


Research Article

DARK SAGE: Next-generation semi-analytic galaxy evolution with multidimensional structure and minimal free parameters

Adam R. H. Stevens^{1,2} , Manodeep Sinha^{2,3}, Alexander Rohl^{4,5}, Mawson W. Sammons⁶, Boryana Hadzhiyska^{7,8}, César Hernández-Aguayo⁹ and Lars Hernquist¹⁰

¹International Centre for Radio Astronomy Research, The University of Western Australia, Crawley, WA 6009, Australia, ²Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia, ³Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia, ⁴School of Computer and Mathematical Sciences, The University of Adelaide, Adelaide, SA 5000, Australia, ⁵Department of Mathematics and Statistics, The University of Western Australia, Crawley, WA 6009, Australia, ⁶International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia, ⁷Miller Institute for Basic Research in Science, University of California, Berkeley, CA 94720, USA, ⁸Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA, ⁹Max-Planck-Institut für Astrophysik, D-85741 Garching, Bayern, Germany and ¹⁰Institute for Theory and Computation, Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138, USA

Abstract

After more than five years of development, we present a new version of DARK SAGE, a semi-analytic model (SAM) of galaxy formation that breaks the mould for models of its kind. Included among the major changes is an overhauled treatment of stellar feedback that is derived from energy conservation, operates on local scales, affects gas gradually over time rather than instantaneously, and predicts a mass loading factor for every galaxy. Building on the model's resolved angular momentum structure of galaxies, we now consider the heating of stellar discs, delivering predictions for disc structure both radially and vertically. We add a further dimension to stellar discs by tracking the distribution of stellar ages in each annulus. Each annulus–age bin has its own velocity dispersion and metallicity evolved in the model. This allows DARK SAGE to make structural predictions for galaxies that previously only hydrodynamic simulations could. We present the model as run on the merger trees of the highest-resolution gravity-only simulation of the MillenniumTNG suite. Despite its additional complexity relative to other SAMs, DARK SAGE only has three free parameters, the least of any SAM, which we calibrate exclusively against the cosmic star formation history and the $z = 0$ stellar and HI mass functions using a particle-swarm optimisation method. The DARK SAGE codebase, written in C and PYTHON, is publicly available at <https://github.com/arhstevens/DarkSage>.

keywords: galaxies: evolution; galaxies: haloes; galaxies: interactions; galaxies: ISM

(Received 7 December 2023; revised 14 February 2024; accepted 16 February 2024)

1. Context

Research into galaxy formation remains among the top priorities in the field of astronomy (National Academies of Sciences, Engineering, and Medicine 2021). With the likes of the James Webb Space Telescope in full operation and construction of the Square Kilometre Array underway, this is a promising decade for the field. The endgame of these multi-billion-dollar enterprises is to develop a comprehensive working theory for how the Universe works, of which galaxy formation is a major, crucial part. But the flood of incoming high-quality observations is only as powerful as the models we use to interpret those data. Galaxy formation is a highly complex problem to model, which requires us to ‘solve’ it with numerical simulations. While every model is flawed, simulations are the closest that extragalactic astrophysicists will ever come to conducting a controlled experiment. Having a variety

of galaxy formation models, frameworks, and simulations — alongside observations — is paramount to test and develop our theoretical understanding of not just galaxies themselves, but cosmology too.

Two methods for simulating the formation of galaxies in a cosmological context have stood out in recent decades: semi-analytic models (SAMs) and hydrodynamic simulations. Hydrodynamic simulations simultaneously and self-consistently account for gravity, fluid dynamics, and all astrophysical processes deemed relevant to the formation of galaxies. SAMs separate the formation of the universe's large-scale structure from that of galaxies by first constructing a universe with only gravity, and subsequently evolving galaxies in the gravitationally bound structures in that universe — referred to as ‘haloes’ — with relatively macroscopic descriptions of baryonic astrophysics. In truth, hydrodynamic simulations are still ‘semi-analytic’ at the level of a particle, which is generally comparable in mass to a giant molecular cloud or globular cluster of stars. Or, to frame this in reverse, all baryonic processes are ‘subgrid’ at the level of the (sub)halo in a SAM. A holistic summary of these methods of galaxy formation modelling can be found in a number of review articles (e.g. Somerville & Davé 2015; Vogelsberger et al. 2020).

Corresponding authors: Adam R. H. Stevens and Manodeep Sinha; Emails: adam@a4e.org, msinha@swin.edu.au

Cite this article: Stevens ARH, Sinha M, Rohl A, Sammons MW, Hadzhiyska B, Hernández-Aguayo C and Hernquist L. (2024) DARK SAGE: Next-generation semi-analytic galaxy evolution with multidimensional structure and minimal free parameters. *Publications of the Astronomical Society of Australia* 41, e053, 1–35. <https://doi.org/10.1017/pasa.2024.14>

In this paper, we present a new version of the SAM known as ‘DARK SAGE’. DARK SAGE has been in development since 2015, with two main versions previously released (Stevens *et al.* 2016, 2018). The model started as a modified version of SAGE (Croton *et al.* 2016). While the code architecture is still based on SAGE, the physical models implemented in DARK SAGE are almost entirely unique. SAGE was itself built from the Croton *et al.* (2006) version of the Munich SAM — from which the likes of GAEA (Xie *et al.* 2020, and references therein) and L-GALAXIES (Henriques *et al.* 2020) also stem — with many versions preceding that too (Kauffmann *et al.* 1999; Springel *et al.* 2001; De Lucia, Kauffmann & White 2004). While the new DARK SAGE that we present here brings together five years of development since its last release, it builds on over 25 years of research in the community.

Outside of the main model papers (Stevens *et al.* 2016, 2018), DARK SAGE has been used to investigate the effects of environment and feedback on the gas content of galaxies (Stevens & Brown 2017), research the angular momentum of galaxies at both high redshifts (Okamura, Shimasaku & Kawamata 2018) and high stellar masses (Porras-Valverde *et al.* 2021), explain the origin of H I-excess galaxies (Lutz *et al.* 2018), confirm the origin of the H I size–mass relation (Stevens *et al.* 2019b), predict the results of H I intensity-mapping experiments (Wolz *et al.* 2022), develop an H I-dependent halo occupation distribution model for cosmology applications (Qin *et al.* 2022), and more.^a The significant updates that we present in this paper will not only expand the range of use cases for DARK SAGE, but also improve the predictive power for the use cases it was originally designed for.

We have endeavoured to adopt a new philosophy with how SAMs should be designed with this new version of DARK SAGE. Our main goal is to re-derive the core aspects of galaxy evolution physics from as near to first principles as possible; there are fewer *ad hoc* descriptions left, and the ones we do use are updated to modern measurements where possible, often over broader mass and redshift ranges. This has meant increasing the coupling between physical prescriptions, which results in the elimination of a large number of free parameters. Simultaneously, we have *increased* the detail and complexity of the galaxies we evolve, building their structure numerically in more than one dimension. Feedback from stars and AGNs — arguably the most important processes in galaxy evolution — are derived from the law of energy conservation. While the model is certainly not free of phenomenology, our philosophy has meant we can capture the uncertainty in our model almost completely with only three free parameters (a factor 3–6 lower than comparable models tabulated in Table 2), each with well defined boundaries, which we successfully calibrate to observational statistics of galaxies with an objective, automated method.

DARK SAGE is a publicly available, version-controlled code (Stevens *et al.* 2017b), open for the community to use and modify as they desire. The aim of this paper is to describe how the model works, and should serve especially well for those using the code itself. While we present some predictions of the model here, we deliberately do not linger on them. Science questions and specific results will be better served as the subject of future, focussed papers.

^aFor a manually updated list of papers that have used results from any of the SAGE family of models, see <https://ui.adsabs.harvard.edu/public-libraries/JslFrCAISE2wzNqL5trCwg>.

2. ‘The next generation’ N-body simulations

The results of DARK SAGE in this paper are produced by running the model on the merger trees of the main gravity-only MillenniumTNG N-body simulation. The MillenniumTNG simulations^b (Barrera *et al.* 2023; Bose *et al.* 2023; Contreras *et al.* 2023; Delgado *et al.* 2023; Ferlito *et al.* 2023; Hadzhiyska *et al.* 2023a,b; Hernández-Aguayo *et al.* 2023; Kannan *et al.* 2023; Pakmor *et al.* 2023) are the latest addition to ‘The Next Generation’ (TNG) simulation suite, which began with IllustrisTNG^c (Weinberger *et al.* 2017; Marinacci *et al.* 2018; Naiman *et al.* 2018; Nelson *et al.* 2018, 2019b; Pillepich *et al.* 2018a,b; Springel *et al.* 2018). Each simulation in the TNG suite has paired N-body and magnetohydrodynamic versions with identical initial conditions, run with the same code [AREPO (Springel 2010), although N-body runs have also been completed with GADGET-4 (Springel *et al.* 2021)], and all carry the same Λ CDM cosmology: $\Omega_m = 0.3089$, $\Omega_\Lambda = 0.6911$, $\Omega_b = 0.0486$, $h = 0.6774$, $\sigma_8 = 0.8159$, $n_s = 0.9667$ (Planck Collaboration 2016). We assume this cosmology for all results in this paper (including observational data that we compare to, where we have adjusted published data accordingly). The MillenniumTNG run used in this paper (hereafter ‘MTNG’, referred to as ‘MTNG740-DM-1’ in other papers) has a comoving box length of $500 h^{-1}$ cMpc (nearly 740 cMpc) with 4320^3 particles of mass $1.329 \times 10^8 h^{-1} M_\odot$ ($= 1.962 \times 10^8 M_\odot$), and was run with AREPO.

There are many advantages to pairing a SAM like DARK SAGE with MTNG:

- Both the volume and resolution are high enough to simultaneously capture a large number of galaxy clusters with small enough galaxies to study environmental effects on galaxy evolution.
- Simulations with different mass resolutions have been run, allowing for future resolution studies.
- Objects from the SAM can be directly compared against the galaxies formed in the hydrodynamical run (Pakmor *et al.* 2023) in future studies.
- Galaxy catalogues from at least one other SAM have already been produced (Barrera *et al.* 2023, another descendant of the Munich model), allowing for additional future model comparison.
- With an equivalent volume and eight times the number of particles, MTNG is a natural successor to the original Millennium simulation (Springel *et al.* 2005), which many SAM results in the literature are based on, including previous iterations of DARK SAGE (Stevens *et al.* 2016, 2018).
- If a galaxy of a given stellar mass being ‘resolved’ in a simulation requires it to have a minimum number of particles, then by virtue of the fact that only a fraction of baryon particles in a hydrodynamic-simulation halo will become star particles, semi-analytic galaxies built in a simulation of fixed volume and dark-matter particle number are, in principle, resolved (and therefore complete) to lower stellar masses than a hydrodynamic simulation with those same specifications.

Haloes and their substructure (‘subhaloes’) are identified in each snapshot of MTNG with SUBFIND-HBT, described in sec-

^b<https://www.mtng-project.org/>

^c<https://www.tng-project.org/>

tion 7 of Springel et al. (2021). A minimum of 20 particles must be connected in a friends-of-friends (FoF) group to form a halo. When using the terms ‘halo mass’ or ‘virial mass’, we mean the mass enclosed by a sphere whose mean density is 200 times the critical density of the (simulated) universe (denoted by subscript ‘200c’). This has a one-to-one relationship with both virial radius and virial velocity:

$$M_{200c} = \frac{100 H^2(z) R_{200c}^3}{G} = \frac{V_{200c}^3}{10 G H(z)}, \quad (1)$$

where $H(z)$ is the redshift-dependent fractional expansion rate of the universe or ‘Hubble parameter’.

The merger trees for MTNG were originally constructed with the new GADGET-4 tree builder (see section 7.4 of Springel et al. 2021). This is similar to the format of LHALOTREE, originally presented in Springel et al. (2001) and used for the Millennium simulation (see the supplementary material of Springel et al. 2005) and is now in HDF5 format. DARK SAGE requires input merger trees in the original binary format of LHALOTREE. We thus converted the MTNG trees back to the old LHALOTREE format for this project.

While the *final* results in this paper are presented with the full MTNG box, all preliminary results and model calibration were performed with a smaller simulation for practical purposes. ‘Mini-MTNG’, as we refer to it, had identical parameters, including mass resolution, but only included 540^3 particles in a $62.5 h^{-1}$ cMpc box (MTNG’s equivalent of ‘Mini-Millennium’).

3. Definitions and design of baryonic reservoirs

Baryonic matter [gas, stars, and black holes (BHs)] in DARK SAGE galaxies and haloes is broken into discrete components. Indeed, the primary function of a SAM is to calculate how mass and metals move between these components. For basic context and to familiarise the reader with the nomenclature used throughout this paper, we briefly describe each component in this section.

3.1. Gas components

The gas reservoirs in DARK SAGE galaxies and haloes include:

- The interstellar medium (ISM) or ‘cold gas’, which refers to gas in a galaxy’s disc. Each disc is broken into N_{ann} annular segments, each fixed *a priori* in terms of specific angular momentum. Each segment has fields for its mass, metallicity, and radial boundaries. The ISM is assumed to be axisymmetric and thin, but its radial structure is built numerically over time in the model. We use the subscript ‘cold’ to refer to ISM properties in equations, even though this medium includes some ionised gas.
- The ‘hot gas’ reservoir, which represents the bulk of the circumgalactic medium (CGM), that is, the gas in the halo that is distinct from that in the galaxy. This reservoir is always assumed to be homogenous in terms of its temperature and metallicity, and therefore is described entirely by two fields: its total mass and total metal content. It is assumed to follow an analytic density profile. ‘Hot gas’ represents that which is available to be accreted onto the galaxy, subject to cooling (i.e. be it through the ‘hot mode’ or ‘cold mode’).
- The ‘fountain’ reservoir, which accounts for gas in the CGM that has recently been reheated out of the ISM and is in the

process of mixing with the hot gas. It is assumed to follow the same density profile and have the same specific energy as the hot-gas reservoir. We label this reservoir ‘fount’ in equations.

- The ‘outflowing’ reservoir, abbreviated to ‘outfl’ in equations. Similar to the fountain reservoir, this represents gas in the CGM that has been reheated due to feedback. But its specific energy is higher, such that it will escape the halo after an appropriate time.
- The ‘ejected’ reservoir, denoted by subscript ‘ejec’, which represents gas that has been expelled from the halo’s virial radius due to feedback (after transiting the outflow reservoir). It is reincorporated into the CGM over time.
- The ‘local intergalactic medium’ (LIGM), which represents the remaining gas in the halo that is outside the virial radius of all its subhaloes. This is a single field per halo; in practice, it is associated with the central subhalo. The LIGM primarily provides a means of tracking metal enrichment in cosmological gas that can eventually be accreted by the central galaxy.

Whenever we generically refer to the ‘CGM’ in text and equations, we mean the sum of the hot, outflowing, and fountain reservoirs:

$$m_{\text{CGM}} \equiv m_{\text{hot}} + m_{\text{fount}} + m_{\text{outfl}}. \quad (2)$$

The temperature and density profile of all 3 CGM components is assumed to be the same.

3.2. Stellar components

The stellar components in DARK SAGE galaxies and haloes include:

- The stellar disc, which — like the ISM — is broken into a series of annuli. These annuli are fixed in terms of specific angular momentum in the same way as the ISM, but the stellar disc need not always be coplanar with the ISM. We denote fields concerning the stellar disc with the subscript ‘*,disc’.
- The instability-driven bulge, which — as its name suggests — is built from disc instabilities driving stellar mass into a galaxy’s centre. This component has zero angular momentum by definition, and is therefore supported against gravity entirely by random motions.
- The merger-driven bulge, which is more akin to a classical galaxy bulge or ellipsoid. This constitutes stars that enter the bulge as a *direct* result of galaxy mergers. We calculate a nominal specific angular momentum for the merger-driven bulge.
- Intrahalo stars, otherwise referred to as the stellar halo or ‘IHS’ for short. This is built up through tidal stripping and the disruption of satellite galaxies before they have the chance to merge. By definition, this exclusively tracks stars inside the virial radius of a subhalo. Satellites can have a non-zero IHS.
- Local intergalactic stars (LIGS), the stellar analogue for the LIGM. It is built up similarly to IHS, but accounts exclusively for tidally stripped stellar mass in a halo that falls *outside* the virial radius. This field is always zero for satellites by construction.

One major infrastructural change we have made to the stellar components is that each one is now broken into N_{age} age bins. This means there are $(N_{\text{ann}} + 4) \times N_{\text{age}}$ stellar-mass and

stellar-metallicity fields for each galaxy (the ‘+4’ accounts for the spheroidal components and the LIGS).

The introduction of stellar-age bins serves several purposes that expand DARK SAGE’s capabilities. For one, it means it is possible to reconstruct the formation history of any disc segment or bulge component of a galaxy (or the galaxy as a whole) from a single-snapshot output. This means DARK SAGE can now predict age gradients in stellar discs and even three-dimensional age–metallicity–radius maps of *populations* of galaxies, which can be tested against modern observations (e.g. with survey data from integral field spectrographs). In this sense, DARK SAGE galaxies have ‘multidimensional’ structure, à la the title of this paper. Component-wise star formation histories of DARK SAGE galaxies are trivially reconstructed from a single-snapshot output, allowing for easy comparison with the outcome of models applied to the spectral energy distributions of survey galaxies by component (notably, e.g. Robotham, Bellstedt & Driver 2022). The age bins also allow us to add detail to our treatment of stellar evolution and feedback, where we now model the distribution of supernovae as a function of time since birth for stellar populations (Section 8).

N_{age} is a user-defined integer. In this work, we take $N_{\text{age}} = 30$, which is high enough to ensure that the delayed feedback scheme is sufficiently converged with N_{age} (not shown here) and to recover the shape of galaxies’ star formation histories, but not so high that it prohibitively slows the code. The age bracket of each age bin is interpolated from the snapshot times of the underlying N -body simulation (see Section 3.5); N_{age} is therefore capped by the number of time steps in the merger trees. If $N_{\text{age}} = N_{\text{snap}} - 1$, then each age bin has a one-to-one correspondence with snapshot intervals in the mergers trees. Later, in Figure 7, we show the age bins used in this work.

Whenever we refer to ‘the bulge’ without further specificity (or use subscript ‘*,bulge’), we mean the combination of the instability- and merger-driven bulges. Whenever we refer to the total stellar content of a galaxy (or shorthand ‘**’), we mean the combination of the stellar disc and both bulges, but *not* intracluster stars (nor the LIGS).

3.3. Disc annuli

We maintain the design of Stevens *et al.* (2016) in keeping the annular boundaries for DARK SAGE discs as

$$\frac{j_{\text{outer}}^{(i)}}{\text{kpc km s}^{-1}} = 1.4^{i-1} h^{-1}, \quad (3a)$$

$$j_{\text{inner}}^{(i)} = \begin{cases} 0, & i = 1 \\ j_{\text{outer}}^{(i-1)}, & i \geq 2 \end{cases}, \quad (3b)$$

$$i \in [1, 2, \dots, N_{\text{ann}} - 1, N_{\text{ann}}], \quad (3c)$$

with $N_{\text{ann}} = 30$. Throughout this paper, we use the superscript ‘(i)’ to refer specifically to the i th annulus. When no such superscript appears for a disc property, we are referring to that property for the *whole* disc. For convenience, in parts of this paper, we will also refer to the mean properties of an annulus, for example,

$$\langle j^{(i)} \rangle \equiv \frac{j_{\text{inner}}^{(i)} + j_{\text{outer}}^{(i)}}{2}. \quad (4)$$

The above applies for both stellar and gas discs. Whenever we refer to ‘the disc’, we mean the combination of the ISM and stellar disc.

3.4. BH components

We expand the consideration of where massive BHs can exist in DARK SAGE. Each subhalo now is now granted several BH reservoirs:

- The central BH of the galaxy maintains the definition as in previous versions of DARK SAGE. It is what powers AGN feedback. Each galaxy is effectively assumed to hold a single central BH; when two galaxies merge, their BHs are assumed to merge too.
- The ‘intra halo black holes’ reservoir (IHBH) tracks the total number and summed mass of BHs that end up inside the halo but *outside* the galaxy itself. This reservoir grows through the disruption of satellite galaxies inside R_{200c} of their parent halo. This reservoir has no effect on galaxy evolution.
- The reservoir for ‘local intergalactic black holes’ (LIGBH) similarly tracks the number and summed mass of BHs stripped from satellites outside the virial radius. Satellite galaxies themselves always have zero LIGBHs by construction.

3.5. Time-steps

In the vernacular of this paper, a full ‘time-step’ refers to the interval between snapshots in the merger trees of the N -body simulation that DARK SAGE is run on. MTNG has 264 time-steps, starting at $z \simeq 29.3$ and finishing at $z = 0$ (the initial conditions were produced at $z = 63$). The separation between each pair of the first 31 snapshots is $\Delta \log_{10}(a) = 0.014$, where a is the cosmological expansion factor. The interval in $\Delta \log_{10}(a)$ is halved for the next 33 snapshots and halved again for the remainder.

The movement of mass, metals, angular momentum, et cetera, between reservoirs in DARK SAGE is ultimately governed by a series of coupled differential equations. These equations are solved numerically and sequentially. Rather than setting the ‘unit’ of time used in the numerical solution to that of the time-step in the merger trees, we break each time-step into three ‘sub-time-steps’ in DARK SAGE.

The use of sub-time-steps helps mitigate the fact that we must chose an order in which we apply our astrophysical processes; in principle, most processes should happen simultaneously (though, some have a natural order in which they should be done). This way, we loop through all physical processes multiple times over a time-step, rather than once. Each sub-time-step within the same full time-step covers the same interval in cosmic time. We denote the cosmic-time interval of a sub-time-step as Δt throughout this paper.

In earlier versions of DARK SAGE, there were 10 sub-time-steps between each time-step. Those versions were based on the original Millennium simulation, which only had 64 snapshots in its merger trees, versus 265 in MTNG. The *total* number of sub-time-steps has therefore moderately increased from those versions (640 then versus 792 now).

Rather than explicitly writing out the network of coupled differential equations in DARK SAGE altogether, we present equations in this paper in the form that they are solved within a

sub-time-step. This best reflects how the equations are written in the codebase.

4. Baryon accretion, gas cooling, and disc formation

When a halo grows in mass — that is, when stepping between snapshots in the underlying merger trees — mass is added to the CGM (and/or occasionally the IHS; see below) of the central to ensure the baryonic mass fraction of the halo remains in line with the cosmic baryon fraction, $f_{\text{bary}}^{\text{cosmic}}$, modulo a suppression factor, f_{γ}^{+} , associated with photoionisation heating (Efstathiou 1992; Gnedin 2000). That is, we enforce

$$f_{\text{bary}}^{+} f_{\text{bary}}^{\text{cosmic}} M_{200\text{c}} = m_{\text{ejec}}^{\text{cen}} + \sum_{\text{gal}}^{R_{\text{gal}} < R_{200\text{c}}} \left(m_{*}^{\text{gal}} + m_{\text{cold}}^{\text{gal}} + m_{\text{CGM}}^{\text{gal}} + m_{\text{IHS}}^{\text{gal}} + m_{\text{BH}}^{\text{gal}} + m_{\text{IHBH}}^{\text{gal}} \right), \quad (5)$$

where the sum over ‘gal’ includes the central galaxy and all satellites within the virial radius. There is an important subtlety in how the ejected reservoir of the central is handled in equation (5). The entire ejected reservoir is defined to lie *outside* $R_{200\text{c}}$. It is included in the right-hand side of equation (5) because that gas *was* inside $R_{200\text{c}}$; this gas is ejected by stellar and AGN feedback (described in Sections 8 and 9), and there is no physical reason why that gas should be automatically replaced by other baryons (which would happen if it were not in the equation). In principle, this means that, even in the absence of photoionisation heating, the cosmic baryon fraction only provides an *upper limit* on the halo baryon fraction in DARK SAGE. We can straightforwardly calculate the baryon fraction inside the virial radius of a halo as

$$f_{\text{bary}}^{200\text{c}} = f_{\gamma}^{+} f_{\text{bary}}^{\text{cosmic}} - \frac{m_{\text{ejec}}^{\text{cen}}}{M_{200\text{c}}}. \quad (6)$$

The f_{γ}^{+} factor is adopted from DARK SAGE’s predecessors (Croton et al. 2006, 2016), which follow the ‘filtering mass’ fitting function of Kravtsov, Gnedin & Klypin (2004, see their appendix B). This depends on both halo mass and redshift and is primarily relevant during the epoch of reionisation ($z \gtrsim 6$, for all haloes) and for haloes with $M_{200\text{c}} < 10^{12} M_{\odot}$ (at all epochs). $f_{\text{bary}}^{\text{cosmic}}$ is set by the cosmology of the underlying N -body simulation ($f_{\text{bary}}^{\text{cosmic}} = \Omega_{\text{b}}/\Omega_{\text{m}} = 0.1573$ in this work).

In previous versions of the model, when a halo grew, pristine cosmological gas (i.e. with close to zero metallicity) was added to the hot-gas reservoir. This still happens as the norm, but with our new framework, preference is first given to mass in the local intergalactic reservoirs to be accreted into the halo. Should a halo have a non-zero LIGM and/or LIGS, mass is proportionally and respectively transferred from these to the hot-gas reservoir and the IHS. Should this be insufficient to raise the halo baryon fraction to $f_{\gamma}^{+} f_{\text{bary}}^{\text{cosmic}}$, then pristine cosmological gas is added to the hot gas to top it up. See Sections 11 and 12 for how the LIGM and LIGS are built up in the first place.

Gas must first pass through the CGM before reaching the ISM. The process for calculating how much mass is transferred from the CGM to the ISM in a given sub-time-step is similar in concept to previous incarnations of the model (which harken back to White & Frenk 1991) but is not identical.

Given that we assume hot gas to be spherically symmetric and homogenised, the cooling time for gas at a given radius in the CGM is calculated as

$$t_{\text{cool}}(R) \equiv \frac{e_{\text{hot}}^{\text{therm}}}{\dot{e}_{\text{hot}}^{\text{rad}}(R)} = \frac{3 T_{\text{hot}} \bar{\mu} m_{\text{p}} k_{\text{B}}}{2 \rho_{\text{hot}}(R) \Lambda(T_{\text{hot}}, Z_{\text{hot}})}, \quad (7)$$

where $e_{\text{hot}}^{\text{therm}}$ is the specific thermal energy of the CGM, $\dot{e}_{\text{hot}}^{\text{rad}}(R)$ is the rate of specific energy loss through radiation, T_{hot} is the CGM temperature, m_{p} is the mass of a proton, $\bar{\mu} = 0.59$ under the assumption that hot gas is fully ionised, k_{B} is the Boltzmann constant, $\rho_{\text{hot}}(R)$ is the hot-gas density profile, $\Lambda(T_{\text{hot}}, Z_{\text{hot}})$ is the tabulated cooling function from Sutherland & Dopita (1993), and Z_{hot} is the hot-gas metallicity. We maintain the assumption that the CGM temperature is the same as the virial temperature:^d

$$T_{\text{hot}} = T_{\text{CGM}} = T_{\text{vir}} \equiv \frac{\bar{\mu} m_{\text{p}} V_{200\text{c}}^2}{2 k_{\text{B}}}. \quad (8)$$

This formula for T_{CGM} is defined for centrals; after infall, satellites are assumed to have fixed T_{CGM} . We further maintain the definition of the ‘cooling radius’ as being that where the cooling time equals the halo dynamical time, that is,

$$t_{\text{cool}}(R_{\text{cool}}) = t_{\text{dyn}} \equiv \frac{R_{200\text{c}}}{V_{200\text{c}}} = \frac{1}{10 H(z)}. \quad (9)$$

However, what we *have* changed is the analytic profile for hot gas.

While previous iterations of SAGE and DARK SAGE assumed the CGM to follow the density profile of a singular isothermal sphere, we instead now adopt the so-called ‘beta’ profile. Nominally, a beta profile follows the expression

$$\rho_{\beta}(R) = \rho_0 \left[1 + \left(\frac{R}{R_{\text{c}}} \right)^2 \right]^{-3\beta/2}, \quad (10)$$

where $\rho_0 \equiv \rho_{\beta}(0)$ and R_{c} is the ‘core’ radius. For the application of this profile to the GALFORM SAM, modulo changes to the profile from feedback, Benson et al. (2003) assume that $\beta = 2/3$ and R_{c} is a fixed fraction of the virial radius of a halo. We follow this idea, but instead of allowing for the profile shape to be modified by galaxy evolution processes, we set an explicit redshift dependence on the latter, that is, $R_{\text{c}} = c_{\beta}(z) R_{200\text{c}}$. Finally, with recognition that the CGM is defined in this paper to be *within* $R_{200\text{c}}$ [i.e. a volume integral of $\rho_{\text{CGM}}(R)$ out to $R_{200\text{c}}$ must return m_{CGM}], and that each component of the CGM is assumed to follow the same density profile, we implement

$$\rho_{\text{hot}}(R) = \frac{m_{\text{hot}}}{m_{\text{CGM}}} \rho_{\text{CGM}}(R), \quad (11a)$$

$$\rho_{\text{CGM}}(R) = \frac{m_{\text{CGM}} C_{\beta}(z)}{4\pi c_{\beta}^2(z) R_{200\text{c}}^3} \left[1 + \left(\frac{R}{c_{\beta}(z) R_{200\text{c}}} \right)^2 \right]^{-1}, \quad (11b)$$

$$C_{\beta}(z) \equiv \left\{ 1 - c_{\beta}(z) \tan^{-1} \left[c_{\beta}^{-1}(z) \right] \right\}^{-1}. \quad (11c)$$

For the functional form of $c_{\beta}(z)$, we use the fitting function of Stevens et al. (2017a) to gas haloes from the EAGLE^e simulations:

$$c_{\beta}(z) = \max \left[0.05, 0.2 \exp(-1.5z) - 0.039z + 0.28 \right]. \quad (12)$$

^dIn practice, the temperature of the CGM should vary with radius. But cosmological hydrodynamic simulations have shown that it remains within a factor of 2 of T_{vir} out to $R_{200\text{c}}$, even with variation in how feedback is modelled (Stevens et al. 2017a).

^eEvolution and Assembly of GaLaxies and their Environment (Crain et al. 2015; Schaye et al. 2015).

The lower limit of 0.05 is imposed by hand; this represents the 16th percentile of haloes at the highest redshift ($z \simeq 4$) studied in Stevens *et al.* (2017a). As a point of comparison, in their models, Benson *et al.* (2003) and Font *et al.* (2008) adopt *default* (but not fixed) c_β values of 0.07 and 0.1, respectively.

For the CGM of satellite galaxies, we still assume a truncation radius. While it does not make complete physical sense to define R_{200c} for a satellite subhalo, we adopt M_{200c} as the smaller of either the present total subhalo mass or its pre-infall virial mass, then solve for an equivalent R_{200c} through equation (1). This ensures that as satellites are tidally stripped, their CGM truncation radius decreases too, which is physically expected. By contrast, as noted above, T_{vir} is deliberately *not* recomputed after infall; if it were, we would be implicitly assuming that tidal stripping causes a decrease in the average CGM temperature, which seems non-physical.

Combining equations (7–12), one can solve for R_{cool} . Because of the update in the hot-gas density profile, this solution is different from earlier versions of DARK SAGE.

Two modes of gas accretion onto a galaxy can take place; the ‘cold mode’ of accretion takes place when $R_{\text{cool}} > R_{200c}$, otherwise the galaxy is in the ‘hot mode’ (White & Frenk 1991). In concept, these are unchanged from earlier models. That is, in the cold mode, gas in the CGM that is en route to the ISM is limited more by the time it takes to fall onto the ISM than any time required to cool. In this instance, it becomes somewhat of a misnomer to call the CGM ‘hot gas’, but we nevertheless maintain the nomenclature for consistency with the literature. For a sub-time-step Δt , the mass deposited onto the ISM in the cold mode is taken as

$$\Delta m_{\text{hot} \rightarrow \text{cold}}^{\text{cold-mode}} = m_{\text{hot}} \frac{\Delta t}{t_{\text{dyn}}} - \Delta m_{\text{heat}}^{\text{radio}}. \quad (13)$$

The term $\Delta m_{\text{heat}}^{\text{radio}}$ represents an offset in accretion onto the ISM from radio-mode AGN feedback (defined and discussed in Section 9). While this term applies for both the cold and hot mode, it is normally only important in practice for the hot mode.

For the hot mode, the flux of mass cooling through R_{cool} is taken as equal to the mass deposition rate onto the galaxy (based on the results of Bertschinger 1989). To calculate this, we solve

$$\Delta m_{\text{hot} \rightarrow \text{cold}}^{\text{hot-mode}} = 4\pi \rho_{\text{hot}}(R_{\text{cool}}) R_{\text{cool}}^2 \left(\frac{dt_{\text{cool}}}{dR} \Big|_{R=R_{\text{cool}}} \right)^{-1} \Delta t - \Delta m_{\text{heat}}^{\text{radio}}. \quad (14)$$

Assuming equation (11), the solution can be written as

$$\Delta m_{\text{hot} \rightarrow \text{cold}}^{\text{hot-mode}} = \frac{m_{\text{hot}} R_{\text{cool}} C_\beta(z)}{2 R_{200c} t_{\text{dyn}}} \Delta t - \Delta m_{\text{heat}}^{\text{radio}}. \quad (15)$$

With $\Delta m_{\text{hot} \rightarrow \text{cold}}$ determined, the final thing to consider is the fraction of that mass that each annulus of the gas disc receives. As each annulus is bound by fixed values of specific angular momentum, this means defining a probability distribution function (PDF) of j for the cooling gas. We base this on the PDF of j for an exponential disc with a constant rotational velocity, that is,

$$\text{PDF}_{\text{cool}}(j) = \frac{4j}{j_{\text{cool}}^2} \exp\left(\frac{-2j}{j_{\text{cool}}}\right), \quad (16)$$

where $\int_0^\infty \text{PDF}_{\text{cool}}(j) dj = 1$ by definition, and j_{cool} is the mean of the PDF. Equation (16) here is simply a more correct and flexible way of writing equation 6 of Stevens *et al.* (2016). We emphasise that even though equation (16) is the same as an exponential disc

with a constant rotational velocity, that does *not* mean it is the only type of disc consistent with this PDF. That is, for any arbitrary surface density profile $\Sigma(r)$, one can find a rotational velocity profile $v_{\text{circ}}(r)$ that gives the same PDF of j as in equation (16). We are therefore not making any explicit assumptions about the form of $\Sigma(r)$ or $v_{\text{circ}}(r)$ as gas cools/accretes onto a galaxy; in fact, we solve for these profiles (using additional information per Section 5). Even *if* one were to make assumptions about the surface density profile or velocity profile of freshly cooled gas, one would still end up with a predictive model for the resulting disc structure after several time-steps, as both the value of j_{cool} and the vector along which accretion takes place change between time-steps.

What then sets j_{cool} ? The answer to this also depends on the accretion mode. In the cold mode, gas is assumed to be efficiently transported from the halo outskirts to the galaxy via cold streams (as shown to occur in hydrodynamic simulations, e.g. Kereš *et al.* 2005; Dekel & Birnboim 2006; van de Voort *et al.* 2011). In the absence of evidence to the contrary, we set

$$j_{\text{cool}}^{\text{cold-mode}} = j_{\text{halo}}, \quad (17)$$

maintaining the status quo from several SAMs (e.g. Lagos *et al.* 2018; Henriques *et al.* 2020). j_{halo} is the specific angular momentum of the halo inside the virial radius, which we take from the input merger trees. For the hot mode, we modify the fitting function of Stevens *et al.* (2017a) — which was first implemented in DARK SAGE in Stevens *et al.* (2018) — by rewriting it in terms of j :

$$\log_{10} \left(\frac{j_{\text{cool}}^{\text{hot-mode}}}{2 R_{200c} V_{200c}} \right) = 0.23 \log_{10}(\lambda) - k_d, \quad (18)$$

where we adopt the Bullock *et al.* (2001) approximation for halo spin,

$$\lambda = \frac{j_{\text{halo}}}{\sqrt{2} R_{200c} V_{200c}}. \quad (19)$$

While the slope of equation (18) originates from equation 19 of Stevens *et al.* (2017a), the normalisation is set by hand to recover the result of Stevens *et al.* (2017a) that $(j_{\text{cool}}^{\text{hot-mode}}/j_{\text{halo}}) \simeq 1.4 \simeq \sqrt{2}$ when one only considers gas inside the virial radius. Rearranging the above, one can then solve for k_d to force this outcome. Doing so gives

$$k_d \simeq \log_{10}(\lambda^{-0.77}). \quad (20)$$

To match the low-redshift sample used in Stevens *et al.* (2017a), we calculate the average of $\lambda^{-0.77}$ for haloes with $10^{11.5} \leq M_{200c}/M_\odot \leq 10^{12.5}$ at $z = 0$ in Mini-MTNG. This returns $k_d = 1.15$.^f Because of the explicit dependence on halo spin in equation (18) — which is known to have a scatter of ~ 0.26 dex and only systematically depend on halo mass to second order in N -body simulations (Bullock *et al.* 2001; Knebe & Power 2008) — we naturally have a scatter in $j_{\text{cool}}^{\text{hot-mode}}/j_{\text{halo}}$ in DARK SAGE of ~ 0.2 dex. This is shown as a function of halo mass in Figure 1. As can be seen by comparing the curves and shaded regions in the top panel for all centrals with those exclusively in the hot mode, $M_{200c} \simeq 10^{12} M_\odot$ marks the transition between haloes predominantly in the cold mode and those in the hot mode. There are, nevertheless, some haloes in the hot mode all the way down to the resolution limit of the simulation.

^fIn Stevens *et al.* (2018), k_d was set to 1.0. This was a calculation error.




Figure 1. Top panel: Host halo mass versus the ratio of the specific angular momentum of gas cooling from the CGM onto the ISM of the central galaxy in DARK SAGE to that of the halo from MillenniumTNG at $z = 0$. For haloes accreting in the ‘hot mode’, this ratio is determined by equation (18). The scatter and mild slope in this relation is driven entirely by the distribution of halo spins, as seen in the bottom panel. For cold-mode haloes, it is assumed that $j_{\text{cool}}/j_{\text{halo}} = 1$. Running percentiles and means are shown for the full population (thinner lines in the foreground) and when exclusively selecting those in the hot mode (thicker, behind). The dotted line in the top panel marks where $j_{\text{cool}}/j_{\text{halo}} = 1.4$, the average value predicted for hot-mode haloes by Stevens et al. (2017a). Satellite subhaloes are excluded from this plot.

With the above pieces in place, we can calculate the amount of cooling gas that goes into each annulus of the new gas disc. For the i th annulus, the mass received is

$$\Delta m_{\text{hot} \rightarrow \text{cold}}^{(i)} = \Delta m_{\text{hot} \rightarrow \text{cold}} \int_{j_{\text{inner}}^{(i)}}^{j_{\text{outer}}^{(i)}} \text{PDF}_{\text{cool}}(j) \, dj. \quad (21)$$

If this is the first cooling episode in a halo or m_{cold} was 0 at the start of the sub-time-step, this completes the relevant description of how a gas disc is formed. If this new gas is adding to a pre-existing disc in the galaxy (as is the case for most galaxies at most snapshots), there is one more step to consider.

While the *magnitudes* of the cooling gas and halo’s specific angular momenta are not always assumed to be the same, their *directions* are. The orientation of a galaxy’s disc in DARK SAGE is always tracked by its angular momentum vector. When a new gas disc forms, the direction of its angular momentum vector is set parallel to that of the halo, that is, $\hat{j}_{\text{cold}} = \hat{j}_{\text{cool}} = \hat{j}_{\text{halo}}$. However, \hat{j}_{halo} changes over time for any given halo, meaning sequential accretion episodes are unlikely to be parallel. When a pre-existing gas disc is present, a vector sum of \vec{j}_{cold} and \vec{j}_{cool} is done to define the new orientation of the resulting gas disc. Both the pre-existing and freshly cooled gas discs are then projected onto this new plane, which uses the same discretised j structure per equation (3a). This means that the mass in the first annulus of the resulting disc will be the sum of the first annuli of the pre-existing and freshly cooled discs *plus* some fraction of their second (and possibly their third, fourth, etc.) annuli, depending on the angles involved. This aspect of DARK SAGE remains unchanged from Stevens et al. (2016).

It is worth mentioning that \hat{j}_{halo} is typically a noisy quantity in an N -body simulation (e.g. Benson 2017; Contreras, Padilla & Lagos 2017). While this induces some randomness to the precise orientation and magnitude of the angular momentum of cooling

gas in a DARK SAGE at a particular instant, this randomness should be averaged out after several time-steps with our method.

5. Computing potentials and converting from specific angular momentum to radius

Recall that disc annuli in DARK SAGE galaxies are fixed in terms of specific angular momentum. While this design choice was deliberate, it nevertheless remains crucial to compute the *radii* of annuli for several of the model’s modules. Annulus radii must be (and are) regularly updated in the model, at least once per sub-time-step per galaxy. To convert j to r , one must first compute the potential of the galaxy and its (sub)halo in the plane of the disc, $\Phi(r)$. In Stevens et al. (2016), this was achieved by approximating the potential as spherically symmetric and the disc to have precisely centripetal motion at all radii, then iteratively solving for $r_{\text{outer}}^{(i)}$ for each annulus through

$$j_{\text{outer}}^{(i)} = r_{\text{outer}}^{(i)} v_{\text{circ}}(r_{\text{outer}}^{(i)}). \quad (22)$$

Stevens et al. (2018) later dropped the assumption of perfect centripetal motion by adding a radially variant factor for the fraction of gravity balanced by circular motion (as opposed to random radial motion), $f_{\text{circ}}(r)$:

$$v_{\text{circ}}^2(r) = f_{\text{circ}}(r) v_{\text{cent}}^2(r) = f_{\text{circ}}(r) r \frac{d\Phi}{dr}, \quad (23)$$

where subscripts ‘circ’ and ‘cent’ are respective shorthands for ‘circular’ and ‘centripetal’. As noted in Stevens et al. (2018), the spherical-potential approximation is accurate in the outskirts of a galaxy, but becomes exponentially inaccurate towards the centre. Moreover, solving equation (22) iteratively was a bottleneck in the code. For these reasons we have overhauled the j -to- r conversion.

With the new version of the model, we treat the shape of the potential as a combination of spherically symmetric components (the dark-matter halo, stellar bulges, intracluster stars, and hot gas; shorthand ‘sph’) and two axisymmetric, infinitesimally thin discs (one each for stars and cold gas). This can be written in terms of contributions to the centripetal velocity:

$$v_{\text{cent}}^2(r) = v_{\text{sph}}^2(r) + v_{*,\text{disc}}^2(r) + v_{\text{gas, disc}}^2(r). \quad (24)$$

The challenge comes in the implementation of the two lattermost terms. While one can numerically solve for a disc’s potential as a function of radius from first principles, the only way to self-consistently do this in the framework of DARK SAGE is by maintaining the iterative j -to- r process *and* increasing the computational demand of each iteration. For practical reasons, we therefore employ an approximation of a different kind, but one that remains faithful to the general shape of a disc potential. Namely, we assume the analytic approximation of Obreschkow et al. (2009, their equation 37):

$$v_{\text{disc}}^2(r) = \frac{G m_{\text{disc}}}{r_s} \left[1 + 4.8 \exp \left(-0.35 \frac{r}{r_s} - 3.5 \frac{r_s}{r} \right) \right] \times \left[\frac{r}{r_s} + \left(\frac{r_s}{r} \right)^2 + 2\sqrt{\frac{r_s}{r}} \right]^{-1}. \quad (25)$$

This approximation works well for exponential discs with scale radius r_s . Clearly, it is less accurate for discs with non-parametric disc structure as in DARK SAGE. Nevertheless, after testing, we found this to be the best compromise between accuracy and practicality.

Because there is no requirement for DARK SAGE discs to be precisely exponential, we need to calculate an analogue for r_s to use in equation (25). For exponential discs, there is a one-to-one relation between the scale radius and any radius that encloses x per cent of the disc's mass, r_x . For each of $x = 10, 20, \dots, 80, 90$, we calculate the effective exponential scale radius that corresponds to the disc's r_x . We then set r_s to be the average of these nine scale radii. This is a small shift from Stevens *et al.* (2018), who only used r_{50} and r_{90} to inform r_s , which relied on an otherwise unnecessary and now non-existent free parameter. The calculation of r_s and subsequent application of equation (25) are *separately* done for *each* of the stellar and gas discs, fulfilling the right-hand side of equation (24).

The contribution to the potential from all other baryons and dark matter is then straightforwardly calculated as:

$$v_{\text{sph}}^2(r) = \frac{G m_{\text{sph}}(<r)}{r}, \quad (26a)$$

$$m_{\text{sph}}(<r) = 4\pi \int_0^r \left(\rho_{\text{DM}}(R) + \rho_{\text{CGM}}(R) + \sum_{\otimes} \rho_{\otimes}(R) \right) R^2 dR, \quad (26b)$$

where $\otimes = \text{i-bulge, m-bulge, IHS}$. Each of the profiles that contributes to $M_{\text{sph}}(<r)$ is analytic. Per Stevens *et al.* (2016, appendix B), dark matter is assumed to follow an NFW profile (Navarro, Frenk & White 1996) with a baryon-influenced mass-concentration relation (Di Cintio *et al.* 2014; Dutton & Macciò 2014):

$$\rho_{\text{DM}}(R) = \frac{M'_{\text{DM}} c_{\text{NFW}}^2(z)}{4\pi R_{200c}^2 R} \left[1 + \frac{c_{\text{NFW}}(z) R}{R_{200c}} \right]^{-2} \times \left[\ln(1 + c_{\text{NFW}}(z)) - \frac{c_{\text{NFW}}(z)}{1 + c_{\text{NFW}}(z)} \right]^{-1}, \quad (27a)$$

$$c_{\text{NFW}}(z) = \left(1 + 3 \times 10^{-5} \exp \left\{ 3.4 \left[\log_{10} \left(\frac{m_*}{M_{200c}} \right) + 4.5 \right] \right\} \right) c_{\text{DMO}}(z), \quad (27b)$$

$$c_{\text{DMO}}(z) = 10^{a_{\text{NFW}}(z)} + \left(\frac{M_{200c} h}{10^{12} M_{\odot}} \right)^{b_{\text{NFW}}(z)}, \quad (27c)$$

$$a_{\text{NFW}}(z) = 0.520 + 0.385 \exp(-0.617 z^{1.21}), \quad (27d)$$

$$b_{\text{NFW}}(z) = -1.01 + 0.026 z. \quad (27e)$$

In principle, the dark-matter mass of the halo should be readily calculated as $M_{\text{DM}} = (1 - f_{\text{bary}}^{200c}) M_{200c}$ (equation 5). But, again for the sake of computation efficiency, we approximate the matter bound in subhaloes to also be distributed according to the same NFW profile as the dark matter (despite knowing that this is not the most accurate description of how satellites are distributed in DARK SAGE — see Qin *et al.* 2022). We therefore effect

$$M'_{\text{DM}} \equiv M_{200c} - (m_* + m_{\text{cold}} + m_{\text{CGM}} + m_{\text{IHS}} + m_{\text{BH}} + m_{\text{IHBH}}). \quad (28)$$

Stellar spheroids (the instability-driven bulge, merger-driven bulge, and intrahalo stars) all follow Hernquist (1990) profiles that

are truncated at the virial radius:

$$\rho_{\otimes}(R) = \frac{m_{\otimes} a_{\otimes} (R_{200c} + a_{\otimes})^2}{2\pi R (R + a_{\otimes})^3 R_{200c}^2}. \quad (29)$$

While we have not changed the functional form of the stellar-spheroid profiles, we now *calculate* each of their scale radii, a_{\otimes} , rather than assuming a scaling relation as done previously. These are discussed in the appropriate sections for each component later in the paper (see equations 71 & 121).

For the final piece of equation (23), we assume a function of the form

$$f_{\text{circ}}(r) = 1 - \exp(-r/r_{\lambda_*}). \quad (30)$$

In Stevens *et al.* (2018), the scale radius r_{λ_*} took the form $r_{\lambda_*} = r_{s,*} / \{3 [1 - m_{*,\text{bulge}} / (m_* + m_{\text{cold}})]\}$,⁸ based loosely on the observational results of Bellstedt *et al.* (2017). Instead, we now calculate r_{λ_*} in a more self-consistent way, but maintain the assumption from Stevens *et al.* (2018) that f_{circ} should scale with the local stellar spin parameter, which is defined for an annulus as

$$\lambda_*^{(i)} \equiv \frac{\langle v_{\text{circ}}^{(i)} \rangle}{\sqrt{\langle v_{\text{circ}}^{(i)} \rangle^2 + \sigma_*^{(i)2}}}. \quad (31)$$

First, we initialise r_{λ_*} as $r_{s,*}/3$. Then, for every subsequent sub-time-step, we update r_{λ_*} to

$$r_{\lambda_*} = \frac{-r\{\lambda_* = 0.5\}}{\ln(0.5)}, \quad (32)$$

where $r\{\lambda_* = 0.5\}$ is smallest radius where λ_* rises to 0.5 (interpolated from the annular $\lambda_*^{(i)}$ values). The factor of $-1/\ln(0.5) \simeq 1.44$ translates the factor-of-two scale radius into an exponential scale radius appropriate for equation (30).

With all of the above, we can analytically calculate $v_{\text{circ}}(r)$ — and therefore j — for any r . To go in the opposite direction, that is, from j to r for annuli, we calculate $j(r)$ for an array of r values. Using this array, we interpolate the j boundaries of the disc annuli to find their corresponding r .

We approximate the surface density of a given disc component to be constant *within* an annulus, that is,

$$\Sigma_{\text{disc}}^{(i)} \equiv \frac{m_{\text{disc}}^{(i)}}{\pi (r_{\text{outer}}^{(i)2} - r_{\text{inner}}^{(i)2})}. \quad (33)$$

With this, it follows that the mass in an annulus has an *average* radius of

$$\langle r^{(i)} \rangle \equiv \sqrt{\frac{r_{\text{inner}}^{(i)2} + r_{\text{outer}}^{(i)2}}{2}}. \quad (34)$$

To be consistent with equation (4), we define the average circular velocity of an annulus as

$$\langle v_{\text{circ}}^{(i)} \rangle \equiv \frac{\langle j^{(i)} \rangle}{\langle r^{(i)} \rangle}. \quad (35)$$

Finally, for simplicity and convenience, we define an approximate average potential of each annulus as

$$\langle \Phi^{(i)} \rangle \equiv \frac{\Phi(r_{\text{inner}}^{(i)}) + \Phi(r_{\text{outer}}^{(i)})}{2}. \quad (36)$$

⁸Note the correction from equation (A5) of Stevens *et al.* (2018), where m_{cold} was accidentally missed from this expression.




Figure 2. Circular velocity profiles (top) and stellar velocity dispersion profiles (bottom) of DARK SAGE discs at $z=0$. We sample one random central galaxy with $\text{sSFR} = 10^{-10.5} \text{ yr}^{-1}$ and $m_{*,\text{bulge}} < 0.3 m_*$ for each stellar mass in intervals of 0.1 dex between $10^{8.5}$ and $10^{11} M_{\odot}$ with a 0.04-dex search window. The colour of the curves transitions from dark blue to bright yellow as we move up in mass.

To demonstrate the resolved kinematics of galaxy discs in DARK SAGE, we plot circular velocity profiles and stellar velocity dispersion profiles for random example galaxies across a range of stellar masses in Figure 2. A description of how $\sigma_{*,\text{disc}}$ builds up over time can be found in Section 7.4.

One relatively fundamental scaling relation of galaxies is the baryonic Tully–Fisher relation (McGaugh et al. 2000), a tighter relation than in its original stellar form of Tully & Fisher (1977). This relates the peak rotational velocity of a galaxy (a proxy for its potential) with its baryonic mass (in principle, another proxy for potential, if baryonic mass and dark-matter mass scale) for rotation-supported (disc-dominated) systems. Given that we have built the velocity curves of DARK SAGE galaxies from their potential and the stellar and H I masses of galaxies are constrained by their observed mass functions at $z=0$ (Section 13), the baryonic Tully–Fisher relation should fall out of DARK SAGE as a natural prediction. As we show in Figure 3 by comparing to data from the SPARC^h sample of galaxies (Lelli et al. 2019), it does! Though, we caution that our sample selection to match the SPARC is crude, intended only as a proof of concept here. The galaxies masses of the SPARC sample also have increasingly larger systematic uncertainties due to the adoption of a morphology-independent mass-to-light ratio (see Lelli et al. 2016, 2019). For these reasons, we should not expect a perfect match. For example, the scatter at low masses for DARK SAGE of ~ 0.02 dex appears low compared to the intrinsic scatter of 0.040 ± 0.006 dex in SPARC (cf. Lelli et al. 2019), but we found in testing other cuts that the scatter in DARK SAGE varied up to 0.03 dex.

^hSpitzer Photometry and Accurate Rotation Curves




Figure 3. The Baryonic Tully–Fisher relation for DARK SAGE galaxies: that is, the maximum rotational velocity of each galaxy as a function of its stellar + neutral-gas mass. Hex bins show the number density of DARK SAGE galaxies with $m_{*,\text{bulge}} < 0.1 m_*$, $m_{\text{H I}} \geq 10^8 M_{\odot}$, and $m_{\text{neutral}}/m_* \geq 10^{-2}$. The running solid lines are the median (thick) and 16th and 84th percentiles (thin) of the DARK SAGE selection, calculated in bins along the y-axis. The DARK SAGE cuts are a crude attempt at a comparable selection to the SPARC sample of galaxies (Lelli et al. 2019), compared here.

6. Gas chemistry and vertical structure of discs

6.1. Velocity dispersion and disc height

Previously in DARK SAGE, it was assumed that the velocity dispersion of gas in galaxy discs, σ_{cold} , was a universal constant for all cosmic time. We now introduce a redshift dependence to σ_{cold} , but maintain the (false but simple) assumption that for a given snapshot this is constant across all annuli of all galaxies. The function we employ for this is empirical, based on the average redshift evolution for observations of neutral gas shown in figure 6 of Übler et al. (2019):

$$\frac{\sigma_{\text{cold}}(z)}{\text{km s}^{-1}} = 11.0 + 11.3 z. \quad (37)$$

Our redshift-zero value of 11 km s^{-1} is what had been previously assumed at all redshifts. The redshift coefficient is taken from equation 1 of Übler et al. (2019). Results from Pillepich et al. (2019) suggest it might be sensible to add a stellar-mass dependence to equation (37), but we leave such modifications for future model developments.

We assume that the vertical structure of discs follows a square-hypersecant profile, typical for those that are self-gravitating and isothermal (see Benítez-Llambay et al. 2018). For a given annulus, we invoke

$$\rho_{\text{cold}}^{(i)}(\zeta) = \frac{\Sigma_{\text{cold}}^{(i)}}{2 \zeta_s^{(i)}} \text{sech}^2\left(\frac{\zeta}{\zeta_s^{(i)}}\right), \quad (38a)$$

$$\zeta_s^{(i)} = \frac{\sigma_{\text{cold}}^2(z)}{G \Sigma_{\text{cold}}^{(i)}}. \quad (38b)$$

Here, ζ represents height from the mid-plane of the disc, where $\zeta_s^{(i)}$ is the local scale height. By definition, the surface density of an annulus must be $\Sigma_{\text{cold}}^{(i)} = \int_{-\infty}^{\infty} \rho_{\text{cold}}^{(i)}(\zeta) d\zeta$. The average three-dimensional density of an annulus is therefore

$$\langle \rho_{\text{cold}}^{(i)} \rangle = \frac{1}{\Sigma_{\text{cold}}^{(i)}} \int_{-\infty}^{\infty} \rho_{\text{cold}}^{(i)2}(\zeta) d\zeta = \frac{G}{3} \left(\frac{\Sigma_{\text{cold}}^{(i)}}{\sigma_{\text{cold}}(z)} \right)^2. \quad (39)$$

We do not differentiate between the dispersions (or scale heights) of ionised and neutral gas in DARK SAGE gas discs, despite the theoretical expectation and empirical evidence that suggest these should differ (Übler *et al.* 2019). Relatively speaking, this is not a great source of uncertainty in the model; the observed systematic difference between the neutral and ionised gas velocity dispersions of galaxies is much less than, for example, the observed scatter in either dispersion at fixed redshift, which we do not model.

When stars are formed (see below), they inherit the velocity dispersion of the gas that formed them. We record a stellar velocity dispersion for each stellar-age bin within each annulus. By construction, this means that old stars carry a higher velocity dispersion than young stars in galaxy discs. As the stellar-age bins are typically wider than the average snapshot interval in the merger trees, each episode of star formation requires updating the velocity dispersion of that age bin. This is done by taking a mass-weighted average of the squares of the velocity dispersions of the new and pre-existing stars.

6.2. Hydrogen fraction

In another update to the model, we adjust the mass fraction of hydrogen, X , in a given gas disc annulus. Naturally, where metallicities are high, one expects to find a lower hydrogen fraction. We adopt the fitting function of Stevens *et al.* (2021, their equation B2), which accurately recovers the hydrogen fraction of gas cells in IllustrisTNG from metallicity alone (i.e. without needing prior knowledge of the helium mass fraction):

$$X = \begin{cases} 0.753 - 1.26 Z, & Z \geq 0.025 \\ \min [0.76, 0.762 - 2.3 Z + 24.2 Z^2], & Z < 0.025 \end{cases}. \quad (40)$$

The helium mass fraction is then straightforwardly recovered as

$$Y = 1 - X - Z. \quad (41)$$

We impose a metallicity floor of $Z = 10^{-10}$ on all baryon reservoirs, which accounts for the trace amounts of lithium formed shortly after the Big Bang. Due to the ‘primordial lithium problem’ — see Fields (2011) — this is lower than what one would expect the initial metallicity of the Universe to be based on Big Bang nucleosynthesis alone. Metallicity is then built up from stellar enrichment (see below). In practice, the precise value of the metallicity floor is irrelevant, provided it is negligible relative to the yield of metals produced by a stellar population, which is of order 10^{-2} .

6.3. Neutral fraction

We introduce a new method to DARK SAGE for calculating the fraction of the ISM that is neutral/ionised. Previously, this had been assumed to be a constant for all annuli and all galaxies. We now instead invoke a model of photoionisation equilibrium within each annulus, providing a predictive framework for the neutral

fraction radial profiles of galaxies. This is done in two stages, considering the situations when (i) local sources of ionising radiation are dominant and (ii) the universal ionising background is instead dominant.

In what follows, we seek a solution for

$$f_{\text{neutral}}^{(i)} \equiv \frac{m_{\text{neutral}}^{(i)}}{m_{\text{cold}}^{(i)}} = \frac{m_{\text{H I}}^{(i)} + m_{\text{H}_2}^{(i)}}{X^{(i)} m_{\text{cold}}^{(i)}}. \quad (42)$$

In practice, we calculate this as

$$f_{\text{neutral}}^{(i)} = \min [f_{\text{neutral}}^{*(i)}, f_{\text{neutral}}^{\text{bg}(i)}]. \quad (43)$$

6.3.1. Local photoionisation

Let us start with the approximation demonstrated by Rahmati *et al.* (2013) that the local photoionisation rate of hydrogen in a galaxy scales linearly with its star formation rate. Writing this in terms of ionisation and star formation rate densities, we have

$$\dot{n}_{\text{ion}} \simeq q_{\gamma} \rho_{\text{SFR}}, \quad (44)$$

where $q_{\gamma} = 2 \times 10^{53} \text{ s}^{-1} (\text{M}_{\odot} \text{ yr}^{-1})^{-1} = 6.31 \times 10^{60} \text{ M}_{\odot}^{-1}$ (Rahmati *et al.* 2013). The question then becomes: how do we implement this on an annulus-by-annulus basis?

Similar to equation (39), we can calculate the average three-dimensional SFR density in terms of an annulus’s SFR surface density, given $\rho_{\text{SFR}}^{(i)}(\zeta) \propto \rho_{\text{cold}}^{(i)}(\zeta)$:

$$\langle \rho_{\text{SFR}}^{(i)} \rangle = \frac{1}{\Sigma_{\text{cold}}^{(i)}} \int_{-\infty}^{\infty} \rho_{\text{cold}}^{(i)}(\zeta) \rho_{\text{SFR}}^{(i)}(\zeta) d\zeta. \quad (45)$$

Approximating ρ_{SFR} in equation (44) as $\langle \rho_{\text{SFR}}^{(i)} \rangle$ above and solving the integral, we take

$$\dot{n}_{\text{ion}}^{(i)} = \frac{q_{\gamma} G \Sigma_{\text{cold}}^{(i)} \Sigma_{\text{SFR}}^{(i)}}{3 \sigma_{\text{cold}}^2(z)}. \quad (46)$$

To solve for $f_{\text{neutral}}^{(i)}$, we need to balance ionisation with the recombination of atoms. The recombination rate density of hydrogen can be written as

$$\dot{n}_{\text{recom}} = n_e n_p \alpha, \quad (47)$$

where n_e and n_p are the respective number densities of free electrons and free protons, and $\alpha = 4 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$ is the recombination coefficient for hydrogen that we adopt. For an annulus, we consider *average* three-dimensional number densities, consistent with our earlier assumptions that there is no vertical variation in the neutral fraction (or ionisation–recombination balance):

$$\dot{n}_{\text{recom}}^{(i)} = \langle n_e^{(i)} \rangle \langle n_p^{(i)} \rangle \alpha, \quad (48a)$$

$$\langle n_e^{(i)} \rangle = \frac{4}{3 X^{(i)} + 1} \langle n_p^{(i)} \rangle, \quad (48b)$$

$$\langle n_p^{(i)} \rangle = \left(1 - f_{\text{neutral}}^{(i)} \right) X^{(i)} \frac{\langle \rho_{\text{cold}}^{(i)} \rangle}{m_p} \quad (48c)$$

$$= \left(1 - f_{\text{neutral}}^{(i)} \right) X^{(i)} \frac{G}{3 m_p} \left(\frac{\Sigma_{\text{cold}}^{(i)}}{\sigma_{\text{cold}}(z)} \right)^2. \quad (48d)$$

The difference between $\langle n_p^{(i)} \rangle$ and $\langle n_e^{(i)} \rangle$ is driven by extra free electrons sourced from ionised helium. For simplicity, we approximate that for every hydrogen atom that is split into a free electron

and proton, there is a helium atom split into a free electron and He II ion. In reality, it takes more energy to ionise a helium atom than a hydrogen atom, and the cross section of the two elements are not identical either, but these minutiae are unimportant for the detail modelled here. We also reasonably neglect any contributed free electrons from He III or ions of other elements, given their low cosmic abundances.

Under photoionisation equilibrium,

$$\dot{n}_{\text{ion}}^{(i)} = \dot{n}_{\text{recom}}^{(i)}. \quad (49)$$

Equation (48) serves as the right-hand side of equation (49). If we take equation (46) as the left-hand side, we are implicitly treating $f_{\text{neutral}}^{(i)}$ as $f_{\text{neutral}}^{*(i)}$. After combining these equations, we can rearrange and solve for the neutral fraction (when local photoionisation dominates) in terms of knowns:

$$f_{\text{neutral}}^{*(i)} = 1 - \frac{1}{X^{(i)}} \sqrt{\frac{3(3X^{(i)} + 1)q_{\gamma}m_{\text{p}}^2\sigma_{\text{cold}}^2(z)\Sigma_{\text{SFR}}^{(i)}}{4\alpha G\Sigma_{\text{cold}}^{(i)3}}}. \quad (50)$$

For convenience, we note that the constants in this equation can be reduced to

$$\frac{3q_{\gamma}m_{\text{p}}^2}{4\alpha G} = 0.2495 \text{ cm}^{-6} \text{ g}^2 \text{ s}^3. \quad (51)$$

One final decision remains: how do we quantify $\Sigma_{\text{SFR}}^{(i)}$ for equation (50)? While we could tie this to the sub-time-step's instantaneous SFR in the annulus, this would require a simultaneous solution for $\Sigma_{\text{SFR}}^{(i)}$ and $f_{\text{neutral}}^{(i)}$, because, as we will show in Section 7.3, local star formation explicitly depends on local molecular fraction, which requires knowing the neutral fraction first. But a better option exists. What equation (44) really symbolises is that *young* stars produce the majority of local photoionising radiation in galaxies. The age bins we have in each DARK SAGE disc annulus effectively tell us how many young stars are present. That is, to be self-consistent with our treatment of stellar evolution (where any stars formed in the same age bin are treated as if they have an identical formation time — see the next section), we use the annulus's average SFR for the age bin that the present sub-time-step falls in to calculate $\Sigma_{\text{SFR}}^{(i)}$ for equation (50). This is both physically motivated and circumvents a circular solution for $f_{\text{neutral}}^{*(i)}$.

We note that AGN should, in principle, add an additional source of local photoionisation. This is not something we have considered for this version of DARK SAGE. We leave it for future developments of the model.

6.3.2. Photoionisation background

We adopt a redshift-dependent, universal photoionising radiation background for neutral hydrogen to place a ceiling on the neutral fraction of the ISM of DARK SAGE galaxies. When this source of photoionisation dominates over local sources, we impose

$$\dot{n}_{\text{ion}} = \Gamma_{\text{HI}}^{\text{eff}}(z) n_{\text{H,neutral}}. \quad (52)$$

Here, the effective rate of hydrogen photoionisation, $\Gamma_{\text{HI}}^{\text{eff}}(z)$, is taken from Table D1 of Faucher-Giguère (2020). In practice, for DARK SAGE disc annuli, we effect

$$\dot{n}_{\text{ion}}^{(i)} = \Gamma_{\text{HI}}^{\text{eff}}(z) f_{\text{neutral}}^{\text{bg}(i)} X^{(i)} m_{\text{p}}^{-1} \langle \rho_{\text{cold}}^{(i)} \rangle. \quad (53)$$

Combining this with equations (48 & 49), after some omitted algebra, we end up with a quadratic equation for $f_{\text{neutral}}^{\text{bg}(i)}$ in terms of

knowns. The requirement that $f_{\text{neutral}}^{\text{bg}(i)} \leq 1$ by definition eliminates one of the quadratic solutions, leaving

$$f_{\text{neutral}}^{\text{bg}(i)} = 1 + \gamma - \sqrt{\gamma^2 + 2\gamma}, \quad (54a)$$

$$\gamma \equiv \frac{3(3X^{(i)} + 1)m_{\text{p}}\Gamma_{\text{HI}}^{\text{eff}}(z)}{8\alpha GX^{(i)}} \left(\frac{\sigma_{\text{cold}}(z)}{\Sigma_{\text{cold}}^{(i)}} \right)^2. \quad (54b)$$

6.4. Molecular fraction

We return to using a formula based on the mid-plane pressure of each disc annulus, $P_{\text{mid}}^{(i)}$, as the default for calculating the molecular fraction of gas. Similar to the implementation in Stevens et al. (2016), we use the equations from Blitz & Rosolowsky (2004) and Elmegreen (1989):

$$f_{\text{H}_2}^{(i)} = \left[\left(\frac{P_{\text{mid}}^{(i)}}{5.93 \times 10^{-13} \text{ Pa}} \right)^{-0.92} + 1 \right]^{-1}, \quad (55a)$$

$$P_{\text{mid}}^{(i)} = \begin{cases} \frac{\pi G \Sigma_{\text{cold}}^{(i)}}{2} \left(\Sigma_{\text{cold}}^{(i)} + \frac{\sigma_{\text{cold}}(z) \Sigma_{*}^{(i)}}{\sigma_{*}^{(i)}} \right), & \theta_{\text{disc}} \leq 10^\circ \\ \frac{\pi G}{2} \Sigma_{\text{cold}}^{(i)2}, & \theta_{\text{disc}} > 10^\circ \end{cases}. \quad (55b)$$

The only practical changeⁱ from Stevens et al. (2016) is how σ_{cold} and $\sigma_{*}^{(i)}$ are calculated, which is now consistent with that described above in this paper. θ_{disc} is the angular misalignment between the stellar and gas discs.

Other prescriptions for $f_{\text{H}_2}^{(i)}$ are available in the code, such as those based on the works of McKee & Krumholz (2010) and Gnedin & Draine (2014). We do not adopt them in this paper, though, as we found in testing that extra free parameters were likely required to calibrate the model to observations with these prescriptions.

In the top two panels of Figure 4, we show the H I and H₂ fractions of DARK SAGE galaxies as a function of stellar mass. These relations are predictions of the model, and are compared against the representative xGASS and xCOLD GASS samples of galaxies (Saintonge et al. 2017; Catinella et al. 2018). We note that this raw comparison is simply for a rough reference; to make a fair, accurate comparison between and any simulations and this dataset requires deeper considerations, as has been explored in detail in Stevens et al. (2019a, 2021).

7. Disc instabilities and star formation

The one physical process that arguably is the linch-pin of the entire DARK SAGE model is gravitational disc instabilities. The handling of said instabilities drives the majority of star formation, radial redistribution of disc mass, disc heating, bulge growth, and BH growth. While the triggering of an instability in the new DARK

ⁱIn truth, there is one other, minor change. Namely, we have removed the rogue factor of h^2 that appeared in the normalising pressure in Stevens et al. (2016), that is, the equivalent of the denominator in the innermost brackets in equation (55a) here. This was an old error in the code, which was discovered during the proofing stage of that paper. The text for that paper was modified at the last second before publishing to be consistent with the code. This was largely inconsequential. The code has since been fixed.




Figure 4. The planes of stellar mass versus H I fraction (top panel), H₂ fraction (middle panel), and specific star formation rate (bottom panel) in DARK SAGE galaxies at $z = 0$. Representative observational samples with $m_* \geq 10^9 M_\odot$ are compared from xGASS for H I (Catinella *et al.* 2018), xCOLD GASS for H₂ (Saintonge *et al.* 2017), and SDSS for sSFR (the volume-limited sample used in Brown *et al.* 2017). Thick lines represent medians. Thin lines are 16th and 84th percentiles. For xGASS and xCOLD GASS, the lines assume non-detections carry their upper-limit value, while the shaded region covers the 16th–84th interpercentile range assuming non-detections have zero mass of that type. SDSS contours approximately encapsulate 38, 68, and 95 per cent of the sample (thicker to thinner). DARK SAGE hexbins and solid lines in the bottom panel use the time-step-averaged SFR (almost instantaneous), while the dashed lines represent the average SFR in the last age bin (~ 1 Gyr).

SAGE is very similar to that of Stevens *et al.* (2016), some important changes have been made to how instabilities are resolved.

Previously in DARK SAGE, stars could form through one of three possible channels. The ‘passive’ channel prescribed each annulus to form stars at a rate proportional to its local H₂ content. The ‘instability’ channel represented stars formed in bursts driven by local gravitational instabilities. The ‘merger’ channel accounted for extra star formation triggered by gas compression during a galaxy merger. While the former two were physically well defined to operate in a single annulus, the prescription for merger-driven starbursts was an *ad hoc*, purely phenomenological formula, designed originally to work on global galaxy scales instead (Somerville, Primack & Faber 2001): see section 3.9.3 of Stevens *et al.* (2016). Because this goes against the ethos of the new model, we have removed the merger channel for star formation. We maintain the passive and instability channels. Elimination of the merger-burst channel does not prevent elevated star formation from taking place. With how mergers are resolved (Section 12), the sudden addition of gas to the descendant’s ISM lowers its stability, making it likely that additional star formation occurs through instabilities directly after the merger.

For more than the above reason, instabilities play a more central role in DARK SAGE star formation than they used to (despite their already having been the most important channel for many galaxies, as shown in figure 1 of Stevens & Brown 2017). While passive star formation still occurs, we make an important change in calculating this *after* resolving instabilities, rather than before. This means that passive star formation is now exclusive to the gas that is stable on the scale of an annulus but still has molecular clouds present. In hindsight, DARK SAGE should have been originally set up this way.

7.1. Calculating (in)stability

Each annulus of each disc of each galaxy is routinely checked for its gravitational stability according to Toomre’s Q criterion. We calculate this local stability parameter first for each of the stellar and gas components separately, following the respective equations

$$Q_*^{(i)} = \frac{\kappa^{(i)} \sigma_{*,\text{disc}}^{(i)}}{3.36 G \Sigma_*^{(i)}}, \quad (56a)$$

$$Q_{\text{cold}}^{(i)} = \frac{\kappa^{(i)} \sigma_{\text{cold}}(z)}{\pi G \Sigma_{\text{cold}}^{(i)}}, \quad (56b)$$

$$\kappa^{(i)} \equiv \sqrt{\frac{2 \langle v_{\text{circ}}^{(i)} \rangle (j_{\text{outer}}^{(i)} - j_{\text{inner}}^{(i)})}{\langle r^{(i)} \rangle^2 (r_{\text{outer}}^{(i)} - r_{\text{inner}}^{(i)})}} \quad (56c)$$

(Toomre 1964; Binney & Tremaine 1987; Pringle & King 2007). Provided $\theta_{\text{disc}} \leq 10^\circ$, $Q_*^{(i)}$ and $Q_{\text{cold}}^{(i)}$ are then combined into a total stability, $Q^{(i)}$, through

$$Q^{(i)} = \left(\frac{1}{\min [Q_{\text{cold}}^{(i)}, Q_*^{(i)}]} + \frac{W^{(i)}}{\max [Q_{\text{cold}}^{(i)}, Q_*^{(i)}]} \right)^{-1}, \quad (57a)$$

$$W^{(i)} \equiv \frac{2 \sigma_*^{(i)} \sigma_{\text{cold}}(z)}{\sigma_*^{(i)2} + \sigma_{\text{cold}}^2(z)} \quad (57b)$$

(Romeo & Wiegert 2011). If $Q^{(i)} < 1$, the annulus is gravitationally unstable.

By definition, it is not physically possible for something to remain in an unstable state for an extended period of time. This demands that something must happen to bring the annulus back to $Q^{(i)} = 1$. We therefore seek to resolve instabilities as soon as they occur.

Physically, one would expect instabilities to produce star formation from the collapse of unstable gas. Instabilities can also cause a rearrangement of mass in a galaxy. In principle, that rearrangement of mass can break the symmetry of a galaxy’s gravitational potential, for example, through creating a bar, which can create a radially variant torque on the disc, thereby redistributing angular momentum (see, e.g. Athanassoula 2003). This would, in turn, establish radial migration in the galaxy’s disc and alter the orbits of stars, which could also affect the stellar velocity dispersion (and disc scale height) of the disc and create a pseudobulge (see Kormendy & Kennicutt 2004, and references therein). In DARK SAGE, we do not explicitly model bars nor asymmetric gravitational potentials, while pseudobulges are inherently considered to be part of the disc (for relevant discussion, see Stevens et al. 2016). We nevertheless *do* model the end result that instabilities have on discs, namely in the form of star formation, radial migration, and disc heating. We do this in a way that is agnostic to the details of *why* each aspect of instability resolution happens the way it does, as it ultimately does not matter for how the model is constructed.

There are two ways to return an unstable disc annulus to stability: either mass must be removed from the annulus or its velocity dispersion must be raised. We have already prescribed σ_{cold} to be a strict function of z , meaning the former must be how gaseous instabilities are resolved. Cold gas can be removed from an annulus in one of two ways: by transferring it to adjacent annuli, or by induced star formation triggering stellar feedback. Stellar mass can only be removed by transferring it to adjacent annuli. However, we allow the velocity dispersion of stars to also rise. Each phase therefore has two modes of instability resolution. The fractional importance of each mode is controlled by a single free parameter, dubbed $f_{\text{move}}^{\text{gas}}$.

The first step to resolving an instability is deciding which baryon phase the onus is on to solve it. A return to stability ($Q^{(i)} \geq 1$) can be guaranteed once *both* $Q_{\text{cold}}^{(i)} \geq 1 + W^{(i)}$ and $Q_{*}^{(i)} \geq 1 + W^{(i)}$. If one phase already satisfies this requirement, the onus is entirely on the other phase to resolve the instability. If neither criterion is met, both phases have to do some work.

If $\theta_{\text{disc}} > 10^\circ$, the stability of the gas and stellar discs are treated entirely independently. Respective instabilities are then triggered when $Q_{\text{cold}}^{(i)} < 1$ and $Q_{*}^{(i)} < 1$.

7.2. Resolving gas instabilities

Once a gas instability is identified, it is resolved in two steps. Both steps involve removing gas from the annulus. The amount of unstable gas that must be removed from that annulus to return it to stability is

$$m_{\text{cold,unstable}}^{(i)} = m_{\text{cold}}^{(i)} \left(1 - \frac{Q_{\text{cold}}^{(i)}}{Q_{\text{cold,min}}^{(i)}} \right), \quad (58)$$

where $Q_{\text{cold,min}}^{(i)}$ is greater than $Q_{\text{cold}}^{(i)}$ (for a stable disc, $Q_{\text{cold}}^{(i)} = Q_{\text{cold,min}}^{(i)}$ by our definition) and typically equal to 1 or $1 + W^{(i)}$, but can vary based on the circumstance of the instability trigger as

described in Section 7.1. The first step in resolving this instability is to redistribute gas in the disc. This is done by moving a controlled fraction of the unstable mass to the two adjacent annuli, in proportion such that angular momentum is conserved. Mathematically, we can write this as

$$\Delta m_{\text{cold}}^{(i) \rightarrow (i \pm 1)} = \frac{\langle j^{(i)} \rangle - \langle j^{(i \mp 1)} \rangle}{\pm \langle j^{(i+1)} \rangle \mp \langle j^{(i-1)} \rangle} f_{\text{move}}^{\text{gas}} m_{\text{cold,unstable}}^{(i)}. \quad (59)$$

Whenever mass is transferred to adjacent annuli in this way, it carries with it metals in proportion to its metallicity. Strictly, equation (59) only applies for $i \in [2, N_{\text{ann}} - 1]$. For the outermost annulus, mass transfer can only happen inwards, meaning we abandon conservation of angular momentum, and simply adopt

$$\Delta m_{\text{cold}}^{(N_{\text{ann}}) \rightarrow (N_{\text{ann}} - 1)} = f_{\text{move}}^{\text{gas}} m_{\text{cold,unstable}}^{(N_{\text{ann}})} \quad (60)$$

For the innermost annulus, any inward-moving gas is assumed to be accreted by the BH. We approximate the gas accreted onto the BH to have lost all its angular momentum. We still adhere to angular momentum conservation for outward- and inward-moving gas in this case, meaning

$$\Delta m_{\text{cold}}^{(1) \rightarrow (2)} = \frac{\langle j^{(1)} \rangle}{\langle j^{(2)} \rangle} f_{\text{move}}^{\text{gas}} m_{\text{cold,unstable}}^{(1)}, \quad (61a)$$

$$\Delta m_{\text{cold} \rightarrow \text{BH}}^{(1)} = \frac{\langle j^{(2)} \rangle - \langle j^{(1)} \rangle}{\langle j^{(2)} \rangle} f_{\text{move}}^{\text{gas}} m_{\text{cold,unstable}}^{(1)}. \quad (61b)$$

Note that the growth of the BH through this channel is independent of its present mass. This is how DARK SAGE is able to seed zero-mass BHs. We discuss the consequences of this ‘quasar’ form of BH accretion in Section 9.

The movement of unstable gas through the disc to the innermost annulus and subsequently onto the BH is the primary means by which BH accretion occurs. The only other way for BHs to gain a substantial amount of mass is when they merge together. An important test of this framework is whether DARK SAGE can recover the well-known relation between a galaxy’s central BH mass and bulge mass. What makes this particularly exciting is that this is now a prediction of the model; as we show in Section 13, this relation has no explicit bearing on the calibration of DARK SAGE’s free parameters. By design, instabilities and mergers *can* simultaneously grow both the BH and bulge in DARK SAGE (if both gas and stars are involved), though this does not guarantee quantitative agreement with the empirical relation between these two properties.

We show in Figure 5 that DARK SAGE predicts a BH–bulge mass relation that overlaps favourably with observations (cf. Scott, Graham & Schombert 2013). Interestingly, DARK SAGE’s relation is not unimodal; this is seen in the intensity of the hexbins and clarified by the misalignment of equivalent running percentiles when the DARK SAGE data are binned along each of the two axes. No cuts have been applied to DARK SAGE here. A deeper assessment of this relation relative to observations is left for future work.

After moving gas, part of the remaining unstable gas is then consumed in star formation, with the last of it reheated out of the disc by the instantaneous component of the stellar feedback associated with those new stars:

$$\Delta m_{\text{cold} \rightarrow *}^{(i) \text{ i-SF}} + \Delta m_{\text{cold} \rightarrow \text{fount}}^{(i) \text{ i-SN}} + \Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i) \text{ i-SN}} = (1 - f_{\text{move}}^{\text{gas}}) m_{\text{cold,unstable}}^{(i)}. \quad (62)$$




Figure 5. Black hole–bulge mass relation for DARK SAGE galaxies at $z = 0$. The x -axis is the sum of the merger- and instability-driven bulge masses. Because this is not a simple unimodal relation in DARK SAGE, we show running percentiles (thick for median, thin for 16th and 84th) when the model’s data are binned along the x -axis (solid) and y -axis (dotted). Neither median traces the high-number-density ridge of DARK SAGE galaxies. Compared is a compilation of observational data from Scott et al. (2013). While previous versions of DARK SAGE calibrated to these data, this outcome is now a prediction of the model.

The balance between each of the terms on the left-hand side of this equation is described in Section 8. Note that $\Delta m_{\text{cold} \rightarrow *}^{(i)\text{-SF}}$ represents the *net* production of stellar mass at the end of the present sub-time-step, and is therefore less than the annulus’s instability-driven star formation rate multiplied by Δt (this is also explained in Section 8). The superscripts i-SF and i-SN are intended to differentiate instability-driven star formation and supernovae from the passive channel of SF (Section 7.3).

After a gas instability is resolved, we recalculate $Q_*^{(i)}$. This is necessary, as the stellar mass in the annulus will have changed, which inevitably (temporarily) raises not only the probability of a stellar instability triggering but also the mass of unstable stellar mass.

7.3. Passive star formation

After resolving instabilities (including what follows in Section 7.4), we allow for a second mode of star formation that depends explicitly on the local molecular gas surface density. This ‘passive’ mode of star formation in DARK SAGE obeys

$$\Sigma_{\text{p-SFR}}^{(i)} = \epsilon_{\text{H}_2} \Sigma_{\text{H}_2}^{(i)}, \quad (63)$$




Figure 6. The resolved molecular Kennicutt–Schmidt relation for DARK SAGE galaxies at $z = 0$. The dashed DARK SAGE line is the median relation for all annuli. The hexbins are grey-scaled according to the logarithmic number of annuli within them. For a fairer comparison to observations, the hexbins only count annuli with $\Sigma_* > 10 M_\odot \text{pc}^{-2}$ from galaxies with $10^{8.9} < m_*/M_\odot < 10^{11}$ and $0.02 < m_{\text{H}_2}/m_{\text{H}_1} < 1.13$ in (sub)haloes with at least 200 particles at one point in their history. The solid lines represent the median (thick) and 16th and 84th percentiles (thin) for those binned annuli *after* weighting each annulus by its area. This weighting further improves the comparison to observations, which typically use pixels of fixed length $\simeq 1$ kpc. The best-fitting relation from observations (Bigiel et al. 2011) and equivalent scatter is shown by the shaded region. Contours encapsulate 38, 68, and 95 per cent of pixels across the HERACLES and VERTICO surveys that are detected in both axes.

where ϵ_{H_2} is the star formation efficiency of H_2 , equivalent to the inverse of a depletion timescale. In a change from Stevens et al. (2016), ϵ_{H_2} is no longer treated as a free parameter. Instead, this is fixed to $\epsilon_{\text{H}_2} = (2.35 \text{ Gyr})^{-1}$, consistent with the empirical best-fitting relation from Bigiel et al. (2011). Of course, there must be a limit to the amount of star formation in a given annulus set by the total mass of gas present, accounting for the gas that must be removed from the annulus due to feedback from the largest stars that form in that population and effectively go supernova immediately. In practice, we therefore effect

$$\Delta m_{\text{cold} \rightarrow *}^{(i)\text{p-SF}} = \min \left[\pi \left(r_{\text{outer}}^{(i)2} - r_{\text{inner}}^{(i)2} \right) \left(1 - f_{\star}^{\text{return}} \right) \Sigma_{\text{p-SFR}}^{(i)} \Delta t, \right. \\ \left. m_{\text{cold}}^{(i)} - \Delta m_{\text{cold} \rightarrow \text{fount}}^{(i)\text{p-SN}} - \Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i)\text{p-SN}} \right]. \quad (64)$$

Many of the terms here will be defined in Section 8.

Despite equation (63) explicitly tying star formation and H_2 surface density together, DARK SAGE predicts a non-trivial resolved molecular Kennicutt–Schmidt (RMKS) relation, that is, Σ_{H_2} versus Σ_{SFR} . This is predictive because (i) equation (63) only applies to stars formed in stable annuli, (ii) the majority of star formation happens in unstable annuli, and (iii) there is no explicit tying of the molecular fraction to disc stability. In Figure 6, we show the RMKS relation for (a relevant subset of) DARK SAGE galaxies at $z = 0$. This is compared to the relation fitted to observations by Bigiel et al. (2011) as well as data from the

HERACLES and VERTICO surveys (Leroy et al. 2009; Brown et al. 2021, respectively). Note that, despite HERACLES and VERTICO respectively probing field and cluster environments, these two datasets have been demonstrated to have consistent RMKS relations (Jiménez-Donaire et al. 2023).[†]

7.4. Resolving stellar instabilities

Similar to gas, there are two steps in resolving a stellar instability. But these steps are different in detail. The mass of stars deemed necessary to move out of the annulus to restore stability is set by

$$m_{*,\text{unstable}}^{(i)} = m_{*,\text{disc}}^{(i)} f_{\text{move}}^{*(i)} \left(1 - \frac{Q_*^{(i)}}{Q_{*,\text{min}}^{(i)}} \right), \quad (65)$$

$$f_{\text{move}}^{*(i)} = 1 - \frac{\sigma_{\text{cold}}(z)}{\sigma_{*,\text{disc}}^{(i)}} (1 - f_{\text{move}}^{\text{gas}}). \quad (66)$$

By construction, it is always true that $\sigma_{\text{cold}}(z) \leq \sigma_{*,\text{disc}}^{(i)}$ and therefore $f_{\text{move}}^{*(i)} \in (0, 1)$. The unstable mass of stars is moved to adjacent annuli in the same j -conserving fashion as done for gas:

$$\Delta m_{*,\text{disc}}^{(i) \rightarrow (i \pm 1)} = \frac{\langle j^{(i)} \rangle - \langle j^{(i \pm 1)} \rangle}{\pm \langle j^{(i \pm 1)} \rangle \mp \langle j^{(i-1)} \rangle} m_{*,\text{unstable}}^{(i)} \quad (67)$$

when $i \in [2, N_{\text{ann}} - 1]$. For the outermost annulus, we again neglect angular momentum conservation by invoking

$$\Delta m_{*,\text{disc}}^{(N_{\text{ann}}) \rightarrow (N_{\text{ann}} - 1)} = m_{*,\text{unstable}}^{(N_{\text{ann}})}. \quad (68)$$

Inward-moving stars from the innermost annulus are transferred to the instability-driven bulge. This bulge component has no angular momentum by design. As such, we effect

$$\Delta m_{*,\text{disc}}^{(1) \rightarrow (2)} = \frac{\langle j^{(1)} \rangle}{\langle j^{(2)} \rangle} m_{*,\text{unstable}}^{(1)}, \quad (69a)$$

$$\Delta m_{*,\text{disc} \rightarrow \text{i-bulge}}^{(1)} = \frac{\langle j^{(2)} \rangle - \langle j^{(1)} \rangle}{\langle j^{(2)} \rangle} m_{*,\text{unstable}}^{(1)}. \quad (69b)$$

As stars transfer from one annulus to the next (or to the instability-driven bulge), they carry with them the same age–metallicity–velocity dispersion distribution. The metallicity and velocity dispersion squared of each age bin in the receiving stellar component are updated with a mass-weighted sum of the moving stars and that already in the receiving annulus.

By itself, the motion of stellar mass prescribed above does not raise $Q_*^{(i)}$ to its minimum stable value. To complete this, we invoke ‘heating’ of the stellar disc. That is, the stellar velocity dispersion of the disc annulus is raised by a factor equal to

$$f_{\text{move}}^* + (1 - f_{\text{move}}^*) \frac{Q_{*,\text{min}}^{(i)}}{Q_*^{(i)}}.$$

This change applies to each age bin individually. The functional form of equation (66) is something we have put in by hand intentionally so that disc heating does not continue to occur unregulated in a perpetually unstable disc. The increase to the annular stellar velocity dispersion is applied after moving the unstable mass of stars.

[†]These data show that environment does, however, affect other local gas scaling relations (Watts et al. 2023).

It is because stars are collisionless and gas collisional that we have different f_{move} values — that is, different radial-migration rates — for gas and stars. The functional form of equation (66) ensures that stellar velocity dispersions in the disc cannot grow unchecked. The more stellar discs are heated, the less they will be heated at the next instability.

The size of an instability-driven bulge is set by the velocity dispersion of the stars it acquires. This is calculated from the Jeans equation, under the assumptions that the instability-driven bulge is both spherically symmetric and isothermal:

$$\frac{d\Phi}{dR} = - \frac{\sigma_{\text{i-bulge}}^2}{\rho_{\text{i-bulge}}} \frac{d\rho_{\text{i-bulge}}}{dR}. \quad (70)$$

Substituting equation (29) — and its derivative, where $\otimes \rightarrow \text{i-bulge}$ — into (70) and rearranging gives

$$a_{\text{i-bulge}} = 3 \left(\frac{1}{\sigma_{\text{i-bulge}}^2} \frac{d\Phi}{dR} - \frac{1}{R} \right)^{-1} - R. \quad (71)$$

In principle, this expression should be solvable at any R . However, because of the approximation of isothermality (i.e. treating $\sigma_{\text{i-bulge}}$ as radially invariant), equation (71) does not return identical answers at any two radii. For that reason, starting at small radii and working outwards incrementally, we continuously solve for $a_{\text{i-bulge}}$, averaging as we go, until reaching a converged answer.

8. New stellar feedback model

We introduce a brand-new model for stellar feedback in DARK SAGE. Gone are the days of explicitly prescribing a mass loading factor that may or may not be modulated by local surface density. Instead, we now use a purely energy-based argument to calculate the amount of reheated (and ejected) gas from each episode of star formation, in each annulus. Gone also are the assumptions of instantaneous recycling, where we have now implemented a delayed, self-consistent scheme for stellar mass loss, metal enrichment, and associated feedback, which we outline below.

8.1. Energy available from feedback

Let us begin with the assumption that the birth masses of stars, \mathcal{M}_* , for any population are distributed according to a Chabrier (2003) initial mass function (IMF, ϕ), where we impose

$$\phi(\bar{\mathcal{M}}_*) = \begin{cases} 0, & \bar{\mathcal{M}}_* < 0.1 \\ \frac{A_*}{\bar{\mathcal{M}}_*} \exp \left[-\frac{1}{2} \left(\frac{\log_{10}(\frac{\bar{\mathcal{M}}_*}{0.08})}{0.69} \right)^2 \right], & 0.1 \leq \bar{\mathcal{M}}_* < 1 \\ k_* \bar{\mathcal{M}}_*^{-2.3}, & 1 \leq \bar{\mathcal{M}}_* \leq 100 \\ 0, & \bar{\mathcal{M}}_* > 100 \end{cases}, \quad (72a)$$

$$\bar{\mathcal{M}}_* \equiv \mathcal{M}_*/M_{\odot}, \quad (72b)$$

where A_* and k_* are constants. Recognising that (i) $\phi(\bar{\mathcal{M}}_*)$ must be continuous at $\bar{\mathcal{M}}_* = 1$ and (ii)

$$\int_0^{\infty} \bar{\mathcal{M}}_* \phi(\bar{\mathcal{M}}_*) d\bar{\mathcal{M}}_* = \int_{0.1}^{100} \bar{\mathcal{M}}_* \phi(\bar{\mathcal{M}}_*) d\bar{\mathcal{M}}_* = 1 \quad (73)$$

leads to the derivation that $k_* \simeq 0.238$ and $A_* \simeq 0.843$.

To calculate the mass fraction of a stellar population returned to the ISM over a specified period of time, one must also know both the lifetimes (t_{life}) and end-of-life remnant masses (\mathcal{M}_{rem}) of stars as a function of their birth mass. For the former, we impose the textbook standard

$$\frac{t_{\text{life}}}{\text{Gyr}} = 10 \bar{\mathcal{M}}_{\star}^{-5/2} \quad (74)$$

(Harwit 1988). This is readily derived from the assertions that the luminosity of a star scales as $\mathcal{M}_{\star}^{3.5}$ and that the Sun has an approximate lifetime of 10 Gyr.

For the remnant masses of stars, we roughly follow Madau & Dickinson (2014) in effecting

$$\bar{\mathcal{M}}_{\text{rem}}(\bar{\mathcal{M}}_{\star}) = \begin{cases} \bar{\mathcal{M}}_{\star}, & \bar{\mathcal{M}}_{\star} < 1 \\ 0.444 + 0.084 \bar{\mathcal{M}}_{\star}, & 1 \leq \bar{\mathcal{M}}_{\star} \leq 7 \\ 0.419 + 0.109 \bar{\mathcal{M}}_{\star}, & 7 < \bar{\mathcal{M}}_{\star} < 8 \\ 1.4, & 8 \leq \bar{\mathcal{M}}_{\star} \leq 50 \\ \bar{\mathcal{M}}_{\star}, & \bar{\mathcal{M}}_{\star} > 50 \end{cases} \quad (75a)$$

$$\bar{\mathcal{M}}_{\text{rem}} \equiv \mathcal{M}_{\text{rem}}/M_{\odot}. \quad (75b)$$

The term for $\bar{\mathcal{M}}_{\text{rem}}(7 < \bar{\mathcal{M}}_{\star} < 8)$ is something we have manually introduced as a means of linearly interpolating between the adjacent mass ranges, which more closely follow Madau & Dickinson (2014). The value of $\mathcal{M}_{\text{rem}} = 1.4 M_{\odot}$ for $8 \leq \bar{\mathcal{M}}_{\star} \leq 50$ assumes that the remnant neutron star or BH has a mass close to the upper limit of white dwarfs (à la Chandrasekhar 1931). The basis for $\bar{\mathcal{M}}_{\text{rem}} = \bar{\mathcal{M}}_{\star}$ when $\bar{\mathcal{M}}_{\star} > 50$ is that the collapse of the star is so rapid, its entire mass is committed to the remnant BH (which we still class as ‘stellar mass’ in the model). Given how quickly the IMF diminishes at high stellar masses (equation 72a), variations on the assumptions for $\bar{\mathcal{M}}_{\text{rem}}(\bar{\mathcal{M}}_{\star} \gtrsim 50)$ are inconsequential for our treatment of stellar evolution and feedback.

The mass returned to the ISM^k for a given star equals $\mathcal{M}_{\star} - \mathcal{M}_{\text{rem}}$. We approximate that that mass is returned instantaneously at $\Delta t = t_{\text{life}}$ after the star’s birth. We can therefore calculate the fraction of a stellar population’s initial mass that is returned to the ISM over a cosmic-time interval $[t_0 - t_{\text{form}}, t_1 - t_{\text{form}}]$ as

$$f_{\star}^{\text{return}} = \int_{\bar{\mathcal{M}}_1}^{\bar{\mathcal{M}}_0} [\bar{\mathcal{M}}_{\star} - \bar{\mathcal{M}}_{\text{rem}}(\bar{\mathcal{M}}_{\star})] \phi(\bar{\mathcal{M}}_{\star}) d\bar{\mathcal{M}}_{\star}, \quad (76a)$$

$$\bar{\mathcal{M}}_x = \left(\frac{t_x - t_{\text{form}}}{10 \text{ Gyr}} \right)^{-2/5}, \quad (76b)$$

where the latter comes directly from equation (74), with $x = 0, 1$, while t_{form} denotes the time at which the stellar population formed. Note that because we define $t_1 > t_0$ and higher-mass stars die sooner, it follows that $\bar{\mathcal{M}}_1 < \bar{\mathcal{M}}_0$, which is why $\bar{\mathcal{M}}_1$ is the lower bound for the integral. The fraction of the remaining mass of a population (where the returned mass from all stars with $\bar{\mathcal{M}}_{\star} > \bar{\mathcal{M}}_0$ has already been subtracted from the population’s total mass) that is returned to the ISM can likewise be calculated as

$$f_{\text{rem}}^{\text{return}} = f_{\star}^{\text{return}} \frac{m_{\star, \text{pop}}^{\text{form}}}{m_{\star, \text{pop}}^{\text{rem}}}, \quad (77a)$$

$$\frac{m_{\star, \text{pop}}^{\text{form}}}{m_{\star, \text{pop}}^{\text{rem}}} = \left[\int_{0.1}^{\bar{\mathcal{M}}_0} \bar{\mathcal{M}}_{\star} \phi(\bar{\mathcal{M}}_{\star}) d\bar{\mathcal{M}}_{\star} + \int_{\bar{\mathcal{M}}_0}^{100} \bar{\mathcal{M}}_{\text{rem}}(\bar{\mathcal{M}}_{\star}) \phi(\bar{\mathcal{M}}_{\star}) d\bar{\mathcal{M}}_{\star} \right]^{-1}, \quad (77b)$$

where $m_{\star, \text{pop}}^{\text{form}}$ and $m_{\star, \text{pop}}^{\text{rem}}$ represent the formation mass and remaining mass of the stellar population, respectively. This expression is practically the most useful in the design of DARK SAGE, as we only explicitly track the remaining mass of a given population.

When mass is returned from a stellar population to the ISM, metals are returned to the ISM in proportion to the metallicity of that stellar population. In addition to this, new metals are added to the ISM, born from fusion in the stars’ cores and during supernova explosions. We maintain that the mass of new metals produced by a stellar population over its lifetime is a fixed fraction of the population’s birth mass (ignoring the metals already present). This fraction is known as the ‘yield’, which is expected to be ~ 0.03 , based on a Chabrier (2003) IMF combined with a Conroy, Gunn & White (2009) simple stellar population model (Lagos *et al.* 2018). These new metals are deposited into the ISM at a rate proportional to the mass-loss rate of stars in that population. This occurs *before* any reheating from feedback (Section 8.2). Note that we do not track the individual abundances of metal elements, only their summed content.

Let us hypothesise for a moment that we know the *average* energy released by a supernova that is used in heating gas out of the ISM and/or ejecting it out of the CGM, which we label \mathcal{E}_{SN} (note that we make no assumptions about the *total* energy released per supernova). Then we must ask: how many supernovae will there be for a given stellar population and when will they explode? Equation (74) already provides us the answer to the latter — any star that produces a supernova does so at the end of its life. For the former, we assume all stars with $8 \leq \bar{\mathcal{M}}_{\star} \leq 50$ produce a (core-collapse/Type-II) supernova. Additionally, we assume some fraction of stars with $1 \leq \bar{\mathcal{M}}_{\star} \leq 8$ produce a (accretion-driven/Type-Ia) supernova. This means we effectively assume that each star with $1 \leq \bar{\mathcal{M}}_{\star} \leq 8$ has a fixed probability, $p_{\text{SN-Ia}}$, of producing a supernova at the end of its life. The total number of supernovae of each type per unit mass of stars formed in a population, n_{SN} , is thus

$$\frac{n_{\text{SN-II}}}{M_{\odot}^{-1}} = \int_8^{50} \phi(\bar{\mathcal{M}}_{\star}) d\bar{\mathcal{M}}_{\star}, \quad (78a)$$

$$\frac{n_{\text{SN-Ia}}}{M_{\odot}^{-1}} = p_{\text{SN-Ia}} \int_1^8 \phi(\bar{\mathcal{M}}_{\star}) d\bar{\mathcal{M}}_{\star}. \quad (78b)$$

To determine $p_{\text{SN-Ia}}$, we look to observations. Broadly speaking, the rate of Type-II supernovae are observed to exceed that of Type Ia. For example, Tsujimoto *et al.* (1995) note that in the Milky Way, the ratio of Type-Ia to Type-II supernovae — equivalent to $n_{\text{SN-Ia}}/n_{\text{SN-II}}$ — is ~ 0.15 , but that it is more like 0.2–0.3 in the Magellanic Clouds. By contrast, Maggi *et al.* (2016) find $n_{\text{SN-Ia}}/n_{\text{SN-II}} \simeq 0.74$ in the Large Magellanic Cloud (LMC), but they expect this to be above the true average ratio due to the short timescale that their data probe (order 10^4 yr) and the shape of the LMC’s star formation history. In the absence of tighter observational constraints, we choose to adopt the Milky Way value of $n_{\text{SN-Ia}}/n_{\text{SN-II}} = 0.15$. Through equation (78), we can therefore solve for $p_{\text{SN-Ia}} \simeq 0.01$. In line with equation (76), we can calculate the

^kStrictly speaking, the mass is only returned to the ISM for stars in DARK SAGE discs. Otherwise it directly ‘returns’ to the CGM, as described in Section 8.5. But we use the phrase ‘returned to the ISM’ throughout the paper for simplicity.




Figure 7. Cumulative mass fraction of a stellar population returned to the ISM (top, equation 76) and cumulative number of supernovae per unit initial mass of that population (bottom) as a function of star birth mass (left) and time since the population’s formation (right) implemented in our new stellar evolution/feedback model in DARK SAGE. For reference, the vertical lines in the right-hand panels represent the age bins used in DARK SAGE for this work (treating t_{form} as 0; see Section 3.2). For stellar populations born at $t_{\text{form}} > 0$, the vertical lines would effectively be shifted to the left to describe how the evolution of that population is discretised in DARK SAGE.

total energy imparted on the local ISM by a stellar population over a given time interval, which we frame as a birth-mass interval, as

$$\Delta E_{\text{SN}} = \mathcal{E}_{\text{SN}} m_{*,\text{pop}}^{\text{form}} \int_{\bar{\mathcal{M}}_1}^{\bar{\mathcal{M}}_0} \varphi(\bar{\mathcal{M}}_*) d\bar{\mathcal{M}}_*, \quad (79a)$$

$$\varphi(\bar{\mathcal{M}}_*) \equiv \begin{cases} p_{\text{SN-Ia}} \phi(\bar{\mathcal{M}}_*), & \bar{\mathcal{M}}_* < 8 \\ \phi(\bar{\mathcal{M}}_*), & \bar{\mathcal{M}}_* \geq 8 \end{cases} \quad (79b)$$

We show how our mass-return and supernova models look as a function of star birth mass and time since population birth in Figure 7.

All stars are assumed to form at the middle of an age bin. Returned mass, the production of new metals, and feedback from stars born within the same age bin occur instantaneously, always assuming half the age bin has elapsed, regardless of where in that bin the current sub-time-step actually is. This is necessary, as once the model has moved into the time window of the next age bin, all stars formed in the previous age bin are assumed to be part of the same population with a single t_{form} . This is the nature of discretised time; regardless of how many bins one has, one always has to treat the ‘current’ bin in a special way. This also means that if we were to set the number of age bins to 1, we would return to the classical instantaneous recycling approximation (with the caveat that the recycling fraction is solved for, rather than given as an explicit input). Whenever star formation occurs, we check if it is necessary to rescale the local SFR to ensure that the mass consumed in star formation plus that associated with the instantaneous portion of stellar feedback does not exceed the mass of the gas annulus (in line with equation 64).

Several numbers in our model for stellar feedback could have been treated as free parameters. For example, the ratio of Type-Ia to Type-II supernovae for a population of stars, the parameters controlling the shape of the IMF, and the normalisation and power index of stellar life-times as a function of birth mass. Instead, we have opted to fix those numbers, and allow *one* stellar feedback parameter to vary freely in our calibration: \mathcal{E}_{SN} . Our rationale is that variation in most of the other parameters would be degenerate with changes to \mathcal{E}_{SN} to first order for calibration purposes. The exact value of \mathcal{E}_{SN} should, therefore, not be interpreted too literally, as it compensates for uncertainty in numerous assumptions we have made. Nevertheless, we should not expect \mathcal{E}_{SN} to depart too far from the canonical value of 10^{44} J (for a discussion on the energy output of supernovae, see e.g. Smartt 2009). Sure enough, as we show in Section 13, we find the best value of \mathcal{E}_{SN} to be $\sim 0.8 \times 10^{44}$ J.

There are further details one may wish to consider in a stellar feedback model, such as how the metallicity of a stellar population affects the rate of supernovae (e.g. Kistler et al. 2013), how the energetics of Type-Ia and Type-II supernovae differ, and how the IMF may change with metallicity and formation time (e.g. Li et al. 2023). Implicitly, our ignorance of these details is also folded in to the best-fitting value of \mathcal{E}_{SN} .

8.2. Impartation of energy on the ISM

With the energy from stellar feedback determined, the next step is to decide how that energy is used. The standard practice for doing this is to reheat an appropriate mass of gas out of the ISM and/or eject it. There are three major differences between how we do this now versus previous versions of DARK SAGE:

1. We do not impose a mass loading factor. Instead, we *solve* for the amount of gas reheated out of the ISM using energy-conservation arguments, as we outline below.
2. We split this reheated mass into two categories: the fraction with enough energy to escape the halo’s virial radius, and that which will remain bound to the CGM.
3. We do not instantaneously transfer mass to the hot and/or ejected reservoirs, but instead initially transfer them to the new fountain and outflowing reservoirs, thereby enforcing a delay before that gas is available for cooling.

Let us imagine that a supernova outflow begins carrying the mass prescribed by its return fraction of the source stellar population according to its age and the IMF assumptions laid out in Section 8.1. This initial blast then collides with the ISM, picking up mass and driving a wind that removes that mass from the ISM. Once mixed, the gas must also be heated to the temperature of the CGM, which is assumed to be uniform (equation 8). Assuming this wind conserves energy, we have

$$\Delta E_{\text{SN}}^{(i)} = \frac{1}{2} \Delta m_{\text{cold} \rightarrow \text{CGM}}^{(i)\text{SN}} \left(V_{200c}^2 + v_{\text{wind}}^{(i)2} \right). \quad (80)$$

Here, we have defined $v_{\text{wind}}^{(i)}$ in the annulus’s rotating frame as *the equivalent initial wind velocity were the outflow were to hypothetically reach thermal equilibrium with the CGM immediately*. In reality, we would expect the wind to start colder and faster, gradually heating and slowing (i.e. exchanging kinetic energy for thermal energy *in addition* to gravitational potential energy) before fully mixing with the CGM. Defining v_{wind} as we have

simply provides a mathematical shortcut to reach our desired outcome (below). While we *do* model how long it takes an outflow to incorporate into the CGM (or travel beyond; see Section 10), this timescale is *not* informed by v_{wind} .

DARK SAGE is designed such that any feedback-affected gas must have its energy raised to at least that of the hot-gas reservoir. As such, we can solve for $v_{\text{wind}}^{(i)}$ under the condition that

$$\Delta E_{\text{SN}}^{(i)} \rightarrow (e_{\text{hot}} - e_{\text{cold}}^{(i)}) \Delta m_{\text{cold} \rightarrow \text{CGM}}^{(i)\text{SN}},$$

where each e represents the specific energy of the subscripted reservoir. Combining with equation (80) and rearranging gives

$$v_{\text{wind}}^{(i)} = \sqrt{2 \left(e_{\text{fount}} - e_{\text{cold}}^{(i)} \right) - V_{200c}^2}. \quad (81)$$

The specific energy of the hot-gas reservoir can be broken into three components: its specific potential energy, specific thermal energy, and specific kinetic energy (i.e. that which is non-thermal, as noted to be non-negligible by e.g. Lochhaas *et al.* 2021), that is,

$$e_{\text{hot}} = \langle \Phi \rangle_{\text{hot}} + e_{\text{hot}}^{\text{kin}} + e_{\text{hot}}^{\text{therm}}. \quad (82)$$

To calculate the average potential of gas in the CGM, we numerically solve

$$\langle \Phi \rangle_{\text{hot}} = \frac{4\pi}{m_{\text{hot}}} \int_0^{R_{200c}} \rho_{\text{hot}}(R) \Phi(R) R^2 dR. \quad (83)$$

The potential profile, $\Phi(R)$, is found by numerically integrating equation (23).¹ Note that $\Phi(r)$ is defined *in the plane of the disc* in equation (23). As a necessary approximation for computational efficiency, we assume this functional form works for three-dimensional radii, R , as well. In the limit where the ratio of the disc mass to halo mass goes to zero, this approximation is accurate. However, in general, it is imperfect.

The specific kinetic energy of hot gas is informed entirely by bulk rotational motion. We approximate its bulk rotational velocity based on its specific angular momentum (assumed here to be equal to that of the halo) and the average radius of a gas element within it:

$$e_{\text{hot}}^{\text{kin}} = \frac{j_{\text{hot}}^2}{2 \langle R \rangle_{\text{hot}}^2} = \frac{j_{\text{halo}}^2}{2 \langle R \rangle_{\text{hot}}^2}, \quad (84a)$$

$$\langle R \rangle_{\text{hot}} = \frac{4\pi}{m_{\text{hot}}} \int_0^{R_{200c}} \rho_{\text{hot}}(R) R^3 dR \quad (84b)$$

$$= \frac{1 - \ln(2)}{2} R_{200c} c_{\beta}^2(z) C_{\beta}(z). \quad (84c)$$

As we assume $T_{\text{hot}} = T_{\text{vir}}$ throughout the CGM, the thermal specific energy of the hot-gas reservoir is readily found as

$$e_{\text{hot}}^{\text{therm}} = \frac{V_{200c}^2}{2}. \quad (85)$$

For satellites, we use V_{200c} at infall here, in line with our assumption that T_{hot} remains fixed after infall from Section 4.

¹Because we are only interested in *differences* in specific energy between reservoirs, the integration constant returned when calculating potential energy is irrelevant. But for practical reasons in the codebase, we take $\Phi(r_{\text{max}}) = 0$, where generally $r_{\text{max}} \gg R_{200c}$, meaning $\Phi(r \leq R_{200c}) < 0$.

Similar to equation (82), the specific energy of an annulus in the ISM can be broken down as

$$e_{\text{cold}}^{(i)} = \langle \Phi \rangle^{(i)} + e_{\text{cold}}^{(i)\text{kin}}. \quad (86)$$

No thermal term is added for the ISM; we have assumed that the temperature of the ISM is negligible compared to that of the CGM, such that

$$T_{\text{CGM}} - T_{\text{cold}} \simeq T_{\text{vir}}.$$

The potential specific energy of an annulus follows equation (36), while its specific kinetic energy is readily found through contributions from its circular velocity and average vertical motion:

$$e_{\text{cold}}^{(i)\text{kin}} = \frac{\langle v_{\text{circ}}^{(i)} \rangle^2 + \sigma_{\text{cold}}^2(z)}{2}. \quad (87)$$

While the above provides the necessary equations to calculate $\Delta m_{\text{cold} \rightarrow \text{CGM}}^{(i)\text{SN}}$, we are yet to specify how this feedback-affected gas is handled in detail. As noted above, we distribute this gas into two transitory reservoirs: the fountain reservoir and the outflowing reservoir, meaning

$$\Delta m_{\text{cold} \rightarrow \text{CGM}}^{(i)\text{SN}} = \Delta m_{\text{cold} \rightarrow \text{fount}}^{(i)\text{SN}} + \Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i)\text{SN}}. \quad (88)$$

The fountain reservoir is defined to have the same specific energy as the hot-gas reservoir:

$$e_{\text{fount}} = e_{\text{hot}}. \quad (89)$$

We track a timescale for which this gas is reincorporated in the hot-gas reservoir, which we describe in Section 10.1. The outflowing reservoir represents gas with energy to escape the (sub)halo's virial radius, and is gradually transferred to the ejected reservoir (see Section 10.2). How, then, do we decide the fraction of $\Delta m_{\text{cold} \rightarrow \text{CGM}}^{(i)\text{SN}}$ that goes to each of the fountain and outflowing reservoirs?

Consider the ideal that any gas kicked from the ISM with a velocity greater than the halo escape velocity,

$$\langle v_{\text{esc}}^{(i)} \rangle^2 = 2 \left[\Phi(R_{200c}) - \langle \Phi \rangle^{(i)} \right], \quad (90)$$

should be placed in the outflowing reservoir, and that the remainder goes to the fountain reservoir. For an observer external to the galaxy — that is, in the inertial frame of the halo — the appropriate velocity to compare against a given $\langle v_{\text{esc}}^{(i)} \rangle$ is not simply $v_{\text{wind}}^{(i)}$, but rather the vector sum of it and the local rotation velocity (recall that v_{wind} is defined to have already taken care of the fact that the wind slows from heating, which is what allows us to compare it to the escape velocity in the first place). We assume that the energy from a supernova is emitted with spherical symmetry in the rotating frame of the annulus. Any gas on the leading side will therefore be kicked with a heightened speed, and conversely that on the trailing side to a lower speed. One can solve for the angle between these two vectors, where the magnitude of the sum is equal to the escape velocity, through the cosine rule. The fraction of gas expelled from the ISM with a velocity above the escape velocity of the halo is then equivalent to the fraction of possible angles less than this critical angle. Recognising that the range of possible angles follows a $\sin(\theta)$ distribution — where $\theta \in [0, \pi]$ — then provides the ratio of ejected gas to reheated gas (after some omitted mathematics):

$$\frac{\Delta m_{\text{cold} \rightarrow \text{fount}}^{(i)\text{SN}}}{\Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i)\text{SN}}} = \left(\frac{1}{2} - \frac{\langle v_{\text{esc}}^{(i)} \rangle^2 - \langle v_{\text{circ}}^{(i)} \rangle^2 - v_{\text{wind}}^{(i)2}}{4 \langle v_{\text{circ}}^{(i)} \rangle v_{\text{wind}}^{(i)}} \right)^{-1} - 1. \quad (91)$$

Built into equation (91) is the implicit assumption that

$$\langle v_{\text{circ}}^{(i)} \rangle - v_{\text{wind}}^{(i)} < \langle v_{\text{esc}}^{(i)} \rangle < \langle v_{\text{circ}}^{(i)} \rangle + v_{\text{wind}}^{(i)}.$$

However, there is freedom in the solution for $v_{\text{wind}}^{(i)}$ to not obey this restricted range. If $\langle v_{\text{circ}}^{(i)} \rangle + v_{\text{wind}}^{(i)} < \langle v_{\text{esc}}^{(i)} \rangle$, then none of the gas kicked out of the ISM is capable of escaping the halo, meaning we set $\Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i)\text{SN}} = 0$. For the inverse reason, if $\langle v_{\text{circ}}^{(i)} \rangle - v_{\text{wind}}^{(i)} > \langle v_{\text{esc}}^{(i)} \rangle$, we set $\Delta m_{\text{cold} \rightarrow \text{fount}}^{(i)\text{SN}} = 0$.

In our framework, energy conservation dictates that

$$\Delta E_{\text{SN}}^{(i)} = \Delta m_{\text{cold} \rightarrow \text{fount}}^{(i)\text{SN}} (e_{\text{fount}} - e_{\text{cold}}^{(i)}) + \Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i)\text{SN}} (e_{\text{outfl}}^{(i)\text{SN}} - e_{\text{cold}}^{(i)}). \quad (92)$$

The final quantity left to define is $e_{\text{outfl}}^{(i)\text{SN}}$, which is the specific energy of the material presently being added to the outflow reservoir from a given annulus. In truth, this mass has a distribution of specific energies, for which we know the minimum and maximum, based on the picture presented above:

$$e_{\text{outfl}}^{(i)\text{SN}} \geq \Phi(R_{200c}) + e_{\text{hot}}^{\text{therm}},$$

$$e_{\text{outfl}}^{(i)\text{SN}} \leq \langle \Phi^{(i)} \rangle + e_{\text{hot}}^{\text{therm}} + \frac{1}{2} \left[\left(\langle v_{\text{circ}}^{(i)} \rangle + v_{\text{wind}}^{(i)} \right)^2 + \sigma_{\text{gas}}^2 \right].$$

We approximate $e_{\text{outfl}}^{(i)\text{SN}}$ as the average of these two limits:

$$e_{\text{outfl}}^{(i)\text{SN}} = e_{\text{hot}}^{\text{therm}} + \frac{\Phi(R_{200c}) + \langle \Phi^{(i)} \rangle}{2} + \frac{1}{4} \left[\left(\langle v_{\text{circ}}^{(i)} \rangle + v_{\text{wind}}^{(i)} \right)^2 + \sigma_{\text{gas}}^2 \right]. \quad (93)$$

Unlike the other reservoirs, the specific energy of the outflow reservoir is updated as mass is added to it:

$$e_{\text{outfl}}(t + \Delta t) = \max \left[\Phi(R_{200c}) + e_{\text{hot}}^{\text{kin}} + e_{\text{hot}}^{\text{therm}}, \frac{m_{\text{outfl}}(t) e_{\text{outfl}}(t) + \sum_{i=0}^{N_{\text{ann}}} \Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i)\text{SN}} e_{\text{outfl}}^{(i)\text{SN}}}{m_{\text{outfl}}(t) + \sum_{i=0}^{N_{\text{ann}}} \Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i)\text{SN}}} \right]. \quad (94)$$

The first term inside the ‘max’ function ensures the outflowing energy exceeds the fountaining energy, which it must by definition. It must at least carry enough energy to reach R_{200c} and have the equivalent of the thermal and kinetic specific energy of the hot gas. e_{outfl} is needed in other parts of the model, which we come to later in the paper.

We come back to what happens to the feedback-affected gas next in Section 10.

8.3. Mass loading factor

Our new stellar feedback model means we now have an entirely predictive framework for determining the mass loading factors of galaxies experiencing feedback. The mass loading factor of a galaxy is defined as the ratio of its gas reheating rate from stellar feedback to its star formation rate:

$$\eta \equiv \frac{\dot{m}_{\text{cold} \rightarrow \text{fount}}^{\text{SN}} + \dot{m}_{\text{cold} \rightarrow \text{outfl}}^{\text{SN}}}{\text{SFR}}, \quad (95)$$

where this excludes any consideration of gas heating from quasar feedback, but includes all channels of star formation. Here, \dot{m} and SFR refer to the mass transfer rates and star formation rates averaged over the previous *full* time step. This is physically more

informative than the most recent sub-time-step, as it is less susceptible to the order in which we resolve galaxy evolution processes (this is why the sub-time-steps are there in the first place).

DARK SAGE is one of few SAMs in the literature to make a true prediction for the mass loading factor from stellar feedback (for another example, see Lagos, Lacey & Baugh 2013). Comparing DARK SAGE’s mass loading predictions with observations and other models is not trivial; definitions of mass loading factors in the literature are far from uniform. This is in part because both real and hydrodynamic-simulation galaxies are not so easily broken into discrete reservoirs as SAMs define galaxies and haloes to have. For the former, it is typically necessary to measure the outflow rate (the numerator in equation 95) at some radius. Our definition of η is effectively measured at a radius of zero. This means our definition of η should be systematically larger than that calculated from observational studies (and hydrodynamic simulations).

While we make a face value comparison with some observational data from the literature in the left panel of Figure 8, for the above reason and others (see discussion in Nelson et al. 2019a on this topic), we caution against over-interpreting the apparent off-set DARK SAGE has with the data. Points for individual galaxies at $z \simeq 0$ are shown from Chisholm et al. (2017), as they are for $z < 0.2$ galaxies from Heckman et al. (2015), while the points from Sugahara et al. (2017) are bins of galaxies from their $z \simeq 0$ sample.

In the right-hand panel of Figure 8, we compare DARK SAGE’s predicted mass loading factors as a function of maximum halo circular velocity against the *input* functional mass loading factors of other SAMs.^m The assumed functions used in Bower et al. (2012) and Lagos et al. (2018) are similar in slope to what DARK SAGE predicts, but systematically higher and lower, respectively. Meanwhile, the asymptotic function of Guo et al. (2011) and the constant used by Croton et al. (2016) appear both qualitatively and quantitatively inconsistent with our energy-conserving framework. The comparison with Lagos et al. (2018) is particularly interesting, as their fitting function was based on the predictive model of Lagos et al. (2013). There are *many* differences between DARK SAGE and the Lagos et al. (2013) version of GALFORM, placing an elaborative discussion and comparison between the models outside the scope of this paper.

The slope and scatter of η as a function of either m_* or V_{max} are the outcome of our complex network of equations in DARK SAGE, although they ultimately boil down to the diversity of gravitational potentials at fixed mass and the increase in average potential depth at higher masses. The normalisation of η in DARK SAGE is effectively set by the value of \mathcal{E}_{SN} . As we show in Section 13, this is calibrated to other observations in a way that is blind to Figure 8.

8.4. Metal loading

In a departure from previous versions of DARK SAGE, we do *not* assume that gas reheated from the ISM carries the same metallicity as the ISM. We imagine a scenario where the energy from a stellar feedback event expands with spherical symmetry (in the ISM’s frame) into an ISM where all atoms (metals and non-metals) are distributed randomly and are isothermal. If the cross section of each atom were the same, we would expect the same amount of energy to be imparted on each atom on average. While heavier

^mSometimes, SAMs will invoke a mass loading factor as a function of virial velocity rather than maximum circular velocity. This might change the preciseness of a comparison with DARK SAGE, but it does not change our overall interpretation.




Figure 8. Mass loading factor predicted by DARK SAGE as a function of stellar mass (left) and maximal circular velocity (right-hand panel) for central galaxies with $s\text{SFR} \geq 10^{-11} \text{ yr}^{-1}$ at $z = 0$. Greyscale hexbins represent a two-dimensional histogram of DARK SAGE galaxies, where the solid, purple lines follow the running median (thick) and 16th and 84 percentiles (thin). Compared are observations in the left-hand panel (Heckman et al. 2015; Chisholm et al. 2017; Sugahara et al. 2017). Based on the captions of tables 1 and 2 of Heckman et al. (2015), we assume an error of 0.3 dex in stellar mass and $\sqrt{0.2^2 + 0.25^2}$ dex in η . Various other models are compared in the right-hand panel (including Guo et al. 2011; Bower, Benson & Crain 2012; Croton et al. 2016; Lagos et al. 2018). Note that V_{max} is from SUBFIND, which is different to what we derive from the rotation curves we build; this is the fairest quantity to compare to other semi-analytic models.

elements may have a greater cross section for interacting with an outflow when stationary, this is offset by the lighter elements having higher thermal motion at fixed temperature. In the absence of a detailed model for these competing effects, we approximate the energy imparted on each atom to be equal. Because we have implicitly assumed elsewhere that anything reheated carries a common *specific* energy, this approximation must mean some metals are left behind. As oxygen is the most abundant metal in the Universe, and carbon the second-most abundant, the average atomic mass of a metal should be a little less than $16 m_{\text{H}}$. Non-metals (hydrogen and helium) should have an average atomic mass of around $7 m_{\text{H}}/4$. This implies that for every unit mass of non-metals reheated from an annulus, we should see $\sim Z_{\text{cold}}^{(i)}/9$ times that mass of metals reheated. Or, in other words, the metallicity of the locally produced outflow is one tenth that of the local ISM. This is what we enforce.

One potential issue with our treatment of metals in stellar feedback outflows is that some observations find evidence that the metallicity of outflows is typically *elevated* relative to the ISM (e.g. Chisholm, Tremonti & Leitherer 2018), which is the opposite of what we have imposed. However, there is once again more nuance than this in comparing our model to observations. As we explain more in Section 10, the outflow that we follow from a galaxy in DARK SAGE represents that gas that will eventually either redistribute itself within the CGM, with the same energy as the CGM, and/or be ejected from the halo. We essentially consider all the material in an outflow to carry one of two possible energy states. This means we lose the information of an energy or velocity *distribution* in the outflow. In practice, we should expect the low-energy portion of that distribution to never fully reincorporate into the quasi-hydrostatic CGM, but instead to return to the ISM whence it came on a relatively short timescale. One might also hypothesise that an outflow has a metallicity distribution, which is anti-correlated with the velocity distribution, by virtue of metals being heavier. It is therefore conceivable that the metallicity of an outflow observed close to the disc of a galaxy




Figure 9. The relation between stellar mass and ISM metallicity for galaxies at $z = 0$. We only include DARK SAGE galaxies with $m_{\text{neutral}} \geq 10^8 M_{\odot}$ to ensure there is appreciable gas present for a meaningful gas metallicity. DARK SAGE data are plotted in the same way as Figure 8. Compared is the observed relation from Tremonti et al. (2004) at $z \simeq 0.1$, covering their 16–84th interpercentile range, and the running median from Bellstedt et al. (2021) at $z = 0.07$.

might differ significantly from the metallicity of the portion of that outflow that ends up mixing with the CGM. Moreover, the statistics on the metallicity of outflows are weak, and the measurement uncertainties are often dominated by systematics.

One observed relation that is much more enshrined in canon is that between the stellar mass and gas metallicity of galaxies. In previous versions of the model, this relation informed the calibration of the model's free parameters. As we discuss in Section 13, our calibration method has changed significantly, and the mass-metallicity relation of galaxies is no longer used. This means it is now a prediction of the model. We show this relation for DARK SAGE in Figure 9. This is compared to the observation-based relations of Tremonti et al. (2004) and Bellstedt et al. (2021).

We note that there are hidden systematic uncertainties in the compared mass–metallicity relations in Figure 9, which should not make the apparent systematic offset between them and DARK SAGE a point of concern. For Bellstedt et al. (2021), assumptions in their method of modelling galaxies’ spectral energy distributions are a natural source of systematic uncertainty, such as their enforced proportionality in the evolution of Z_{gas} with stellar-mass growth. There is also ~ 1 Gyr more of cosmic evolution from their $z = 0.07$ result to our DARK SAGE result at $z = 0$. For the comparison with Tremonti et al. (2004), whose data represent $z \simeq 0.1$, the largest sources of systematic uncertainty are from the emission line diagnostics used to measure the oxygen abundance (see Kewley & Ellison 2008) and the conversion from oxygen abundance to metallicity. In our comparison, we assume

$$12 + \log_{10}(\text{O}/\text{H}) = 8.69 + \log_{10}(Z_{\text{gas}}/0.0142), \quad (96)$$

where 0.0142 is the estimated metallicity of the Sun (Asplund et al. 2009) (we have corrected the Tremonti et al. 2004 relation for a Chabrier 2003 IMF and our cosmology). For similar reasons, in their comparison of the mass–metallicity relation of the IllustrisTNG simulation with observations, Torrey et al. (2019) ignore the normalisation of the Tremonti et al. (2004) relation, moving it down in their figure 6. While we could have taken the same approach in our Figure 9, we opted not to for the sake of fidelity.

The change to how we model metal loading in DARK SAGE turned out to be crucial to the outcome shown in Figure 9. That is, had the outflows from the ISM carried higher metallicity, Z_{cold} would have been systematically lower in DARK SAGE galaxies. While we outlined above why *some* systematic offset between DARK SAGE and observations as presented in Figure 9 is expected — in this case, an apparent factor of < 1.4 — larger offsets naturally become increasingly harder to waive away.

With the resolved detail of DARK SAGE discs, we can also predict the gas metallicity *gradients* of galaxies. This is a rare quality for a SAMs, which few others have (as another example, see Yates et al. 2021). As a first prediction from the new DARK SAGE, we show the radius-weighted *average* metallicity gradient of DARK SAGE galaxies as a function of stellar mass in Figure 10, where

$$\left\langle \frac{d \log_{10}(Z_{\text{cold}})}{dr} \right\rangle \equiv \left[\sum_{i=1, m_{\text{cold}}^{(i)} \neq 0}^{N_{\text{ann}}-1} \log_{10}(Z_{\text{cold}}^{(i)}) \right] \times \left[\sum_{i=1, m_{\text{cold}}^{(i)} \neq 0}^{N_{\text{ann}}-1} ((r^{(i+1)}) - (r^i)) \right]^{-1}. \quad (97)$$

We choose to explicitly weight by the radial range of each annulus, as observational gradients are typically *implicitly* weighted in the same way, that is, data will be binned by radius and a gradient fitted to those bins. Such observations from integral field spectroscopy (IFS) surveys are compared in Figure 10 (to appear in Acharyya et al. in prep.),^a though we caution that this is still a cursory comparison, and should therefore not be interpreted as a means of falsifying DARK SAGE. While the vast majority of DARK SAGE galaxies have negative metallicity gradients (which is the

^aThe same data were meant to be presented in figure 2 of Sharda et al. (2021), but the points shown in the panel with gradients in units dex kpc⁻¹ were actually from a file in units of dex per effective radius (Acharyya, private communication).




Figure 10. Radius-weighted average of the metallicity gradient in gas discs in DARK SAGE galaxies at $z = 0$. Data for DARK SAGE are presented in the same way as in previous figures. Dotted lines with error bars represent the average relations from three IFS surveys: the ‘Calar Alto Legacy Integral Field Spectroscopy Area’ survey (CALIFA; Sánchez et al. 2014), the ‘Mapping nearby Galaxies at Apache Point Observatory’ survey (MaNGA; Belfiore et al. 2017), and the ‘Sydney-AAO Multi-object Integral-field spectrograph’ survey (SAMI; Poetrodjojo et al. 2021), which have been homogenized (Sharda et al. 2021; Acharyya et al. in prep.).

expected norm), DARK SAGE is also able to produce galaxies with positive metallicity gradients. The relation between gas metallicity gradient and stellar mass is relatively flat in DARK SAGE, as it is in observations. But this is dependent on how the metallicity gradient is calculated. Though not shown here, reproducing Figure 10 with a *mass-weighted* average gradient across the annuli in a DARK SAGE disc results in an increasingly steep relation towards low stellar masses. Metallicity gradients normalised by the stellar disc half-mass radius look very similar (but for a re-scaled y-axis).

8.5. Direct impartation of energy on the CGM

The instability-driven bulge, merger-driven bulge, and intrahalo stars (collectively ‘spheroid stars’) can produce feedback. While no stars form in spheroids in DARK SAGE, the delayed feedback scheme means that spheroid stars still produce Type-Ia supernovae. We assume that spheroid supernovae cannot reheat gas from the disc. Instead, the metals and energy they produce are imparted directly onto the CGM. Summing over all stellar-age bins, the energy from spheroid supernovae, $\Delta E_{\text{SN}}^{\text{sph}}$, is balanced by generating an outflow from the CGM. In practice, this means transferring mass from the hot and fountain reservoirs to the outflowing reservoir. Because we have defined $e_{\text{fount}} = e_{\text{hot}}$, we transfer mass *proportionally* out of the hot and fountain reservoirs, assuming this gas carries the same specific energy as that pre-computed for the outflowing reservoir:

$$\Delta m_{\text{hot} \rightarrow \text{outfl}}^{\text{SN}} = \min \left[m_{\text{hot}}, \frac{\Delta E_{\text{SN}}^{\text{sph}}}{e_{\text{outfl}} - e_{\text{hot}}} \frac{m_{\text{hot}}}{m_{\text{hot}} + m_{\text{fount}}} \right], \quad (98a)$$

$$\Delta m_{\text{fount} \rightarrow \text{outfl}}^{\text{SN}} = \min \left[m_{\text{fount}}, \frac{\Delta E_{\text{SN}}^{\text{sph}}}{e_{\text{outfl}} - e_{\text{hot}}} \frac{m_{\text{fount}}}{m_{\text{hot}} + m_{\text{fount}}} \right]. \quad (98b)$$

8.6. Comparative remarks

Although vastly different in how the feedback is *applied*, a clear parallel can be drawn between our model for delayed stellar feedback and that in the MERAXES SAM (Mutch *et al.* 2016). This is not entirely by coincidence; part of the inspiration for our model design came from Mutch *et al.* (2016). Similar designs of delayed feedback and chemical enrichment are also incorporated into the GALACTICUS model (Benson 2012) and in De Lucia *et al.* (2014), although again different in implementation to DARK SAGE. We refer the reader to the other papers for more context.

In previous incarnations of SAGE and DARK SAGE, the returning of mass from stars to the ISM and the production of new metals was done *after* reheating gas from feedback. Implicitly, this meant it was guaranteed that some (local, metal-rich) cold gas would forcibly remain after an episode of star formation. We have deliberately changed this. Now, mass is returned and metals added to the ISM *before* we apply feedback to the ISM. Provided there is sufficient energy, this means that returned gas can immediately find its way to the CGM.

We explored the possibility of a distributive feedback model in DARK SAGE, where the energy from supernovae in one annulus could affect the gas in adjacent annuli. Our default assumption was that energy would be distributed from an annulus centre isotropically. For infinitesimally wide annuli, one could imagine that most of the energy from feedback should be directed towards neighbouring annuli instead of out of the disc. However, with our set-up in DARK SAGE, it is almost always true that the width of a disc annulus is much greater than the disc scale height, implying a negligible fraction of feedback energy should transfer through the disc. As such, in the end, we reasonably approximated the feedback taking place in an annulus to be independent from other annuli in a galaxy.

9. Active galactic nuclei

We maintain a two-mode model for feedback from active galactic nuclei (AGNs) in DARK SAGE, but have changed the way both modes operate.

9.1. Radio mode

In earlier versions of DARK SAGE, the model for radio feedback followed Croton *et al.* (2016). In this mode, hot gas is passively accreted directly onto the BH, based on the BH's Bondi (1952) accretion rate. Combined with the 'maximal cooling flow' model of Nulsen & Fabian (2000), the accretion rate can be derived as

$$\dot{m}_{\text{BH}}^{\text{radio}} = \frac{15\pi G \bar{\mu} m_{\text{p}} k_{\text{B}} T_{\text{vir}}}{16 \Lambda(T_{\text{vir}}, Z_{\text{hot}})} m_{\text{BH}} \quad (99)$$

(cf. Croton *et al.* 2006), where m_{BH} is the BH's mass. In Croton *et al.* (2016), an additional free parameter was added to equation (99). However, with the changes we have made to the model, we no longer need this. We note that we explicitly do not allow $\dot{m}_{\text{BH}}^{\text{radio}}$ to exceed the BH's Eddington (1926) accretion rate in the code, but it is typically orders of magnitude below this anyway.

A half-century-old standard, we assume that mass accreting onto a BH quasi-statically moves through equatorial, circular orbits — all the while losing energy and angular momentum — until reaching the orbit of lowest energy, where it is captured (see Lynden-Bell 1969; Bardeen 1970; Bardeen, Press & Teukolsky 1972; Novikov & Thorne 1973; Shakura & Sunyaev 1973; Stevens

2015). The power emanating from the accretion disc — which drives an AGN — is then given directly by the accretion rate:

$$\dot{E}_{\text{AGN}} = \epsilon_{\text{AGN}} \dot{m}_{\text{BH}} c^2, \quad (100)$$

where c is the speed of light in a vacuum and ϵ_{AGN} is the ratio of inertial (gravitational) mass lost by that accreted to its rest mass (at infinity), often referred to as the 'radiative efficiency'. Note that \dot{m}_{BH} represents the *rest-mass* accretion rate of the BH, but m_{BH} represents the *inertial* mass of the BH. Over a sub-time-step, the growth of a BH is therefore

$$\Delta m_{\text{BH}} = (1 - \epsilon_{\text{AGN}}) \dot{m}_{\text{BH}} \Delta t. \quad (101)$$

The difference between rest mass and inertial mass was not properly accounted for in previous versions of the model. This means BH feedback is now even more self-regulatory. That is, for a fixed rest mass of accreted material, a BH can either grow by a lot and produce a weak AGN (low ϵ_{AGN}), or grow less and produce a stronger AGN (high ϵ_{AGN}). But the less the BH grows, the weaker any future radio-mode feedback will be. In principle, to conserve inertial mass (energy), that lost by the material accreted onto the BH should be added somewhere. Without doing so, the baryonic mass of the halo within which that BH resides will drop. In practice though, this is inconsequential not only because this comprises a negligible fraction of a halo's baryonic mass, but also because it is automatically accounted for by DARK SAGE's infall prescription anyway. At the next time step, baryonic mass will be added to the CGM of that halo to top it back up to the cosmic baryon fraction (modulo any suppression from photoionisation feedback, and including ejected gas in the accounting — see Section 4). We implicitly approximate rest mass and inertial mass as equivalent for all other components of galaxies and haloes (as literally all cosmological simulations and SAMs do), which we do not expect to have any adverse effects whatsoever.

Under the above picture, the value of ϵ_{AGN} depends exclusively on the spin of the BH (neglecting charge) and the relative orientation of the accretion disc to the spin vector (Bardeen *et al.* 1972). While $\epsilon_{\text{AGN}} = 0.0572$ for a non-rotating Schwarzschild (1916) BH, ϵ_{AGN} can theoretically range from 0.0378 to 0.3210, given the anti-parallel and parallel cases for a Kerr (1963) BH rotating at the Thorne (1974) limit, and can reach 0.4226 if the Thorne limit is breached. However, we do not explicitly consider BH spin in DARK SAGE (for an example of a SAM that does, see Fanidakis *et al.* 2011). The addition of such a feature would be an entire project in and of itself. In the meantime, we opt to treat ϵ_{AGN} as a free parameter that is applied uniformly for all BHs.

Equation (100) applies to both modes of accretion, but the way that that energy is handled depends on the mode. For the radio mode, the energy is used as a preventative form of feedback. This has been standard for SAMs since Croton *et al.* (2006) and Bower *et al.* (2006). After the gross cooling rate of a galaxy is calculated (Section 4), we calculate how much of that cooling would be prevented by the energy released from the BH accretion disc:

$$\Delta m_{\text{heat}}^{\text{radio}} = \frac{\dot{E}_{\text{AGN}}^{\text{radio}} \Delta t}{e_{\text{hot}} - \langle e_{\text{cold}} \rangle} = \frac{\Delta E_{\text{AGN}}^{\text{radio}}}{e_{\text{hot}} - \langle e_{\text{cold}} \rangle}, \quad (102a)$$

$$\langle e_{\text{cold}} \rangle \equiv \sum_{i=1}^{N_{\text{ann}}} e_{\text{cold}}^{(i)} \int_{j_{\text{inner}}^{(i)}}^{j_{\text{outer}}^{(i)}} \text{PDF}_{\text{cool}}(j) dj, \quad (102b)$$

providing the final pieces to equations (13–15).

The term ‘radio mode’ is derived from the presence of radio jets that are generated from significantly sub-Eddington rates of BH accretion. These pervade the CGM and are what keeps the gas hot, or, rather, prevents it from cooling. While we have no *explicit* consideration of AGN jets in DARK SAGE (but see Raouf et al. 2017, where there is an explicit consideration of this in the RADIO-SAGE model), the radio mode still implicitly relies on there being jets. The influence of these jets persists over extended timescales. For this reason, AGN memory in a SAM is important; without it, a model can be vulnerable to spurious, non-physical increases and decreases in radio-mode heating.

In Croton et al. (2016), radio-mode memory was achieved by introducing a ‘heating radius’, which never decreased with time and effectively prevented gas in the CGM internal to it from cooling. This radius would increase when periods of radio AGN activity exceeded what the galaxy had had previously. This model was adopted in previous versions of DARK SAGE, and has been adopted in other SAMs too (e.g. Lagos et al. 2018).

For the new DARK SAGE, we achieve radio-mode memory differently. We abandon any consideration of a heating radius and instead move mass from the hot reservoir to the fountain reservoir equal to that offset from cooling:

$$\Delta m_{\text{hot} \rightarrow \text{fount}}^{\text{radio}} = \Delta m_{\text{heat}}^{\text{radio}}. \quad (103)$$

This does not affect energy conservation, as the fountain and hot reservoirs are defined to have the same specific energy. But this *does* create a lasting impact by effectively disabling the AGN-affected gas from cooling for a dynamical time. This is not only less contrived than the Croton et al. (2016) heating radius model, but it also allows the suppression of cooling from radio-mode feedback in a galaxy to wax and wane moderately over time, rather than effectively forcing it to stay the same or increase.

Provided $m_{\text{BH}} > 0$ and $m_{\text{hot}} > 0$, radio-mode accretion and feedback will occur in a DARK SAGE galaxy at every sub-time-step.

9.2. Quasar mode

Whereas the radio mode is coupled to the CGM and is primarily responsible for the quenching of high-mass galaxies, the quasar mode is coupled to the ISM and is the primary growth channel for BHs. In earlier versions of SAGE and DARK SAGE, quasar-mode accretion and feedback were triggered in two instances: from disc instabilities and when galaxy mergers occurred. Instability-driven accretion in DARK SAGE requires gas to first flow into the innermost annulus; any unstable gas there that is not consumed/removed by a starburst and its associated feedback (Section 7.2) is accreted onto the central BH. The previous treatment of merger-driven BH accretion in DARK SAGE was *ad hoc* and not well motivated theoretically; its phenomenological design (see Kauffmann & Haehnelt 2000; Croton et al. 2016) was better suited for SAMs that only treat *global* galaxy properties. We have now abandoned this channel of quasar-mode accretion entirely, leaving only the instability channel. Though, this is not to say that mergers no longer affect when BHs grow. The additional gas placed in the relevant annuli after a merger inevitably drives an instability that typically cascades all the way into the galaxy centre. These merger-induced instabilities play a major role in BH growth. Effectively, all we have done is retire a redundant prescription; this has no adverse consequences on the model’s predictions and it means the model has fewer free parameters.

$\dot{m}_{\text{BH}}^{\text{quasar}}$ is set by the amount of unstable gas in the innermost annulus of a galaxy and one free parameter that controls what fraction of this gas is consumed in a starburst ($f_{\text{move}}^{\text{gas}}$, see Section 7.2). Unlike the radio mode, there is nothing to explicitly prevent quasar-mode accretion from exceeding the Eddington (1926) rate. Equation (100) then gives the power that drives quasar-mode feedback.

Quasar-mode feedback is ejective, not preventative. The energy is first dumped into the innermost annulus of the ISM, working outward until all the energy output over a given sub-time-step is used. Quasar feedback on the ISM differs from stellar feedback in that all feedback-affected gas is sent to the outflowing reservoir, with none going to the fountain reservoir. The physical justification for these choices is that the gas nearest the BH should be maximally affected before energy from it starts coupling to the outer ISM. Using the pre-computed specific energy of the outflowing reservoir (Section 8.2), we effect

$$\Delta m_{\text{cold} \rightarrow \text{outfl}}^{(i)\text{quasar}} = \min \left[m_{\text{cold}}^{(i)}, \frac{\Delta E_{\text{left}}^{(i)}}{e_{\text{outfl}} - e_{\text{cold}}^{(i)}} \right], \quad (104a)$$

$$\Delta E_{\text{left}}^{(i)} = \begin{cases} \Delta E_{\text{AGN}}^{\text{quasar}}, & i = 1 \\ \Delta E_{\text{AGN}}^{\text{quasar}} - \sum_{k=1}^{i-1} \Delta E_{\text{used}}^{(k)}, & i \geq 2 \end{cases}, \quad (104b)$$

$$\Delta E_{\text{used}}^{(k)} = \Delta m_{\text{cold} \rightarrow \text{outfl}}^{(k)\text{quasar}} (e_{\text{outfl}} - e_{\text{cold}}^{(k)}). \quad (104c)$$

This equation is solved for each annulus i until either reaching $i = N_{\text{ann}}$ or $\Delta E_{\text{left}}^{(i)} = 0$. Whenever the former is the trigger to stop, there must still be some non-zero excess energy

$$\Delta E_{\text{excess}}^{\text{quasar}} = \Delta E_{\text{AGN}}^{\text{quasar}} - \sum_{i=1}^{N_{\text{ann}}} \Delta E_{\text{used}}^{(i)}. \quad (105)$$

Identical to our treatment of excess energy from stellar feedback, we use as much of this energy as possible to eject gas from the hot reservoir:

$$\Delta m_{\text{hot} \rightarrow \text{outfl}}^{\text{quasar}} = \min \left[m_{\text{hot}}, \frac{\Delta E_{\text{excess}}^{\text{quasar}}}{e_{\text{outfl}} - e_{\text{hot}}} \frac{m_{\text{hot}}}{m_{\text{hot}} + m_{\text{fount}}} \right], \quad (106a)$$

$$\Delta m_{\text{fount} \rightarrow \text{outfl}}^{\text{quasar}} = \min \left[m_{\text{fount}}, \frac{\Delta E_{\text{excess}}^{\text{quasar}}}{e_{\text{outfl}} - e_{\text{hot}}} \frac{m_{\text{fount}}}{m_{\text{hot}} + m_{\text{fount}}} \right]. \quad (106b)$$

The main differences from the Stevens et al. (2016) implementation of quasar-mode AGN feedback is our far more comprehensive consideration of the specific energy differences between reservoirs, the allowance for an annulus to have its gas partially ejected (as opposed to it being ‘all or nothing’), and the timescale on which this gas reincorporates back into the hot reservoir (described in the next section).

10. The galactic fountain and gas reincorporation

We introduce a new treatment for the galactic fountain, which affects how gas is reincorporated into the hot-gas reservoir to become available for cooling again after being heated by feedback. Our method is new not only for DARK SAGE but for SAMs in general.

Early in their development, SAMs assumed the reincorporation timescale to scale linearly with the halo dynamical time

(equation 9) (Springel *et al.* 2001), with the proportionality constant often treated as a free parameter (per De Lucia *et al.* 2004). Motivated by the difficulty that this simple model posed in recovering observations at multiple epochs (Henriques *et al.* 2013; Mutch, Poole & Croton 2013), it later became commonplace for reincorporation timescales to have a direct dependence on virial mass or velocity (with various functional forms — e.g. Henriques *et al.* 2013; White, Somerville & Ferguson 2015; Croton *et al.* 2016). However, none of these approaches are in line with our philosophy of having a self-consistent feedback framework that minimises the number of free parameters. Instead, we introduce a framework where we update timescales for mass transfer between the CGM and ejected reservoir self-consistently with the energetics of feedback.

10.1. Fountaining in the CGM

Gas in the fountain reservoir is intended to reincorporate into the hot-gas reservoir to become available for cooling again. We must decide the timescale that that gas remains in the transitory fountain reservoir. In the absence of a detailed consideration of how efficiently gas mixes (beyond the scope of this version of DARK SAGE), we make a simple assumption that gas takes a halo dynamical time to transit the fountain reservoir. We track an average ‘fountain time’ whenever mass is added to the reservoir:

$$t_{\text{fount}}(t + \Delta t) = \max \left[0, \frac{m_{\text{fount}}(t) t_{\text{fount}}(t) + \Delta m_{\Sigma \rightarrow \text{fount}} t_{\text{dyn}}}{m_{\text{fount}}(t) + \Delta m_{\Sigma \rightarrow \text{fount}}} - \Delta t \right], \quad (107)$$

where $\Delta m_{\Sigma \rightarrow \text{fount}}$ is the sum of contributions to the fountain reservoir for the sub-time-step.

For each sub-time-step, we move a fraction of the fountaining mass to the hot reservoir in proportion to the fraction of t_{fount} that has elapsed:

$$\Delta m_{\text{fount} \rightarrow \text{hot}} = \min \left[1, \frac{\Delta t}{t_{\text{fount}}} \right] m_{\text{fount}}. \quad (108)$$

10.2. Ejection of outflowing gas

Gas in the outflowing reservoir is, by definition, eventually going to escape the halo. Again, we need a timescale to decide how long gas added to the outflowing reservoir will remain in it before being transferred to the ejected reservoir. We estimate the instantaneous outflow timescale based on the specific energy of the gas added to the reservoir at that time. If we assume the thermal energy of the outflowing gas remains constant, we can approximate the average velocity of outflowing gas travelling as that which balances energy at the halo centre and at the virial radius:

$$t_{\text{outfl}}^{\text{inst}} = \frac{2 R_{200c}}{v_{\text{outfl}}^{\text{inst}}(0) + v_{\text{outfl}}^{\text{inst}}(R_{200c})}, \quad (109a)$$

$$v_{\text{outfl}}^{\text{inst}}(R) \equiv \sqrt{2 (e_{\text{outfl}}^{\text{inst}} - \Phi(R) - e_{\text{hot}}^{\text{therm}})}. \quad (109b)$$

When mass is added to the outflow reservoir from supernova feedback,

$$e_{\text{outfl}}^{\text{inst}} = e_{\text{outfl}}^{\text{SN}}.$$

Otherwise

$$e_{\text{outfl}}^{\text{inst}} = e_{\text{outfl}}.$$

In the same manner as the fountain reservoir, we track an average outflow time as mass is continually added to the outflowing reservoir, ensuring to reduce this as cosmic time elapses:

$$t_{\text{outfl}}(t + \Delta t) = \max \left[0, \frac{m_{\text{outfl}}(t) t_{\text{outfl}}(t) + \Delta m_{\Sigma \rightarrow \text{outfl}} t_{\text{outfl}}^{\text{inst}}}{m_{\text{outfl}}(t) + \Delta m_{\Sigma \rightarrow \text{outfl}}} - \Delta t \right], \quad (110)$$

where $\Delta m_{\Sigma \rightarrow \text{outfl}}$ represents the additions to the outflowing reservoir from all sources for the sub-time-step. For central galaxies and satellites outside their host’s virial radius, mass is added to their ejected reservoir according to:

$$\Delta m_{\text{outfl} \rightarrow \text{ejec}} = \min \left[1, \frac{\Delta t}{t_{\text{outfl}}} \right] m_{\text{outfl}}. \quad (111)$$

Satellites inside the virial radius of their host do not have an ejected reservoir by construction. For these galaxies, the mass on the left-hand side of equation (111) is instead placed in the fountain reservoir of the central (and its fountain timescale is updated accordingly).

10.3. Effect of halo growth

As a halo grows in mass, the depth of its potential well and thermal energy both increase. This can cause outflowing gas that previously had enough energy to escape the halo to no longer have enough. Whenever a (sub)halo increases in mass, we first check whether *any* of the outflowing gas can still reach the virial radius. This requires

$$\Phi(R_{200c}, t) + e_{\text{hot}}^{\text{therm}}(t) > e_{\text{outfl}}(t - \Delta t).$$

If this inequality is not satisfied, all gas in the outflowing reservoir is transferred to the fountain reservoir. Otherwise, we compute the average specific energy required for the outflowing gas to maintain its outflow timescale.

While we only explicitly track the *average* outflow timescale and energy of gas in the CGM, one would physically expect the outflowing gas to have a *distribution* of energy. The lower-energy gas in this distribution must be that which is moved into the fountain reservoir as the result of halo growth. This should, therefore, *raise* the specific energy of the material that remains in the outflowing reservoir. In principle, this should reduce the outflow time. However, as the halo grows, the assumed temperature of the remaining outflowing gas rises — which means it gets slower, if we assume its total energy is conserved — and the distance it must travel to reach R_{200c} also increases. Both of these effects should raise the outflow time. Because there are competing effects that should both raise and lower t_{outfl} , we take the agnostic stance of keeping the outflow time constant, and then use energy conservation to calculate the fraction of the outflowing gas that should be moved to the fountain reservoir:

$$\Delta m_{\text{outfl} \rightarrow \text{fount}} e_{\text{fount}}(t) + m_{\text{outfl}}(t) e_{\text{outfl}}(t) = m_{\text{outfl}}(t - \Delta t) e_{\text{outfl}}(t - \Delta t), \quad (112a)$$

$$e_{\text{outfl}}(t) = \Phi(R_{200c}, t) + e_{\text{hot}}^{\text{therm}}(t) + \frac{1}{2} v_{\text{outfl}}^2(R_{200c}, t), \quad (112b)$$

$$v_{\text{outfl}}(R_{200c}, t) = \frac{2 R_{200c}(t)}{t_{\text{outfl}}} - \sqrt{2 [e_{\text{outfl}}(t - \Delta t) - \Phi(0, t) - e_{\text{hot}}^{\text{therm}}(t)]}. \quad (112c)$$

10.4. Reincorporation of ejected gas

When gas is added to the ejected reservoir from the outflowing reservoir, we calculate an expected time for that gas to reincorporate back into the halo. From the specific energy of the outflowing reservoir, we can calculate both the speed of the ejected gas at the virial radius, and the radius at which the gas will cease to have kinetic energy and must turn around. We approximate the average speed the gas moves during this period as half its ejected speed at the virial radius. With the understanding that the gas must travel the distance from R_{200c} to R_{turn} twice, we effect

$$t_{\text{reinc}}^{\text{inst}} = \frac{4(R_{\text{turn}} - R_{200c})}{v_{\text{outfl}}(R_{200c})}, \quad (113a)$$

$$\Phi(R_{\text{turn}}) + e_{\text{hot}}^{\text{therm}} = e_{\text{outfl}}. \quad (113b)$$

To solve these simultaneous equations efficiently in the code, we approximate the shape of the potential outside the virial radius with an analytic NFW profile.

As with the other timescales, we average the reincorporation timescale as more mass is added:

$$t_{\text{reinc}}(t + \Delta t) = \max\left[0, \frac{m_{\text{ejec}}(t) t_{\text{reinc}}(t) + \Delta m_{\text{outfl} \rightarrow \text{ejec}} t_{\text{reinc}}^{\text{inst}}}{m_{\text{ejec}}(t) + \Delta m_{\text{outfl} \rightarrow \text{ejec}}} - \Delta t\right]. \quad (114)$$

Ejected mass is then reincorporated directly into the hot-gas reservoir at a rate set by

$$\Delta m_{\text{ejec} \rightarrow \text{hot}} = \min\left[1, \frac{\Delta t}{t_{\text{reinc}}}\right] m_{\text{ejec}}. \quad (115)$$

10.5. Baryon breakdown

With all the baryon reservoirs and the equations for how mass moves between them in DARK SAGE now described, it is informative to see how the baryonic mass inside haloes is broken between them at $z = 0$. We show precisely this in Figure 11.

The stellar mass fraction as a function of halo mass is as one would expect, peaking at $M_{200c} \simeq 10^{12} M_{\odot}$. The CGM fraction increases monotonically, dominating the baryon budget at high halo masses. Of the mass in the CGM, the outflowing reservoir contains little mass at all halo mass scales, showing that the time spent in this transitory reservoir is less than the time spent in the ejected reservoir before reincorporation. A significant portion of the CGM is in the fountain reservoir for $M_{200c} \lesssim 10^{13} M_{\odot}$. This highlights that most of the gas reheated from stellar feedback does not have sufficient energy to escape the halo, and that notably less star formation and stellar feedback operates in the group–cluster regime. At these mass scales, after fountaining gas reincorporates into the hot-gas reservoir, it is offset from cooling from radio-mode AGN feedback.

The age bins in the stellar reservoirs allow us also reconstruct the star formation histories of galaxies on a per-component basis. As an example of DARK SAGE’s abilities in this area, we use the summed birth mass of stars in common stellar components across galaxies within wide stellar-mass bins to build cosmic star formation density histories (CSFHs) in the top panels of Figure 12. Because each stellar component has its own metallicity, we can also reconstruct the average metallicity of the star-forming gas as a function of cosmic time for each component. The averages of

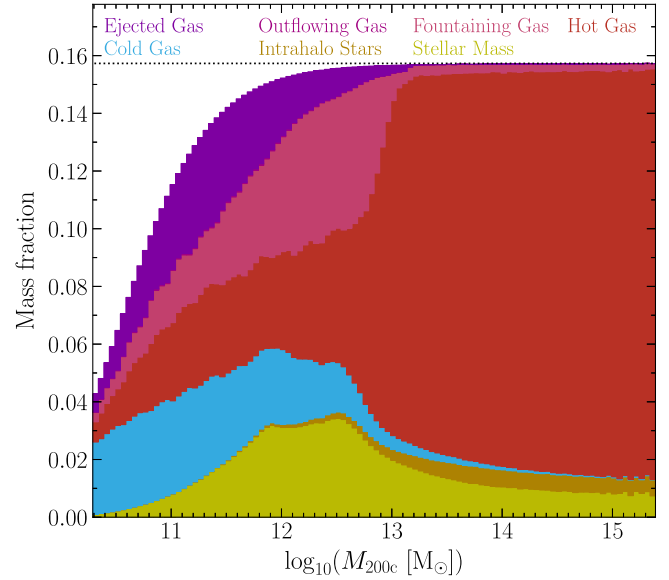


Figure 11. The breakdown of baryons in DARK SAGE haloes at $z = 0$. We calculate the average mass of each reservoir for all haloes in each halo mass bin of width 0.05 dex. Bars are stacked on top of each other, such that the height of all bars in a column, less the ejected gas, gives the total baryon fraction inside the halo. All matter for satellite galaxies inside the virial radius is included. Black holes are reasonably neglected from this plot. The dotted horizontal line represents the cosmic baryon fraction. Photoionization heating suppresses the total baryon fraction below the cosmic value for $M_{200c} \lesssim 10^{13} M_{\odot}$. The ejected gas contributes towards the halo baryon fraction for the purposes of cosmic accretion, even though it is outside the halo (as discussed in Section 4). The outflowing-gas reservoir is high negligible at all halo masses.

these for the binned galaxies are shown in the bottom panels of Figure 12. This plot shows that high-mass galaxies are dominated in stellar mass by merger-driven bulges, while the stellar mass in low-mass galaxies is predominantly in discs. It also shows that spheroidal stars (be they in either bulge or the intrahalo stars; and even the LIGS) have a notably higher average age than disc stars. Except for the highest-mass galaxies, merger-driven bulges tend to be metal-poor relative to discs. There is also a smooth transition in both the age and metallicity of disc stars, getting younger and more metal-rich from the outskirts to the centre and into the instability-driven bulge. These predictions from DARK SAGE can be tested with component-based SED-fitting models applied to large-scale galaxy surveys (e.g. Robotham et al. 2022; also see Bellstedt et al. 2020).

11. Environmental stripping of gas and stars

When a subhalo becomes the satellite of a host halo, several rules for its evolution change. Most notably, baryon reservoirs become prone to environmental effects that can deplete their contents. In DARK SAGE, we now account for the effects of starvation, strangulation, tidal stripping, and ram pressure stripping.

In principle, central haloes can also be affected by the environment of nearby large-scale structure. For example, a halo falling into a cosmic filament from a void should experience tidal and ram pressure effects from that filament. In this section, we describe how environmental effects are handled for both central and satellites galaxies in DARK SAGE.




Figure 12. Top panels: Contributions to the cosmic star formation density history based on the stellar components of galaxies in three wide stellar-mass bins. This is reconstructed from the $z = 0$ output of DARK SAGE using the stellar-age bins in each stellar component after multiplying the mass in those components by $m_{*,\text{pop}}^{\text{form}}/m_{*,\text{pop}}^{\text{rem}}$. Bottom panels: The average metallicity of those stars, equal to the average metallicity of the star-forming gas at the time of formation. The dashed, thin lines are the per-annulus contribution to the disc stars, smoothly changing colour from teal for outer annuli ($i \rightarrow N_{\text{ann}}$) to blue for intermediate annuli, to magenta for inner annuli ($i \rightarrow 1$, i.e. towards the instability-driven bulge, which functions like a ‘zeroth’ annulus).

11.1. Central galaxies

While we have no explicit treatment of large-scale environmental effects on central haloes, this is considered implicitly from how the galaxies are evolved within the merger trees of the N -body simulation. Where an increase in virial mass leads to the addition of new baryons in the halo in order to maintain a target baryon fraction (Section 4), so too are baryons removed when the virial mass decreases to ideally maintain this same fraction. Only the CGM and IHS of a central are stripped in this circumstance. This means that the baryon fraction can become higher than the target value in rare circumstances where removing the CGM and IHS outright is not enough to bring it back down.

When the CGM and IHS are stripped this way, they are done so proportionally to their mass. Those stripped gas and stars are transferred to the central’s LIGM and LIGS, respectively. This means that those baryons will return to the halo preferentially over fresh primordial gas, should the halo later grow (Section 4). Helpfully, this mitigates against potential halo-finding and/or tree-building issues that can lead to artificial, momentary losses in halo mass (Behroozi *et al.* 2013; Srisawat *et al.* 2013; Elahi *et al.* 2019a,b); should this happen in DARK SAGE, baryons will initially be taken out of the halo, but then the *same* baryons (i.e. gas with the same metallicity and stars with the same metallicity and age distribution) will return a snapshot later.

Note that in a change from previous versions of the model, we do not remove the ejected gas of central haloes when they are tidally stripped, as this mass is already considered to be beyond the virial radius.

11.2. Satellite galaxies

Much more detail is given to (and needed for) the environmental effects on satellites. To start with, immediately when a galaxy becomes a satellite, all gas in its LIGM is transferred to the host’s equivalent reservoir. This is simply in line with how we have defined that reservoir; it does not make sense to track a LIGM for a satellite. The same applies for LIGS and LIGBHs.

Satellites are shut off from cosmological gas accretion entirely. This starves the CGM from any replenishment. However, gas can continue to cool from the CGM of the satellite onto its ISM. Likewise, while the satellite remains outside its host virial radius, its CGM will continue to reincorporate gas from the satellite’s ejected reservoir. This is physically justifiable from the results from hydrodynamic simulations that have shown that satellites can still continue to accrete new gas, even while they are being affected by environmental effects that remove gas (Wright *et al.* 2022).

The CGM of a satellite is gradually stripped, an effect often referred to as ‘strangulation’ in the literature. In a change from earlier versions of DARK SAGE, we now also strip the intrahalo stars of a satellite gradually, alongside the CGM (previously, the IHS was instantly transferred to the central upon infall). While this change is inconsequential for galaxy evolution, it allows us to track the stellar haloes of groups after they fall into a cluster. The rate of stripping of the CGM and IHS is based on (as follows) the subhalo mass-loss rate from the N -body simulation, which is caused by gravitational tides. We routinely check the total baryonic mass of the satellite subhalo relative to its expected baryon fraction and total subhalo mass. Should the former be above expectation, we

proportionally remove mass (and metals) from the CGM and IHS to bring the baryonic mass as close to expectation as possible,^o transferring that mass to the respective reservoirs of the central. The decision to strip the CGM and IHS of satellites this way is to be consistent with how those same reservoirs are stripped for centrals when their halo mass drops (Section 11.1). This philosophy extends from Croton et al. (2016).

In truth, the above paragraph is only an accurate description of a satellite that has crossed the virial radius of its host. Until that point, we preference removal of the ejected reservoir of the satellite over the CGM. Any stripped gas and stars of satellites outside R_{200c}^{host} actually go to the respective local intergalactic reservoir. Only if the ejected gas mass of the satellite is zero does DARK SAGE start to strip the CGM. This is expedited by the choice that once the satellite crosses R_{200c}^{host} , its ejected reservoir is transferred to the central's CGM immediately and in its entirety. The assumption is that the gravitational influence of the host halo is sufficiently larger than the satellite, such that any ejected gas can no longer be reincorporated into the satellite's CGM. Because this gas is lost inside the virial radius of the host, there should arguably be no delay in its incorporation to the host's CGM, hence why we do *not* put it into the central's ejected reservoir.

Material from satellite *galaxies* themselves (cold gas and stars in the disc) is also stripped, but follows a different procedure. For this, we calculate the strength of both tides and ram pressure on the satellite galaxy. Previous versions of DARK SAGE only accounted for ram pressure, which only affects the gas. However, we only turn these processes on once $m_* + m_{cold} > m_{hot}$. The reason for this is that the hot-gas halo should initially shield the cold-gas disc from ram pressure, as should the CGM and IHS experience tidal stripping before the baryons in the galaxy disc that are more tightly gravitationally bound. Our exact criterion for switching this shielding off is not theoretically derived; it is simply maintained from Stevens & Brown (2017). Other criteria for the same phenomenon have been used in other SAMs (e.g. Cora et al. 2018).

We approximate the tidal radius of a satellite galaxy based on the formula of Tormen, Diaferio & Syer (1998, cf. their equation 6):

$$r_{tidal} = R_{sat} M_{sat}^{1/3} \Upsilon^{-1/3}, \tag{116a}$$

$$\Upsilon \equiv \begin{cases} 2 M_{halo}(<R_{sat}) \\ -R_{sat} \left. \frac{dM_{halo}(<R)}{dR} \right|_{R=R_{sat}}, & R_{sat} < R_{200c}^{host} \\ 2 M_{200c}^{host}, & R_{sat} \geq R_{200c}^{host} \end{cases} \tag{116b}$$

Strictly, this is only accurate for satellites with circular orbits where $r_{tidal} \ll R_{sat}$ and $M_{sat} \ll M_{halo}(<R_{sat})$ (and therefore $M_{sat} \ll M_{200c}^{host}$).^p However, we only enforce the lattermost criterion for tidal stripping to take place, requiring that $M_{sat} < 0.1 M_{halo}(<R_{sat})$ (and by extension that $M_{sat} < 0.1 M_{200c}^{host}$). With that criterion passed, we remove all gas and stars in each disc annulus where $r_{outer}^{(i)} \geq r_{tidal}$. The stripped gas and stars are respectively transported to the CGM

and IHS of the central if $R_{sat} \leq R_{200c}^{host}$. If the satellite is instead outside the virial radius, the stripped gas and stars go to the respective local intergalactic reservoirs. In the instance that $r_{tidal} < r_{outer}^{(i)}$, we trigger a tidal disruption event where *all* the satellite's baryons are transferred to the CGM and IHS, and the satellite's existence is no longer considered thereafter. As we explain in Section 12, this is not the only means by which a disruption event is triggered. With the exception of tidal disruption events, the stellar bulge of a satellite galaxy is left unaffected by environmental effects in DARK SAGE.

With tidal effects resolved, any remaining gas in the ISM of a satellite is then subject to ram pressure stripping. Consistent with previous versions of DARK SAGE, for each annulus, we calculate the ram pressure exerted on a satellite galaxy as

$$P_{ram} = \rho_{hot}^{cen}(R_{sat}) V_{sat}^2 \tag{117}$$

and compare this against the gravitational restoring force per unit area for each annulus, approximated as

$$P_{restore}^{(i)} = \begin{cases} 2\pi G \Sigma_{cold}^{(i)} (\Sigma_{cold}^{(i)} + \Sigma_*^{(i)}), & \theta_{disc} \leq 10^\circ \\ 2\pi G \Sigma_{cold}^{(i)2}, & \theta_{disc} > 10^\circ \end{cases} \tag{118}$$

(cf. Gunn & Gott 1972). For each annulus where $P_{ram} > P_{restore}^{(i)}$, its gas is transferred to the central's CGM (or LIGM if it is beyond R_{200c}^{host}).

12. Galaxy mergers

Similar to Croton et al. (2016), a galaxy merger or disruption event occurs in DARK SAGE when the baryon fraction of a subhalo rises above a threshold value, f_{bary}^{thresh} . In previous versions of SAGE and DARK SAGE, the subhalo baryon fraction only considered the cold-gas and stellar mass of the satellite galaxy it housed. We now further include the CGM, IHS, and BH masses of the satellite in this calculation. Where previously a threshold fraction of 1.0 was used, we now impose a threshold fraction of 0.5. Recall that the subhalo mass is the total sum of all particles associated with that object from the underlying N -body simulation. Further recall that that mass represents the sum of *both* dark matter and baryons, but that all matter in the N -body simulation *behaves* like dark matter. As a subhalo is tidally stripped, both its dark-matter and hot-gas mass decline. But the stellar and cold-gas content of a satellite typically take longer to be reduced by environmental effects (Section 11). This means the baryon fraction of a subhalo tends to get larger the longer it is a satellite. In principle, it could continue to rise to arbitrarily large numbers greater than 1 — for example, when a subhalo composed of thousands of particles at infall is continually stripped down to the minimum of 20 particles — unless we impose an extra limit. Our physical argument for triggering a merger or satellite disruption event once a subhalo's baryon fraction rises to 0.5 is that the implicit assumption in the N -body simulation — that the satellite galaxy behaves like a dark-matter-dominated object — no longer holds. It presumably must have reached the point where either the un-modelled influence of baryons on the satellite's orbit should have caused it to merge with the central galaxy *or* the stellar and cold-gas content of the galaxy should have been tidally stripped.

The above means that satellite galaxies are generally merged or disrupted *before* their subhalo is 'lost' in the merger trees of the N -body simulation (their being lost in the merger trees triggers

^oWe say 'as close to expectation as possible' rather than definitively saying it will be equal to expectation because it is possible to remove all the CGM and IHS of a satellite and for that to not be enough to restore the expected baryon fraction.

^pSpringel et al. (2008, see their figure 15) show that, on average, equation (116) reproduces the edge of subhaloes in simulations as identified by SUBFIND. This supports the use of this equation as an approximate tidal radius, even when orbits are not circular.

a merger or disruption in the same fashion, but these occasions are rare). This also means there are no ‘orphan’ galaxies in DARK SAGE.

Arguably, it would be better to include orphan galaxies (i.e. detach satellite galaxies from their host subhalo and let them continue to evolve until different criteria are met for merging or disruption), but this requires changes to DARK SAGE’s framework beyond the scope of this version of the model. In practice, this means that the intrahalo and local intergalactic stellar and gas components of DARK SAGE centrals include the content that other models might have otherwise associated with orphan galaxies.

12.1. Merger or disruption event?

In the above situation, how do we then choose between merging the satellite or disrupting it? The answer lies in how long the subhalo has been a satellite. The moment a subhalo becomes a satellite, a predicted merger timescale is calculated. If less than this time has elapsed when the above is triggered, the satellite presumably must have been disrupted. Otherwise, it should have reasonably merged already or be in the process of doing so imminently. In Croton *et al.* (2016), this predicted timescale was based on the dynamical friction model of Binney & Tremaine (1987). However, Poulton *et al.* (2020) has since shown that this timescale — in addition to several other merger timescale formulae commonly used in the literature — only works for a limited range of infall scenarios, and even then carries a factor of ~ 2 uncertainty. Building on this, Poulton *et al.* (2021) built a more sophisticated model by fitting a function to timescales derived from the meticulous tracking of subhalo orbits in cosmological simulations with the ORBWEAVER code. We now adopt this timescale in DARK SAGE.

The Poulton *et al.* (2021) timescale, hereafter t_{merge} , can be calculated as follows (cf. their equations 4–7, 11–12).

$$\frac{t_{\text{merge}}}{5.5 R_{\text{peri}}^{0.2}} = \begin{cases} \sqrt{\frac{R_{200c}^{\text{host}} R_{\text{sat}}^{1.6}}{G M_{\text{halo}}(<R_{\text{sat}})}}, & R_{\text{sat}} < R_{200c}^{\text{host}} \\ \frac{R_{200c}^{\text{host}} R_{\text{sat}}^{0.3}}{\sqrt{G M_{200c}^{\text{host}}}}, & R_{\text{sat}} \geq R_{200c}^{\text{host}} \end{cases}, \quad (119a)$$

$$R_{\text{peri}} = \frac{J_{\text{orb}}^2}{(1 + \varepsilon) G M_{\text{host}} M_{\text{sat}} M_{\text{redu}}}, \quad (119b)$$

$$\varepsilon = \sqrt{1 + \frac{2 E_{\text{orb}} J_{\text{orb}}^2}{(G M_{\text{host}} M_{\text{sat}})^2 M_{\text{redu}}}}, \quad (119c)$$

$$J_{\text{orb}} = M_{\text{redu}} |\vec{R}_{\text{sat}} \times \vec{V}_{\text{sat}}|, \quad (119d)$$

$$E_{\text{orb}} = \frac{M_{\text{redu}} V_{\text{sat}}^2}{2} + M_{\text{sat}} \Phi(R_{\text{sat}}), \quad (119e)$$

$$M_{\text{redu}} \equiv \frac{M_{\text{sat}} M_{\text{host}}}{M_{\text{sat}} + M_{\text{host}}}, \quad (119f)$$

$$M_{\text{host}} \equiv \begin{cases} M_{\text{halo}}(<R_{\text{sat}}), & R_{\text{sat}} < R_{200c}^{\text{host}} \\ M_{200c}^{\text{host}}, & R_{\text{sat}} \geq R_{200c}^{\text{host}} \end{cases}. \quad (119g)$$

Every term in the above equations is known/calculable for each subhalo in DARK SAGE, making this straightforward to implement. R_{sat} represents the distance between the central galaxy’s centre of mass and the satellite galaxy’s centre of mass. R_{peri} represents the expected separation of the galaxies at the pericentre of their orbit with total angular momentum J_{orb} , energy E_{orb} , and ellipticity ε . V_{sat} is the relative speed between the two galaxies. M_{redu} is the ‘reduced’ mass of the two-galaxy system. The mass terms for the central and satellite — M_{host} and M_{sat} , respectively — account for all matter, including dark matter.

When a disruption event occurs, if the satellite galaxy is inside the host’s virial radius, the remaining entirety of the satellite’s CGM, ISM, and ejected gas are transferred to the central’s hot-gas reservoir. Similarly, the entirety of the satellite’s stellar mass (including its IHS) is transferred to the central’s IHS. The size of the central’s IHS is then updated based on where the satellite was disrupted:

$$\langle R \rangle_{\text{IHS}}^{\text{cen}}(t + \Delta t) = \frac{m_{\text{IHS}}^{\text{cen}}(t) \langle R \rangle_{\text{IHS}}^{\text{cen}}(t) + (m_*^{\text{sat}} + m_{\text{IHS}}^{\text{sat}}) R_{\text{sat}}}{m_{\text{IHS}}^{\text{cen}}(t + \Delta t)}, \quad (120)$$

where the average size of the IHS relates to its scale radius, a_{IHS} , of a truncated Hernquist (1990) sphere via

$$\frac{\langle R \rangle_{\text{IHS}}}{a_{\text{IHS}}} = 2 \left(\frac{a_{\text{IHS}}}{R_{200c}} + 1 \right)^2 \ln \left(\frac{R_{200c}}{a_{\text{IHS}}} + 1 \right) - 2 \frac{a_{\text{IHS}}}{R_{200c}} - 3. \quad (121)$$

This equation can be solved for a_{IHS} iteratively. If the disrupted satellite was outside the virial radius, the gas and stars of that satellite are instead transferred to the LIGM and LIGS components of the halo, respectively.

The resolution of a merger is more complicated. How the remaining baryons are handled depends on whether the merger is ‘major’ or ‘minor’.

A minor merger represents when one galaxy vastly outweighs the other. Conversely, a merger is deemed major when the merging systems have comparable mass. In practice, a major merger is defined in DARK SAGE when

$$\min [m_*^{\text{cen}} + m_{\text{cold}}^{\text{cen}}, m_*^{\text{sat}} + m_{\text{cold}}^{\text{sat}}] \geq f_{\text{major}} \times \max [m_*^{\text{cen}} + m_{\text{cold}}^{\text{cen}}, m_*^{\text{sat}} + m_{\text{cold}}^{\text{sat}}] \quad (122)$$

(while the central always has a greater *total* mass than its satellites, it is not strictly guaranteed that it has greater galactic baryonic mass). The treatment of minor mergers remains similar to Stevens *et al.* (2016), but the handling of the ISM in major mergers has notably changed. To more accurately reflect the implicit assumption in the handling of minor mergers that $M_{\text{host}} \gg M_{\text{sat}}$, we set $f_{\text{major}} = 0.1$. Previous versions of the model had this set to 0.3.

12.2. Resolving minor mergers

The ISM of a satellite is assumed to be dispersed uniformly across a subset of the annuli of the central’s disc. The orbital specific angular momentum of the satellite is projected onto the central disc’s rotation axis to find the annulus with equivalent j . Rather than dumping the entire ISM of the satellite into this one annulus of the central, we deposit it across several annuli. This is intended to account for the fact there is actually a *distribution* of orbital j of the satellite’s ISM relative to the central, given that the satellite’s ISM has its own rotation prior to merging. We approximate this by selecting the maximum number of annuli ($l - k + 1 \geq 1$) where

$$\vec{j}_{\text{outer}}^{(k)\text{cen}} > \left(\vec{j}_{\text{orb}} \cdot \hat{j}_{\text{cold}}^{\text{cen}} \right) - j_{\text{spread}}, \quad (123a)$$

$$\vec{j}_{\text{inner}}^{(l)\text{cen}} < \left(\vec{j}_{\text{orb}} \cdot \hat{j}_{\text{cold}}^{\text{cen}} \right) + j_{\text{spread}}, \quad (123b)$$

$$\vec{j}_{\text{orb}} \equiv \vec{R}_{\text{sat}} \times \vec{V}_{\text{sat}}, \quad (123c)$$

$$j_{\text{spread}} \equiv \left(\vec{R}_{\text{sat}} \cdot \hat{j}_{\text{cold}}^{\text{cen}} \right) \left(\hat{j}_{\text{cold}}^{\text{sat}} \cdot \hat{j}_{\text{cold}}^{\text{cen}} \right) V_{200\text{c, infall}}^{\text{sat}} \quad (123d)$$

For each annulus i of the central’s gas disc where $k \leq i \leq l$, we deposit $m_{\text{cold}}^{\text{sat}}/(l - k + 1)$ into that annulus. Metals are deposited in proportion to the satellite’s average ISM metallicity. This is identical to the procedure from Stevens et al. (2016).

In reality, the manner in which we calculate j_{spread} or how we distribute the satellite’s ISM across annuli k through l in the central are of little consequence. The sudden addition of mass to those annuli typically drives a local instability, which promptly redistributes some of that mass over a larger number of annuli (and also removes some of that gas from those annuli due to feedback from instability-driven star formation — see Section 7.2).

The above takes place regardless of whether the orbit of the satellite was prograde or retrograde relative to the central’s gas disc. However, if the orbit was retrograde, this would artificially add angular momentum to the descendant galaxy’s disc, when instead there should be a net loss. To rectify this violation of angular momentum conservation, after a retrograde minor merger, we effectively shrink the descendant’s gas disc by proportionally transferring matter inward, such that the resulting $j_{\text{cold}}^{\text{cen}}$ is as it ought to be, that is,

$$\vec{j}_{\text{cold}}^{\text{cen}}(t + \Delta t) = \left[m_{\text{cold}}^{\text{cen}}(t) + m_{\text{cold}}^{\text{sat}} \right]^{-1} \left[m_{\text{cold}}^{\text{cen}}(t) \vec{j}_{\text{cold}}^{\text{cen}}(t) + m_{\text{cold}}^{\text{sat}} \hat{j}_{\text{cold}}^{\text{cen}} \left(\vec{j}_{\text{orb}} \cdot \hat{j}_{\text{cold}}^{\text{cen}} \right) \right]. \quad (124)$$

This equation is imposed before determining whether any local instabilities are triggered.

We maintain the choice that all stars in the satellite galaxy are transferred to the merger-driven bulge of the central from previous versions of the model. While we would expect collisional gas to settle into the disc structure of the ISM, it would be contrived to force the orbits of collisionless stars acquired from a minor merger to match the stellar disc. The merger-driven bulge exists in the model precisely to account for stars whose orbits are less ordered (there is no analogue for gas). While this principle has not changed, we have added several details in how the merger-driven bulge is updated after a minor merger. For example, we now also transfer any remaining IHS of the satellite to the central’s merger-driven bulge too. That is,

$$m_{\text{m-bulge}}^{\text{cen}}(t + \Delta t) = m_{\text{m-bulge}}^{\text{cen}}(t) + m_{*}^{\text{sat}} + m_{\text{IHS}}^{\text{sat}}. \quad (125)$$

In all likelihood, most of the IHS of the satellite would have already been stripped (per Section 11) and any remaining IHS should be relatively nuclear (and therefore small compared to the size of the central’s bulge).

As another new feature in DARK SAGE, we now approximately track the angular momentum and velocity dispersion of the merger-driven bulge, which in turn is used to calculate its size. After a minor merger, we update the angular momentum of the central by summing the rotational angular momentum of what

the bulge originally had with the total rotational angular momentum of the satellite. That is, we conserve bulge rotation, but neglect any transfer of orbital angular momentum. The assumption is any orbital angular momentum is lost to dynamic friction before stars are added to the bulge. The specific angular momentum of the merger-driven bulge after a minor merger is then:

$$\vec{j}_{\text{m-bulge}}^{\text{cen}}(t + \Delta t) = m_{\text{m-bulge}}^{\text{cen} - 1}(t + \Delta t) \left[m_{\text{m-bulge}}^{\text{cen}}(t) \vec{j}_{\text{m-bulge}}^{\text{cen}}(t) + m_{*, \text{disc}}^{\text{sat}} \vec{j}_{*, \text{disc}}^{\text{sat}} + m_{\text{m-bulge}}^{\text{sat}} \vec{j}_{\text{m-bulge}}^{\text{sat}} \right]. \quad (126)$$

For the merger-driven bulge’s velocity dispersion, we add the mass-weighted squares of the dispersion of each relevant component and a further term that accounts for the radial motion of the satellite relative to the central. That is,

$$\sigma_{\text{m-bulge}}^{\text{cen} 2}(t + \Delta t) = m_{\text{m-bulge}}^{\text{cen} - 1}(t + \Delta t) \left[m_{\text{m-bulge}}^{\text{cen}}(t) \sigma_{\text{m-bulge}}^{\text{cen} 2}(t) + m_{\text{m-bulge}}^{\text{sat}} \sigma_{\text{m-bulge}}^{\text{sat} 2} + m_{i\text{-bulge}}^{\text{sat}} \sigma_{i\text{-bulge}}^{\text{sat} 2} + (m_{*}^{\text{sat}} + m_{\text{IHS}}^{\text{sat}}) \times \left(\vec{V}_{\text{sat}} \cdot \hat{R}_{\text{sat}} \right)^2 + \sum_{i=1}^{N_{\text{ann}}} m_{*, \text{disc}}^{(i)\text{sat}} \sigma_{*, \text{disc}}^{(i)\text{sat} 2} \right]. \quad (127)$$

The CGM and ejected reservoirs of the satellite are added to the hot-gas reservoir of the central. In essence, we assume there is no significant delay for that gas to become available for cooling. There is no impact on whether the satellite was inside or outside the central’s virial radius at the time the merger was triggered.

12.3. Resolving major mergers

During a major merger, the structure of the remnant galaxy is changed significantly from either of its progenitors. First, we define the plane for the new gas disc to exist, chosen to be that of the last orbital plane between the two galaxies. We then calculate the projected specific angular momentum component of each gas disc annulus in each galaxy *relative to the centre of mass of the two-galaxy system* in the direction of this plane. The annulus with the corresponding j in the new disc is where the gas from the old annulus is placed. Note the difference from earlier versions of the model in our choosing the centre-of-mass frame to determine how the new gas disc is built. Not only will the two progenitor discs be mixed, but the surface density profile of the new disc is likely to look entirely different to both progenitors. Similar to the case for minor mergers, if any retrograde gas is added to the new disc, we shrink it to ensure angular momentum conservation is satisfied in the plane. In earlier versions of the model, we triggered a specific starburst when a retrograde merger occurred. We no longer do this, as this is effectively taken care of by the instability prescription.

All stars in the discs and bulges of both progenitors are moved to the merger-driven bulge of the descendant after a major merger, while the two IHS components of the progenitors are summed to form the IHS of the descendant. We set the merger-driven bulge to have a specific angular momentum equal to the orbital specific angular momentum of the two-progenitor system in the centre-of-mass frame. The velocity dispersion of the merger-driven bulge is updated consistently with how it is during minor mergers (where terms for the central that are only present for the satellite in

equation 127 are added). The size of the merger-driven bulge is calculated by solving the Jeans equation, in the same manner as for the instability-driven bulge, except this time accounting for the non-zero angular momentum:

$$a_{\text{m-bulge}} = 3 \left[\frac{1}{\sigma_{\text{m-bulge}}^2} \left(\frac{d\Phi}{dR} - \frac{j_{\text{m-bulge}}^2}{R^3} \right) - \frac{1}{R} \right]^{-1} - R. \quad (128)$$

The CGM and the ejected reservoir of the satellite are instead transferred to the fountain reservoir of the central, and the fountain time is updated when a major merger occurs. The idea here is that the collision of the two comparable mass CGMs should drive turbulence in the gas, causing a delay for it become available for cooling onto the remnant galaxy.

13. Calibration of the model

Our new version of DARK SAGE stands out from most — possibly *all* — other SAMs of galaxy formation in its small number of free parameters. We only allow three parameters to be calibrated in the model:

- $f_{\text{move}}^{\text{gas}}$, the fraction of Toomre-unstable gas that moves to adjacent annuli, as opposed to being consumed in a starburst with feedback (Section 7). This strongly influences star formation rates, the structure and morphology of galaxies, and the rate at which BHs grow. To our knowledge, no other SAM has a parameter that is directly analogous to $f_{\text{move}}^{\text{gas}}$.
- \mathcal{E}_{SN} , the average energy per supernova that is coupled to the ISM/CGM (Section 8). This sets the normalisation for the strength of stellar feedback, which other models usually rely on additional free parameters for.
- ϵ_{AGN} , the radiative efficiency of BHs, which controls the strength of AGN feedback (Section 9). Again, most other models have more than one parameter to control the overall strength of AGN feedback.

Our three free parameters are calibrated against three observables:

- the $z=0$ stellar mass function (SMF) from Driver *et al.* (2022);⁴
- the $z=0$ H I mass function (HIMF) from Jones *et al.* (2018);
- and the CSFH from D’Silva *et al.* (2023).

The best-fitting values for our parameters, following the method we describe below, are provided in Table 1.

We invoke a particle-swarm optimisation (PSO) method to calibrate DARK SAGE. PSO was introduced by Kennedy & Eberhart (1995) and first adopted to calibrate a SAM by Ruiz *et al.* (2015). Using PSO, we find the parameter set that provides the minimum multiplicative reduced χ^2 across the three observables. We deliberately multiply these together instead of adding, as this circumvents issues of the quoted errors carrying different meaning across the constraints, and it ensures that an improvement in one reduced χ^2 of a given *factor* is treated the same as any other

Table 1. Search range of DARK SAGE’s free parameters for the PSO calibration and the output best-fitting values to two significant figures. The latter are used in all results throughout this paper. There is no prior expectation on the value of $f_{\text{move}}^{\text{gas}}$; as it is defined to have a value in (0,1), its search range conservatively covers almost all possible values. \mathcal{E}_{SN} is expected to have a value close to 10^{44} J, based on literature canon, and the search range around this value is deliberately conservatively wide. The search range of ϵ_{AGN} matches the maximum possible range for a black-hole accretion model where mass moves quasi-statically through a centripetal disc until reaching the innermost stable circular orbit, given the range of possible black-hole spin values.

Parameter	Search range	Best fit
$f_{\text{move}}^{\text{gas}}$	0.010–0.999	0.82
$\mathcal{E}_{\text{SN}}/(10^{44} \text{ J})$	0.050–3.000	0.81
ϵ_{AGN}	0.0378–0.4226	0.31

constraint. That is, we minimise

$$\chi_{\text{joint}}^2 \equiv \prod_{\text{var}} \left(\frac{1}{N_{p,\text{var}} - 3} \sum_{p=1}^{N_{p,\text{var}}} \left[\frac{y_{p,\text{var}}^{\text{obs}} - y_{p,\text{var}}^{\text{mod}}}{\sigma_{p,\text{var}}^{\text{obs}}} \right]^2 \right), \quad (129)$$

where var = SMF, HIMF, CSFH. Each y (evaluation) and σ (uncertainty in y) term for point p is calculated in log-space. We match bin sizes in DARK SAGE to the observational data where possible (SMF, HIMF), and interpolate DARK SAGE’s inherent bins otherwise (CSFH), to ensure the fairest comparison of every point. Note that we do not include an uncertainty for DARK SAGE points in equation (129). In principle, we could have applied a Poisson statistic here (as done by Ruiz *et al.* 2015), but this would lead to down-weighting points near and after the knee of the SMF and HIMF, which are arguably the most important — and, in practice, the hardest — parts of the mass functions to get right. Counting uncertainties are also already incorporated into the observational uncertainties.⁵

The PSO code makes use of the PYSWARM package for PYTHON and acts as a wrapper around DARK SAGE. PYSWARM implements a PSO technique similar to that discussed by Ruiz *et al.* (2015). A core difference, however, is that the initial particle positions are randomly drawn from a uniform distribution corresponding to the search space boundary conditions (central column of Table 1), rather than a Latin Hypercube configuration. The PSO code we use was originally developed for the SHARK SAM, with a full description due to be published elsewhere. Minimal edits were made from the PSO codebase in the SHARK repository⁶ to allow the code to interface with DARK SAGE.

Once the initial conditions are set, the particle positions at the subsequent iteration are computed according to

$$\vec{p}_m^{[k+1]} = \vec{p}_m^{[k]} + \vec{v}_m^{[k]}, \quad (130a)$$

$$\vec{v}_m^{[k]} = \frac{1}{2} \vec{v}_m^{[k-1]} + \xi_1 (\vec{p}_m^{[b]} - \vec{p}_m^{[k]}) + \xi_2 (\vec{p}_q^{[l]} - \vec{p}_m^{[k]}). \quad (130b)$$

$\vec{p}_m^{[k]}$ represents the m th particle’s position in the ($f_{\text{move}}^{\text{gas}}$, ϵ_{AGN} , \mathcal{E}_{SN}) parameter space at iteration k . $\vec{v}_m^{[k]}$ is the particle’s corresponding ‘velocity’. $\vec{p}_m^{[b]}$ represents particle m ’s position at the iteration of best fit (lowest χ_{joint}^2) across all previous iterations ($b \leq k$). $\vec{p}_q^{[l]}$

⁴The bins widths used here are half that published in Driver *et al.* (2022). This better captures the shape of the knee of the SMF. These data were provided privately by S. P. Driver. Note as well that these data are at an average redshift of $z \simeq 0.072$, but we have calibrated to our $z=0$ output.

⁵The code was adapted from here: <https://github.com/ICRAR/shark/tree/1dc89e8ad8d76490f783c68a91798430605b6f70/optim>. Note that the SHARK PSO codebase has since been updated.

represents the position for the global best fit for *all* particles across all previous iterations (that of particle q at iteration l). ξ_1 and ξ_2 are random numbers, each drawn from a uniform distribution in $[0, 0.5]$. The initial velocity of each particle, $\vec{v}_m^{[0]}$, is set to a random fraction (< 1) of the length of the parameter space in each dimension, and also in a random direction. Should the velocity of a particle move it to outside the search range, it is manually shifted to the edge of the search range.

We must also specify the range in the observational data to which the model is fitted. We use the simulation’s resolution to inform this decision. Our loose criterion is that if galaxies of a given stellar or H I mass exclusively occupy (sub)haloes that have had at least 200 particles for one or more snapshots throughout their primary progenitor branch in the merger trees, then we are ‘complete’ at that mass, and can fit observations from there upwards. 200 particles is a semi-arbitrary number to require completeness to, but in the context of MTNG haloes having a minimum of 20 particles, our criterion strikes a balance between using as much information (as many haloes) as possible without giving too much trust towards MTNG’s resolution limit. The stellar or H I mass to which we are complete to 200 particles depends on the parameters of the model. After exploring the parameter space, we found we could comfortably rely on being complete for $m_* \geq 10^9 M_\odot$ and $m_{\text{H I}} \geq 10^9 M_\odot$ (erring on the side of caution, especially in the case of the stellar mass). These are the limits we adopt for the SMF and HIMF when calculating χ^2_{joint} . In Figure 13, we show stellar and H I mass distributions of DARK SAGE galaxies as a function of their historical-maximum (sub)halo mass at $z = 0$,[†] and the 200-particle completeness as a function of each mass.

In total, the PSO code went through two rounds of 16 average[‡] iterations with 16 particles, meaning 512 parameter sets of DARK SAGE were fully simulated with mini-MTNG. Each instance of DARK SAGE on mini-MTNG required ~ 2.2 CPUh of walltime. It was important to do more than one run of the PSO code from start to finish to mitigate the possibility of finding a local minimum instead of the global minimum within the defined parameter space. Figure 14 shows the output of each run in the PSO chains against the calibration data, the relative goodness of fit of each run, and the particle positions in the model’s parameter space. We rounded the best-fitting parameters to two significant figures for the final run of the model on MTNG.

We note that DARK SAGE can reproduce any of these three observables to higher accuracy if constrained exclusively by that dataset. In principle, we could also have weighted the constraints in the joint χ^2 calculation had we felt one was more important to reproduce. But in the absence of an obvious reason why one constraint should have greater inherent value over another, we defaulted to having equal weights. Evidently, there is some tension in the model with respect to simultaneously reproducing the SMF and HIMF. This could reflect (i) the simplicity of the molecular-fraction calculation in the model (Section 6.4); (ii) a deeper limitation of the model’s design; (iii) uncertainties in the observational data, bearing in mind that no forward-modelling of any observational uncertainties has been done here; or a combination thereof. We leave a deeper assessment of this for future studies.

[†]Although similar, Figure 13 is *not* the same as a nominal stellar–halo or H I–halo mass relation. Those, instead, tend to have M_{200c} on the x-axis and do not consider subhaloes as separate objects.

[‡]One round went through 15 iterations, the other 17. Both had converged.




Figure 13. Historical maximum of the FoF or subhalo mass for DARK SAGE galaxies versus their stellar mass (top panel) or H I mass (bottom panel) at $z = 0$. We define a galaxy as having 200 times the particle mass of MTNG on the x-axis as being resolved. Calibration to the SMF and HIMF only occurs (approximately) above the respective masses where DARK SAGE is complete to this mass resolution. The long-dashed, red curves show the completeness as a function of stellar and H I mass. The shaded region on the left is where galaxies are not reliably resolved by this definition. The thin, dashed, horizontal lines are the masses above which we calibrate the SMF and HIMF. As in other figures, hexbins and purple lines show the density of DARK SAGE galaxies and the running 16th, 50th, and 84th percentiles.

To highlight the significance of having so few free parameters, especially in the context of DARK SAGE galaxies having multidimensional structure, we compare the number of free parameters and details of galaxy structure with other popular SAMs in the literature in Table 2. We count 15 free parameters for L-GALAXIES based on the parameters listed that are not fixed in table 1 of Henriques et al. (2020), that is, used in their Monte Carlo Markov Chain calibration procedure; 10 for MERAXES based on table 1 of Mutch et al. (2016), discounting the one parameter listed with a

Table 2. Stand-out features of DARK SAGE compared to its previous version and other semi-analytic models in the literature. Each row represents a different semi-analytic model. Columns represent: (i) the number of free parameters that model has; (ii) the model's treatment of radial disc structure; (iii) treatment of radial metallicity variation in galaxy discs, if any; (iv) treatment of stellar-age distributions as function of disc radius; (v) treatment of mass loading from stellar feedback. Only if an entry is 'non-parametric' can that model make predictions for that property. Note that the total number of nominal parameters a model may have is not what we equate as free parameters. For a parameter to count as 'free' by our definition, it has to have been described as being varied in a calibration procedure in the model paper. The most relevant reference for each model is: DARK SAGE 2018 (Stevens *et al.* 2018), L-GALAXIES (Henriques *et al.* 2020), MERAXES (Mutch *et al.* 2016), SAG (Cora *et al.* 2018), SAGE (Croton *et al.* 2016), SHARK (Lagos *et al.* 2018).

Model (year)	(i) $N_{\text{param}}^{\text{free}}$	(ii) $\Sigma_{\text{disc}}(r)$	(iii) $Z(r)$	(iv) $d\Sigma_*(r)/dt_{\text{form}}$	(v) Mass loading
DARK SAGE (this work)	3	Non-parametric	Non-parametric	Non-parametric	Local, non-parametric
DARK SAGE (2018)	8	Non-parametric	Non-parametric	None	Local, parametric
L-GALAXIES (2020)	15	Non-parametric	Non-parametric	None	Local, parametric
MERAXES (2016)	10	Parametric	Flat	Flat	Global, parametric
SAG (2018)	9	Parametric	Flat	None	Global, parametric
SAGE (2016)	8	Parametric	Flat	None	Global, constant
SHARK (2018)	18	Parametric	Flat	Flat	Global, parametric

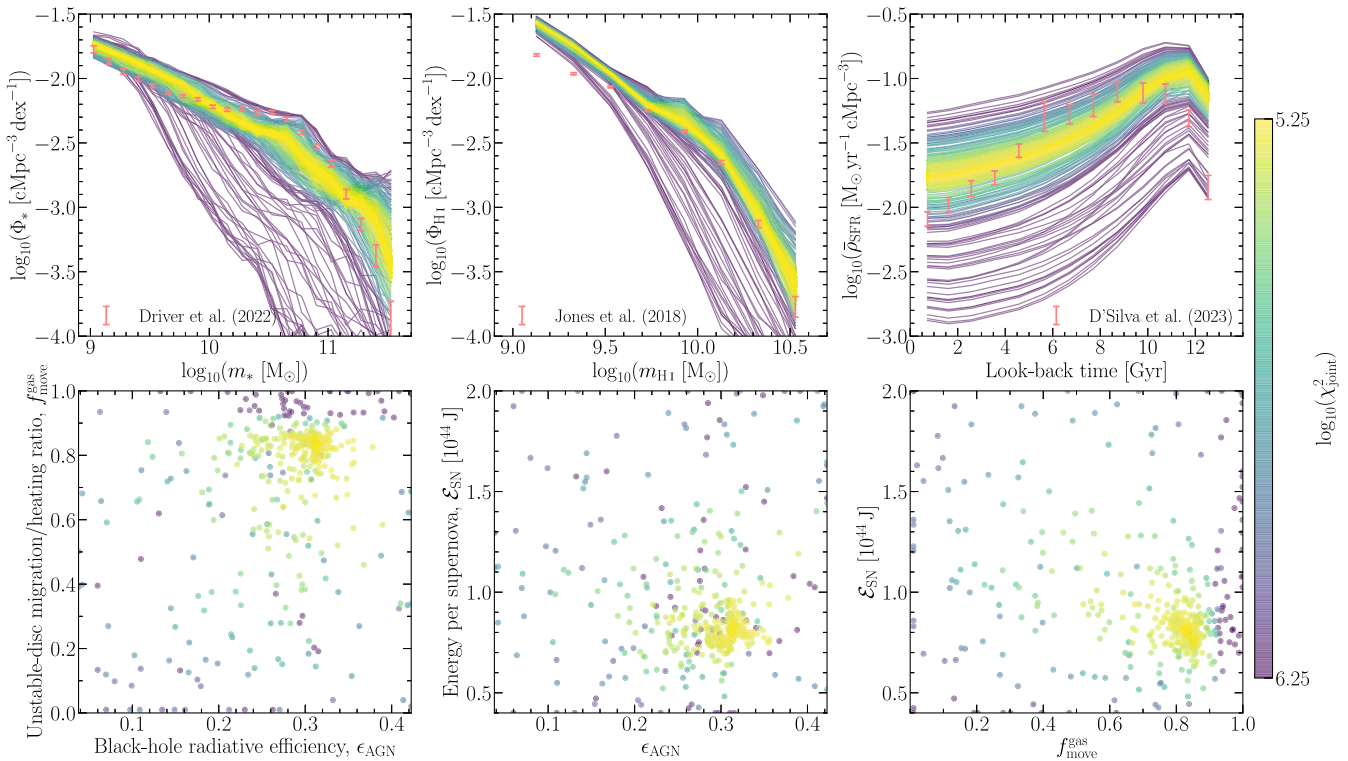


Figure 14. Three observed relations used to constrain the three free parameters in DARK SAGE. *Top left:* stellar mass function at $z = 0$ with observational data at $z \simeq 0.072$ from Driver *et al.* (2022). *Top middle:* H I mass function at $z = 0$ with observational data from Jones *et al.* (2018). *Top right:* cosmic star formation history with observational data from D'Silva *et al.* (2023). Each line in each of the panels corresponds to a run of DARK SAGE with a different $(f_{\text{move}}^{\text{gas}}, \epsilon_{\text{AGN}}, \mathcal{E}_{\text{SN}})$ parameter set. The colour indicates the combined reduced χ^2 of the three constraints (equation 129). The *bottom* panels show where each run lies in DARK SAGE's parameter space. The best fit was found by applying a particle-swarm optimisation code. See Section 13 for further details.

value of zero; 9 for SAG listed in table 1 of Cora *et al.* (2018) in their PSO calibration; 8 for SAGE based on inside knowledge and what is described as explicitly being calibrated/free in Croton *et al.* (2016); and 18 for SHARK after cross-referencing the parameters listed in table 3 of Lagos *et al.* (2018) as explicitly described as free in the text and/or departing from their default value. There are many more SAMs in the literature that we could have added to

Table 2, in principle. However, we found some papers to be opaque in their description of how their model is parametrised, and we therefore felt that we could not make a fair comparison with them. We expect those models fall in the same range as the others in Table 2.

For a deeper dive into the complexity of calibrating (other) SAMs, we refer the reader to Knebe *et al.* (2018).

14. Closing remarks

We have presented an overview of the salient features of the brand-new version of the DARK SAGE SAM of galaxy formation. DARK SAGE bucks the trend of most SAMs by numerically evolving the multidimensional structure of galaxy discs and by only having three free parameters that require calibration against observations.

The predictions from DARK SAGE presented in this paper are but the tip of the iceberg as far as what the model can produce. The code codebase is open to the community, which we hope will serve as a useful resource for research into galaxy formation and cosmology.

Acknowledgements. All plots in this paper were built with the MATPLOTLIB package for PYTHON (Hunter 2007). Our analysis also made extended use of the NUMPY package (Harris et al. 2020) and PYSWARM (<https://github.com/tisimst/pyswarm>).

ARHS acknowledges receipt of the Jim Buckee Fellowship at ICRAR-UWA. Parts of this research were supported by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013.

We used the OzSTAR supercomputer at Swinburne University of Technology to convert the MTNG merger trees and run DARK SAGE for this project. The OzSTAR supercomputing environment is managed through the Centre for Astrophysics & Supercomputing with support from Swinburne IT. The supercomputing program receives continued financial support for operations from Astronomy Australia Limited and the Australian Commonwealth Government through the National Collaborative Research Infrastructure Strategy (NCRIS). The hardware for Ngarrgu Tindebeek was purchased through a grant from the Victorian Higher Education State Investment Fund (VHESIF).

We thank A. Acharyya, S. Bellstedt, J. C. J. D'Silva, S. P. Driver, M. G. Jones, and A. B. Watts for supplying some of the observation-based data/results compared to in our figures. ARHS further thanks C. Power, the SU3 research group at ICRAR-UWA, and the GENESIS research group within ASTRO 3D for their comments in many meetings over the years; R. Tobar for help interfacing the PSO code with DARK SAGE; V. Springel and the wider MillenniumTNG collaboration for green-lighting the inclusion of the simulations in this project.

Finally, we thank the referee for finding several places in the paper where we had accidentally left out information.

Author contributions. ARHS developed the model, led the code development, produced the data, and wrote the paper. MS directly contributed to the code development, assisted in the model design, and interfaced DARK SAGE with MTNG. AR assisted in the development of the stellar feedback model. MWS developed and wrote the code for the initial PSO framework. BH, CHA, and LH co-produced the MTNG simulations and merger trees. All co-authors have read and supplied feedback on this text.

Data Availability Statement. At the time of writing this paper, the MillenniumTNG data are private, but are due to be made publicly available in 2024. We encourage anyone interested in obtaining access to DARK SAGE data sooner to contact MS.

References

Athanassoula, E. 2003, *MNRAS*, 341, 1179
 Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, *ARA&A*, 47, 481
 Bardeen, J. M. 1970, *Nature*, 226, 64
 Bardeen, J. M., Press, W. H., & Teukolsky, S. A. 1972, 178, 347
 Barrera, M., et al. 2023, *MNRAS*, 525, 6312
 Behroozi, P. S., Wechsler, R. H., Wu, H.-Y., Busha, M. T., Klypin, A. A., & Primack, J. R. 2013, *ApJ*, 763, 18
 Belfiore, F., et al. 2017, *MNRAS*, 469, 151

Bellstedt, S., Forbes, D. A., Foster, C., Romanowsky, A. J., Brodie, J. P., Pastorello, N., Alabi, A., & Villaume, A. 2017, *MNRAS*, 467, 4540
 Bellstedt, S., et al. 2020, *MNRAS*, 498, 5581
 Bellstedt, S., et al. 2021, *MNRAS*, 503, 3309
 Benítez-Llambay, A., Navarro, J. F., Frenk, C. S., & Ludlow, A. D. 2018, *MNRAS*, 473, 1019
 Benson, A. J. 2012, *New Astron.*, 17, 175
 Benson, A. J. 2017, *MNRAS*, 471, 2871
 Benson, A. J., Bower, R. G., Frenk, C. S., Lacey, C. G., Baugh, C. M., & Cole, S. 2003, *ApJ*, 599, 38
 Bertschinger, E. 1989, *ApJ*, 340, 666
 Bigiel, F., et al. 2011, *ApJ*, 730, L13
 Binney, J., & Tremaine, S. 1987, *Galactic Dynamics* (Princeton, NJ: Princeton Univ. Press)
 Blitz, L., & Rosolowsky, E. 2004, *ApJ*, 612, L29
 Bondi, H. 1952, *MNRAS*, 112, 195
 Bose, S. et al. 2023, *MNRAS*, 524, 2579
 Bower, R. G., Benson, A. J., Malbon, R., Helly, J. C., Frenk, C. S., Baugh, C. M., Cole, S., & Lacey, C. G. 2006, *MNRAS*, 370, 645
 Bower, R. G., Benson, A. J., & Crain, R. A. 2012, *MNRAS*, 422, 2816
 Brown, T., et al. 2017, *MNRAS*, 466, 1275
 Brown, T., et al. 2021, *ApJS*, 257, 21
 Bullock, J. S., Dekel, A., Kolatt, T. S., Kravtsov, A. V., Klypin, A. A., Porciani, C., & Primack, J. R. 2001, *ApJ*, 555, 240
 Catinella, B., et al. 2018, *MNRAS*, 476, 875
 Chabrier, G. 2003, *PASP*, 115, 763
 Chandrasekhar, S. 1931, *ApJ*, 74, 81
 Chisholm, J., Tremonti, C. A., Leitherer, C., & Chen, Y. 2017, *MNRAS*, 469, 4831
 Chisholm, J., Tremonti, C. A., & Leitherer, C. 2018, *MNRAS*, 481, 1690
 Conroy, C., Gunn, J. E., & White, M. 2009, *ApJ*, 699, 486
 Contreras, S., Padilla, N., & Lagos, C. d. P. 2017, *MNRAS*, 472, 4992
 Contreras, S., et al. 2023, *MNRAS*, 524, 2489
 Cora, S. A. et al. 2018, *MNRAS*, 479, 2
 Crain, R. A., et al., 2015, *MNRAS*, 450, 1937
 Croton, D. J., et al. 2006, *MNRAS*, 365, 11
 Croton, D. J. et al. 2016, *ApJS*, 222, 22
 D'Silva, J. C. J., et al. 2023, *MNRAS*, 524, 1448
 De Lucia, G., Kauffmann, G., & White, S. D. M. 2004, *MNRAS*, 349, 1101
 De Lucia, G., Tornatore, L., Frenk, C. S., Helmi, A., Navarro, J. F., & White, S. D. M. 2014, *MNRAS*, 445, 970
 Dekel, A., & Birnboim, Y. 2006, *MNRAS*, 368, 2
 Delgado, A. M., et al. 2023, *MNRAS*, 523, 5899
 Di Cintio, A., Brook, C. B., Dutton, A. A., Macciò, A. V., Stinson, G. S., & Knebe, A. 2014, *MNRAS*, 441, 2986
 Driver, S. P., et al. 2022, *MNRAS*, 513, 439
 Dutton, A. A., & Macciò, A. V. 2014, *MNRAS*, 441, 3359
 Eddington, A. S. 1926, *The Internal Constitution of the Stars* (Cambridge Univ. Press, Cambridge)
 Efstathiou, G. 1992, *MNRAS*, 256, 43P
 Elahi, P. J., et al. 2019a, *PASA*, 36, 21
 Elahi, P. J., Poulton, R. J. J., Tobar, R. J., Cañas, R., Lagos, C. d. P., Power, C., & Robotham, A. S. G. 2019b, *PASA*, 36, 28
 Elmegreen, B. G. 1989, *ApJ*, 338, 178
 Fanidakis, N., Baugh, C. M., Benson, A. J., Bower, R. G., Cole, S., Done, C., & Frenk, C. S. 2011, *MNRAS*, 410, 53
 Faucher-Giguère, C.-A. 2020, *MNRAS*, 493, 1614
 Ferlito, F., et al. 2023, *MNRAS*, 524, 5591
 Fields, B. D. 2011, *Annu. Rev. Nucl. Part. Sci.*, 61, 47
 Font, A. S., et al., 2008, *MNRAS*, 389, 1619
 Gnedin, N. Y. 2000, *ApJ*, 542, 535
 Gnedin, N. Y., & Draine, B. T. 2014, *ApJ*, 795, 37
 Gunn, J. E., & Gott, J. R. 1972, *ApJ*, 176, 1
 Guo, Q., et al. 2011, *MNRAS*, 413, 101
 Hadzhiyska, B., et al. 2023a, *MNRAS*, 524, 2507
 Hadzhiyska, B., et al. 2023b, *MNRAS*, 524, 2524
 Harris, C. R., et al. 2020, *Nature*, 585, 357

- Harwit, M. 1988, *Astrophysical Concepts*, 2nd Ed. (Springer-Verlag, New York)
- Heckman, T. M., Alexandroff, R. M., Borthakur, S., Overzier, R., & Leitherer, C. 2015, *ApJ*, 809, 147
- Henriques, B. M. B., White, S. D. M., Thomas, P. A., Angulo, R. E., Guo, Q., Lemson, G., & Springel, V. 2013, *MNRAS*, 431, 3373
- Henriques, B. M. B., et al. 2020, *MNRAS*, 491, 5795
- Hernández-Aguayo, C., et al. 2023, *MNRAS*, 524, 2556
- Hernquist, L. 1990, *ApJ*, 356, 359
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, 9, 90
- Jiménez-Donaire, M. J., et al. 2023, *A&A*, 671, A3
- Jones, M. G., Haynes, M. P., Giovanelli, R., & Moorman, C. 2018, *MNRAS*, 477, 2
- Kannan, R., et al. 2023, *MNRAS*, 524, 2594
- Kauffmann, G., & Haehnelt, M. 2000, *MNRAS*, 311, 576
- Kauffmann, G., Colberg, J. M., Diaferio, A., & White, S. D. M. 1999, *MNRAS*, 303, 188
- Kennedy, J., & Eberhart, R. 1995, *IEEE*, 4, 1942
- Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, *MNRAS*, 363, 2
- Kerr, R. P. 1963, *PhRvL*, 11, 237
- Kewley, L. J., & Ellison, S. L. 2008, *ApJ*, 681, 1183
- Kistler, M. D., Stanek, K. Z., Kochanek, C. S., Prieto, J. L., & Thompson, T. A. 2013, *ApJ*, 770, 88
- Knebe, A., & Power, C. 2008, *ApJ*, 678, 621
- Knebe, A., et al. 2018, *MNRAS*, 475, 2936
- Kormendy, J., & Kennicutt, R. C. Jr 2004, *ARA&A*, 42, 603
- Kravtsov, A. V., Gnedin, O. Y., & Klypin, A. A. 2004, *ApJ*, 609, 482
- Lagos, C. d. P., Lacey, C. G., Baugh, C. M., 2013, *MNRAS*, 436, 1787
- Lagos, C. d. P., Tobar, R. J., Robotham, A. S. G., Obreschkow, D., Mitchell, P. D., Power, C., & Elahi, P. J. 2018, *MNRAS*, 481, 3573
- Lelli, F., McGaugh, S. S., & Schombert, J. M. 2016, *AJ*, 152, 157
- Lelli, F., McGaugh, S. S., Schombert, J. M., Desmond, H., & Katz, H. 2019, *MNRAS*, 484, 3267
- Leroy, A. K., et al. 2009, *AJ*, 137, 4670
- Li, J., Liu, C., Zhang, Z.-Y., Tian, H., Fu, X., Li, J., & Yan, Z.-Q. 2023, *Nature*, 613, 460
- Lochhaas, C., et al. 2021, *ApJ*, 922, 121
- Lutz, K. A., et al. 2018, *MNRAS*, 476, 3744
- Lynden-Bell, D. 1969, *Nature*, 223, 690
- Madau, P., & Dickinson, M. 2014, *ARA&A*, 52, 415
- Maggi, P., et al. 2016, *A&A*, 585, 162
- Marinacci, F., et al. 2018, *MNRAS*, 480, 5113
- McGaugh, S. S., Schombert, J. M., Bothun, G. D., de Blok, W. J. G. 2000, *AJ*, 533, L99
- McKee, C. F., & Krumholz, M. 2010, *ApJ*, 709, 308
- Muratov, A. L., Kereš, D., Faucher-Giguère, C.-A., Hopkins, P. F., Quataert, E., & Murray, N. 2015, *MNRAS*, 454, 2691
- Mutch, S. J., Poole, G. B., & Croton, D. J. 2013, *MNRAS*, 428, 2001
- Mutch, S. J., Geil, P. M., Poole, G. B., Angel, P. W., Duffy, A. R., Mesinger, A., & Wyithe, J. S. B. 2016, *MNRAS*, 462, 250
- Naiman, J. P., et al. 2018, *MNRAS*, 477, 1206
- National Academies of Sciences, Engineering, and Medicine 2021. *Pathways to Discovery in Astronomy and Astrophysics for the 2020s* (Washington, DC: The National Academies Press). <https://doi.org/10.17226/26141>.
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1996, *ApJ*, 462, 563
- Nelson, D., et al. 2018, *MNRAS*, 475, 624
- Nelson, D., et al. 2019a, *MNRAS*, 490, 3234
- Nelson, D., et al. 2019b, *ComAC* 6, 2
- Novikov, I. D., & Thorne, K. S. 1973, in *Black Holes*, eds. C. DeWitt & B. S. DeWitt (Paris: Gordon & Breach)
- Nulsen, P. E. J., & Fabian, A. C. 2000, *MNRAS*, 311, 346
- Obreschkow, D., Croton, D., De Lucia, G., Khochfar, S., & Rawlings, S. 2009, *ApJ*, 698, 1467
- Okamura, T., Shimasaku, K., & Kawamata, R. 2018, *ApJ*, 854, 22
- Pakmor, R., et al. 2023, *MNRAS*, 524, 2539
- Pillepich, A., et al. 2018a, *MNRAS*, 473, 4077
- Pillepich, A., et al. 2018b, *MNRAS*, 475, 648
- Pillepich, A., et al. 2019, *MNRAS*, 490, 3196
- Plank Collaboration XIII, 2016, *A&A*, 594, A13
- Poetrodjojo, H., et al. 2021, *MNRAS*, 502, 3357
- Porrás-Valverde, A. J., Holley-Bockelmann, K., Berlind, A. A., & Stevens, A. R. H. 2021, *ApJ*, 923, 273
- Poulton, R. J. J., Power, C., Robotham, A. S. G., & Elahi, P. J. 2020, *MNRAS*, 491, 3820
- Poulton, R. J. J., Power, C., Robotham, A. S. G., Elahi, P. J., & Lagos, C. D. P. 2021, *MNRAS*, 501, 2810
- Pringle, J., & King, A. 2007, *Astrophysical Flows* (Cambridge Univ. Press, Cambridge)
- Qin, F., Howlett, C., Stevens, A. R. H., & Parkinson, D. 2022, *ApJ*, 937, 113
- Rahmati, A., Schaye, J., Pawlik, A. H., & Raičević, M. 2013, *MNRAS*, 431, 2261
- Raouf, M., Shabala, S. S., Croton, D. J., Khosroshahi, H. G., & Bernyk, M. 2017, *MNRAS*, 471, 658
- Robotham, A. S. G., Bellstedt, S., & Driver, S. P. 2022, *MNRAS*, 513, 2985
- Romeo, A. B., & Wiegert, J. 2011, *MNRAS*, 416, 1191
- Ruiz, A. N., et al. 2015, *ApJ*, 801, 139
- Saintonge, A., et al. 2017, *ApJS*, 233, 22
- Sánchez, S. F., et al. 2014, *A&A*, 563, A49
- Schaye, J., et al. 2015, *MNRAS*, 446, 521
- Schwarzschild, K. 1916, *Abh. Konigl. Preuss. Akad. Wissenschaften Jahre*, 1916, 189
- Scott, N., Graham, A. W., & Schombert, J. 2013, *ApJ*, 768, 76
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, 24, 337
- Sharda, P., Krumholz, M. R., Wisnioski, E., Archaryya, A., Federrath, C., & Forbes, J. C. 2021, *MNRAS*, 504, 53
- Smartt, S. J. 2009, *ARA&A*, 47, 63
- Somerville, R. S., & Davé, R. 2015, *ARA&A*, 53, 51
- Somerville, R. S., Primack, J. R., & Faber, S. M. 2001, *MNRAS*, 320, 50
- Springel, V. 2010, *MNRAS*, 401, 791
- Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, *MNRAS*, 328, 726
- Springel, V., et al. 2005, *Nature*, 435, 629
- Springel, V., et al. 2008, *MNRAS*, 391, 1685
- Springel, V., et al. 2018, *MNRAS*, 475, 676
- Springel, V., Pakmor, R., Zier, O., & Reinecke, M. 2021, *MNRAS*, 506, 2871
- Srisawat, C., et al. 2013, *MNRAS*, 436, 150
- Stevens, A. R. H. 2015, *PASA*, 32, 30
- Stevens, A. R. H., & Brown, T. 2017, *MNRAS*, 471, 447
- Stevens, A. R. H., Croton, D. J., & Mutch, S. J. 2016, *MNRAS*, 461, 859
- Stevens, A. R. H., Lagos, C. d. P., Contreras, S., Croton, D. J., Padilla, N. D., Schaller, M., Schaye, J., & Theuns, T. 2017a, *MNRAS*, 467, 2066
- Stevens, A. R. H., Croton, D. J., Mutch, S. J., & Sinha, M. 2017b, *Astrophysics Source Code Library*, record ascl:1706.004
- Stevens, A. R. H., Lagos, C. d. P., Obreschkow, D., & Sinha, M. 2018, *MNRAS*, 481, 5543
- Stevens, A. R. H., et al. 2019a, *MNRAS*, 483, 5334
- Stevens, A. R. H., Diemer, B., Lagos, C. d. P., Nelson, D., Obreschkow, D., Wang, J., & Marinacci, F. 2019b, *MNRAS*, 490, 96
- Stevens, A. R. H., et al. 2021, *MNRAS*, 502, 3158
- Sugahara, Y., et al. 2017, *ApJ*, 850, 51
- Sutherland, R. S., & Dopita, M. A. 1993, *ApJS*, 88, 253
- Thorne, K. S. 1974, *ApJ*, 191, 507
- Toomre, A. 1964, *ApJ*, 139, 1217
- Tormen, G., Diaferio, A., & Syer, D. 1998, *MNRAS*, 299, 728
- Torrey, P., et al. 2019, *MNRAS*, 484, 5587
- Tremonti, C. A., et al. 2004, *ApJ*, 613, 898
- Tsujimoto, T., Nomoto, K., Yoshii, Y., Hashimoto, M., Yanagida, S., & Thielemann, F.-K. 1995, *MNRAS*, 277, 945
- Tully, R. B., Fisher, J. R. 1977, *A&A* 54, 661
- Übler, H., et al. 2019, *ApJ*, 880, 48
- van de Voort, F., Schaye, J., Booth, C. M., Haas, M. R., & Dalla Vecchia, C. 2011, *MNRAS*, 414, 2458
- Vogelsberger, M., Marinacci, F., Torrey, P., & Puchwein, E. 2020, *Nat. Rev. Phys.*, 2, 42

- Watts, A. B., et al. 2023, PASA, 40, 17
Weinberger, R., et al. 2017, MNRAS 465, 3291
Wolz, L., et al. 2022, MNRAS, 510, 3495
White, S. D. M., & Frenk, C. S. 1991, ApJ, 379, 52
White, C. E., Somerville, R. S., & Ferguson, H. C. 2015, ApJ, 799, 201
- Wright, R. J., Lagos, C. d. P., Power, C., Stevens, A. R. H., Cortese, L., & Poulton, R. J. J. 2022, MNRAS, 516, 2891
Xie, L., De Lucia, G., Hirschmann, M., & Fontanot, F. 2020, MNRAS, 498, 4327
Yates, R. M., Henriques, B. M. B., Fu, J., Kauffmann, G., Thomas, P. A., Guo, Q., White, S. D. M., & Schady, P. 2021, MNRAS, 503, 4474