was also corrected for the possibility of amplification by mild gravitational lensing; this effect is estimated by L23 to be 0.15 dex, and the reported mass (and stellar mass density) of this object are therefore reduced by this amount to compensate. The error bars in Fig. 2 include errors in the volume estimates coming from both sample variance and Poisson noise, with the latter always being dominant in the regime considered here¹²⁰.

The discrepancy between the observed high-redshift galaxy candidates and Λ CDM expectations is robust to uncertainties in cosmological parameters in the base Λ CDM model: the precision on each of the relevant parameters is at the $\lesssim 1\%$ level¹². Intriguingly, extensions to the base Λ CDM with enhanced values of σ_8 , n_s and the physical matter density $\Omega_m h^2$ —such as some early dark energy (EDE) models whose aim is to resolve the Hubble tension—predict earlier structure formation and a higher abundance of haloes at fixed mass at high redshift²¹, which would enhance the baryonic reservoirs available for forming early massive galaxies. Taking the best-fit EDE parameters from Smith et al.²², the cumulative co-moving baryonic density contained in galaxies more massive than $M_* = f_b M_{\text{halo}}$ for the most massive L23 galaxy candidate at $z \approx 9.1$ is a factor of 3.1 larger in EDE than in base Λ CDM, which is

non-negligible; the L23 data points would then lie at $\epsilon = 0.74$ instead of $\epsilon = 0.99$. However, this EDE cosmology is in stronger tension with values of $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$ measured at low redshift and predict that the Universe is ≈ 13 billion years old (as opposed to 13.8 billion years in the base Λ CDM model), which is in moderate tension with the measured ages of ultra-faint galaxies and globular clusters²³.

At the redshifts studied here, $z \approx 7\text{--}10$, the Sheth–Tormen mass function overestimates the abundance of massive haloes by 20–50% relative to numerical simulations^{24–27}, meaning their true abundance at high redshift is probably lower than the Sheth–Tormen prediction and the constraints derived here are conservative. However, the lack of detailed comparisons between theory and simulations at high redshifts and high masses points to the importance of continued theoretical work in understanding the universality and applicability of halo mass function parameterizations in regimes relevant for JWST observations (and other forthcoming observatories).

The tension discussed in this paper is straightforward: the masses measured by L23 are only consistent with expectations from the standard cosmological model at the reported redshifts if star formation in the earliest phases of galaxy formation is incredibly efficient ($\epsilon \geq 0.57$). In the low-redshift Universe, such efficiencies are never seen, with $\epsilon \lesssim 0.2$ for all galaxies. The theoretical expectation is that efficiencies do indeed increase at high redshift²⁸, although $\epsilon \geq 0.57$ is still highly extreme and probably implausibly high. If the explanation of the L23 galaxies is indeed a very high star formation efficiency, it implies that the star formation histories of such systems must rise steeply with time, following the behaviour of the baryon reservoirs inside virialized structures in Λ CDM. The results presented here could also be explained if the stellar initial mass function differs substantially from the assumed Salpeter form, the black hole accretion contributes significantly to the galaxies' emitted light or the volumes currently surveyed turn out to be highly atypical.

If none of these explanations holds up and these massive galaxies are spectroscopically confirmed, they will pose a serious challenge for Λ CDM structure formation with parameters given by Planck Collaboration¹² because they signify the existence of a larger reservoir of collapsed baryons than is possible in this model. Forthcoming wider field JWST surveys, along with JWST spectroscopy of massive galaxy candidates, should be able to quickly confirm or refute the existence of this tension. Furthermore, the compatibility of any additional high-redshift galaxies or galaxy candidates discovered in JWST observations with Λ CDM expectations can be assessed in a straightforward way via Fig. 1. If analysis of JWST data continues to reveal the presence of strikingly massive galaxies at very early cosmic epochs, more exciting surprises lie ahead for the fields of galaxy formation and cosmology.

Data availability

Data from L23, including stellar mass estimates and photometric redshifts, are available at <https://github.com/ivolabbe/red-massive-candidates>; this paper uses data from sample_rev3_2207.12446.ecsv, commit 59fbbfa (from 2 January 2023).

Code availability

All calculations that go into the figures in this paper are publicly available at https://github.com/mrbk/JWST_MstarDensity.

References

- Behroozi, P. & Silk, J. The most massive galaxies and black holes allowed by Λ CDM. *Mon. Not. R. Astron. Soc.* **477**, 5382–5387 (2018).
- Steinhardt, C. L., Capak, P., Masters, D. & Speagle, J. S. The impossibly early galaxy problem. *Astrophys. J.* **824**, 21 (2016).
- Adams, N. J. et al. Discovery and properties of ultra-high redshift galaxies ($9 < z < 12$) in the JWST ERO SMACS 0723 field. *Mon. Not. R. Astron. Soc.* **518**, 4755–4766 (2023).
- Atek, H. et al. Revealing galaxy candidates out to z 16 with JWST observations of the lensing cluster SMACS0723. *Mon. Not. R. Astron. Soc.* **519**, 1201–1220 (2023).
- Castellano, M. et al. Early results from GLASS-JWST. III. Galaxy candidates at z 9–15. *Astrophys. J. Lett.* **938**, L15 (2022).
- Donnan, C. T. et al. The evolution of the galaxy UV luminosity function at redshifts $z \approx 8\text{--}15$ from deep JWST and ground-based near-infrared imaging. *Mon. Not. R. Astron. Soc.* **518**, 6011–6040 (2023).
- Finkelstein, S. L. et al. A long time ago in a galaxy far, far away: a candidate z 12 galaxy in early JWST CEERS imaging. *Astrophys. J. Lett.* **940**, L55 (2022).
- Labbé, I., van Dokkum, P., Nelson, E. et al. A population of red candidate massive galaxies ~ 600 Myr after the Big Bang. *Nature* <https://doi.org/10.1038/s41586-023-05786-2> (2023).
- Morishita, T. & Stiavelli, M. Physical characterization of early galaxies in the Webb's First Deep Field SMACS J0723.3-7323. Preprint at *arXiv* <https://doi.org/10.48550/arXiv.2207.11671> (2022).
- Naidu, R. P. et al. Two remarkably luminous galaxy candidates at $z \approx 10\text{--}12$ revealed by JWST. *Astrophys. J. Lett.* **940**, L14 (2022).
- Yan, H., Ma, Z., Ling, C., Cheng, C. & Huang, J.-S. First batch of $z \approx 11\text{--}20$ candidate objects revealed by the James Webb Space Telescope early release observations on SMACS 0723-73. *Astrophys. J. Lett.* **942**, L9 (2023).
- Planck Collaboration. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.* **641**, A6 (2020).
- Press, W. H. & Schechter, P. Formation of galaxies and clusters of galaxies by self-similar gravitational condensation. *Astrophys. J.* **187**, 425–438 (1974).
- Sheth, R. K. & Tormen, G. Large-scale bias and the peak background split. *Mon. Not. R. Astron. Soc.* **308**, 119–126 (1999).
- Mason, C. A., Trenti, M. & Treu, T. The brightest galaxies at cosmic dawn. *Mon. Not. R. Astron. Soc.* **521**, 497–503 (2023).
- Boyer, M. L. et al. The JWST Resolved Stellar Populations Early Release science program. I. NIRCcam flux calibration. *Res. Not. Am. Astron. Soc.* **6**, 191 (2022).
- Nardiello, D. et al. Photometry and astrometry with JWST-I. NIRCcam point spread functions and the first JWST colour-magnitude diagrams of a globular cluster. *Mon. Not. R. Astron. Soc.* **517**, 484–497 (2022).
- Zavala, J. A. et al. Dusty starbursts masquerading as ultra-high redshift galaxies in JWST CEERS observations. *Astrophys. J. Lett.* **943**, L9 (2023).
- Salpeter, E. E. The luminosity function and stellar evolution. *Astrophys. J.* **121**, 161 (1955).
- Trenti, M. & Stiavelli, M. Cosmic variance and its effect on the luminosity function determination in deep high- z surveys. *Astrophys. J.* **676**, 767–780 (2008).
- Klypin, A. et al. Clustering and halo abundances in early dark energy cosmological models. *Mon. Not. R. Astron. Soc.* **504**, 769–781 (2021).
- Smith, T. L. et al. Hints of early dark energy in Planck, SPT, and ACT data: new physics or systematics? *Phys. Rev. D* **106**, 043526 (2022).
- Boylan-Kolchin, M. & Weisz, D. R. Uncertain times: the redshift–time relation from cosmology and stars. *Mon. Not. R. Astron. Soc.* **505**, 2764–2783 (2021).
- Reed, D. et al. Evolution of the mass function of dark matter haloes. *Mon. Not. R. Astron. Soc.* **346**, 565–572 (2003).
- Despali, G. et al. The universality of the virial halo mass function and models for non-universality of other halo definitions. *Mon. Not. R. Astron. Soc.* **456**, 2486–2504 (2016).
- Shirasaki, M., Ishiyama, T. & Ando, S. Virial halo mass function in the Planck cosmology. *Astrophys. J.* **922**, 89 (2021).

27. Wang, Q., Gao, L. & Meng, C. The ultramarine simulation: properties of dark matter haloes before redshift 5.5. *Mon. Not. R. Astron. Soc.* **517**, 6004–6012 (2022).
28. Tacchella, S., Bose, S., Conroy, C., Eisenstein, D. J. & Johnson, B. D. A redshift-independent efficiency model: star formation and stellar masses in dark matter halos at $z \geq 4$. *Astrophys. J.* **868**, 92 (2018).
29. Harris, C. R. et al. Array programming with NumPy. *Nature* **585**, 357–362 (2020).
30. Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).
31. Hunter, J. D. Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* **9**, 90–95 (2007).
32. Murray, S. G., Power, C. & Robotham, A. S. G. HMFcalc: an online tool for calculating dark matter halo mass functions. *Astron. Comput.* **3**, 23–24 (2013).
33. Murray, S. HMF: Halo Mass Function calculator. Astrophysics Source Code Library, record ascl:1412.006 (2014).
34. Pérez, F. & Granger, B. E. IPython: a system for interactive scientific computing. *Comput. Sci. Eng.* **9**, 21–29 (2007).

Acknowledgements

This paper is dedicated to the memory of Steven Weinberg, who would have been thrilled to see how well JWST is working and excited to learn what it will reveal about cosmology and galaxy formation across a variety of cosmic epochs. I thank P. van Dokkum and I. Labbé for sharing data from L23 and S. Finkelstein, P. Kumar and D. Weisz for helpful discussions. I acknowledge support from the University of Texas at Austin through the Faculty Research Assignment program, NSF CAREER award AST-1752913, NSF grants AST-1910346 and AST-2108962, NASA grant 8ONSSC22K0827, and HST-AR-15809, HST-GO-15658, HST-GO-15901, HST-GO-15902, HST-AR-16159 and HST-GO-16226 from the Space Telescope Science Institute, which is operated by AURA, Inc., under NASA contract NAS5-26555. I am grateful to the developers of the Python packages used in preparing this paper: NumPy²⁹, SciPy³⁰, Matplotlib³¹, HMF^{32,33}

and IPython³⁴. This research has made extensive use of NASA's Astrophysics Data System (<http://adsabs.harvard.edu/>) and the arXiv e-Print service (<http://arxiv.org>).

Competing interests

The author declares no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Michael Boylan-Kolchin.

Peer review information *Nature Astronomy* thanks the anonymous reviewers for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023